Essays on Economic Growth and Innovation

by

Todd Lensman

B.A. with Distincton, Economics & Mathematics, Cornell University, 2019

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN ECONOMICS

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

Authored by: Todd Lensman
Department of Economics
May 15, 2025

Certified by: Daron Acemoglu
Institute Professor, Thesis Supervisor

Certified by: Michael Whinston
Professor of Economics and Sloan Fellows Professor of Management, Thesis Supervisor

Certified by: Jacob Moscona
3M Career Development Assistant Professor of Economics, Thesis Supervisor

Accepted by: Isaiah Andrews
Charles E. and Susan T. Harris Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays on Economic Growth and Innovation

by

Todd Lensman

Submitted to the Department of Economics
on May 15, 2025 in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN ECONOMICS

ABSTRACT

A foundational observation by Robert Solow holds that long-run economic growth is primarily driven by the innovation and adoption of new technologies (Solow, 1957). This set of essays provides new theory and evidence to explain how firms choose *which* technologies to innovate and adopt. A point of emphasis, particularly in the first two chapters, is that complementarities across firms play an important role in determining the rate and direction of technological change. These complementarities arise as firms build shared knowledge by innovating (Chapter 1) and from joint consumption of new products (Chapter 2). They provide a new channel through which market structure and property rights affect long-run technological change.

Chapter 1. The first chapter is motivated by the observation that the direction of innovation shapes both current technologies and future innovation opportunities, as firms acquire expertise and create public knowledge through discovery. But how do firms choose which technologies to develop? Do they ever fail to exploit new technological paradigms? I build a new model of innovation and firm dynamics to study a novel link between market structure, the direction of innovation, and economic growth: Expertise in a current technology gives incumbents a comparative advantage at innovating it relative to entrants, who instead favor a new technology with higher growth potential. Each firm's innovation decisions influence others through knowledge spillovers, so the initial market structure can affect the long-run direction of innovation. Concentrating R&D resources in a small number of firms allows faster accumulation of expertise, raising growth when all firms innovate the same technology. But it can lower growth when firms face a technology choice, amplifying the influence of incumbents and potentially delaying or preventing the emergence of the new technology. I provide empirical evidence for the theory using data on firm patenting and R&D expenditures. I also show that it explains the historical development of mRNA vaccines, and I explore its implications for the highly concentrated innovation of artificial intelligence.

Chapter 2. In the second chapter, joint with Rebekah Dix, we observe that innovations often combine several components to achieve outcomes greater than the "sum of the parts." We argue that such combination innovations can introduce an understudied inefficiency—a positive market expansion externality that benefits the owners of the components. We demonstrate the importance of this externality in the market for pharmaceutical cancer treatments, where drug

combination therapies have proven highly effective. Using data on clinical trial investments, we document several facts consistent with inefficiently low private innovation: firms are less likely than publicly funded researchers to trial combinations, firms are less likely to trial combinations including other firms' drugs than those including their own drugs, and firms often wait to trial combinations including other firms' drugs until those drugs experience generic entry. Using microdata on drug prices and utilization, we quantify the externalities that arise from new combinations and find that the market expansion externality often dominates the standard negative business stealing externality, suggesting too little innovation in combination therapies. As a result, firms may have incentives to free ride off others' innovation, which we analyze with a dynamic structural model of innovation decisions. We use the model to design cost-effective policies that advance combination innovation. Redirecting publicly funded innovation toward combinations with high predicted market expansion or consumer surplus spillovers minimizes crowd out of private investments, increasing the rate of combination innovation and total welfare while remaining budget neutral.

Chapter 3. The final chapter, joint with Daron Acemoglu, considers incentives to adopt transformative technologies that promise to accelerate productivity growth across many sectors but also present new risks from potential misuse. We develop a multi-sector technology adoption model to study the optimal regulation of transformative technologies when society can learn about these risks over time. Socially optimal adoption is gradual and typically convex. If social damages are large and proportional to the new technology's productivity, a higher growth rate paradoxically leads to slower optimal adoption. Equilibrium adoption is inefficient when firms do not internalize all social damages, and sector-independent regulation is helpful but generally not sufficient to restore optimality.

JEL Codes. L16, O31, O41

Thesis supervisor: Daron Acemoglu

Title: Institute Professor

Thesis supervisor: Michael Whinston

Title: Professor of Economics and Sloan Fellows Professor of Management

Thesis supervisor: Jacob Moscona
Title: 3M Career Development Assistant Professor of Economics

Acknowledgments

I am extraordinarily grateful to my advisors, Daron Acemoglu, Mike Whinston, and Jacob Moscona, for their boundless generosity with time, encouragement, feedback, and insights over the past six years. From the very beginning of my graduate studies in his class on economic growth, Daron has had a profound impact on my intellectual trajectory. He has taught me how to ask important and empirically-grounded questions and to focus on answering them the *right way* – with depth and clarity about the key economics, despite the challenges. Mike has been an exceptional role model of an economic theorist. Our conversations have thoroughly shaped my understanding of what it means to develop new theory, and I can only hope to emulate the care and precision with which he builds and deconstructs economic arguments. I have also benefited enormously from Jacob's depth of knowledge about innovation, as well as his unmatched ingenuity with empirical economics. His enthusiasm for tackling new questions is both infectious and inspirational.

Much thanks is due to my recent co-authors, Daron Acemoglu, Rebekah Dix, and Maxim Troshkin, who have taught me so much about what it means to produce research. I have also benefited from conversations with many in the broader economics community at MIT: David Atkin, Martín Beraja, Ricardo Caballero, Isa Di Tella, Dave Donaldson, Sarah Gertler, Shin Kikuchi, Danny O'Connor, Nancy Rose, Adam Solomon, Rafael Veiel, Iván Werning, and Chris Wolf. Finally, I thank my undergraduate advisors, Francesca Molinari, Mathieu Taschereau-Dumouchel, and Maxim Troshkin, who introduced me to economics in the first place.

This thesis is dedicated to my family: To my future wife María, thank you for your love, support, and intellectual companionship over the past several years and for many to come. To my parents Debbie and Todd, thank you for your love, sacrifices, and unyielding encouragement from the beginning. And to my sister Kelsey, thank you for the endless inspiration to continue tackling new challenges. Without you all, none of this would have been possible.

Contents

List of Figures

List of Tables

Chapter 1

Technology Choice, Spillovers, and the Concentration of R&D

1.1 Introduction

Rapid progress in artificial intelligence (AI) has sparked intense debate about how firms are developing this technology and whether governments should intervene. Two issues are prominent: First, AI is a general-purpose technology that can be applied to a wide variety of tasks, from predicting the structure of proteins to personalizing online ads. Firms developing AI must choose which applications to target, but there is broad concern that the resulting *direction of innovation* may be not be in society's best interest.[1] Second, many of these decisions are made by just a handful of large, incumbent technology firms. These firms may face different innovation incentives than society or smaller entrants, and they also control vast research and development (R&D) resources that lend them influence over the direction of innovation. Figure 1.1 shows that the *concentration of R&D* is a recent and growing macro phenomenon. In the past decade, the share of US R&D expenditures accounted for by the top five US public firms has nearly doubled to 22.5%. This trend coincides exactly with the rise of the "Big Five" technology firms to the top of the R&D rankings.

Concerns about the direction of innovation and the concentration of R&D are salient for AI, but they apply to many industries with large incumbents, including the aerospace and defense, automotive, and pharmaceutical industries. They also raise several fundamental questions about the relationship between market structure and innovation:

[1] For example, see Acemoglu (2021), Brynjolfsson (2023), and un.org/sg/en/content/sg/statement/2024-01-17/secretary-generals-special-address-the-world-economic-forum-delivered.

Figure 1.1: Share of US R&D Expenditures, Top 5 US Public Firms 1970-2022

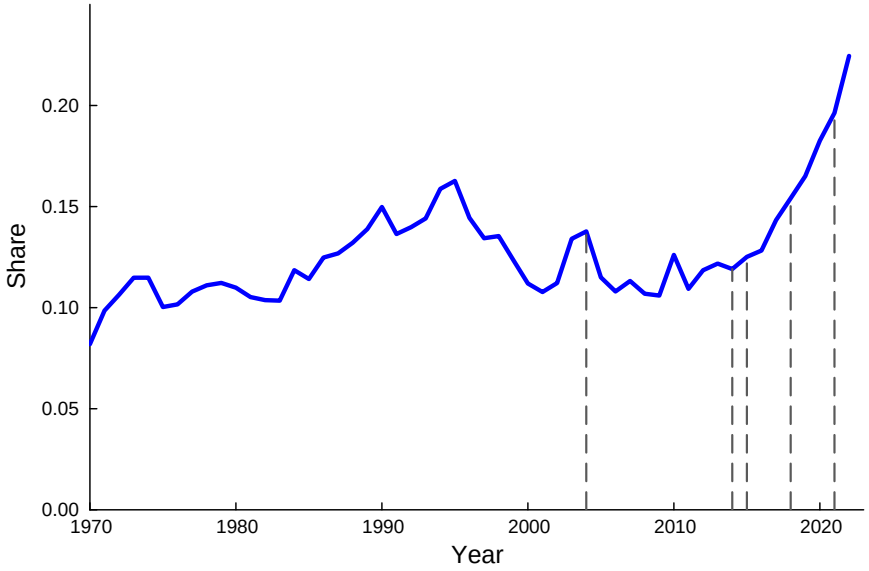

Notes: The vertical lines mark the first year Microsoft (2004), Google (2014), Amazon (2015), Apple (2018), and Meta (2021) each became one of the top five R&D firms. All except for Microsoft have remained in the top five after entry; Microsoft has continuously appeared since 2011. Data on firm R&D expenditures come from Compustat North America, and data on aggregate US R&D expenditures come from the National Science Board.

1. What incentives drive the innovation direction of incumbents, and how do they differ from the incentives faced by entrants or a welfare-minded social planner?

2. How do firms' decisions collectively determine the aggregate direction of innovation?

3. How does the concentration of R&D within large firms affect aggregate innovation?

The central argument of this paper is that our current answers to these questions are incomplete. They overlook how a basic feature of the innovation process — the accumulation of knowledge — drives incumbents and entrants to develop different technologies, while linking these decisions through spillovers to determine the aggregate direction of innovation. I build a new model of directed innovation and firm dynamics to clarify this mechanism. The model reveals how market structure influences innovation not just through competition, but because firms with different accumulated expertise choose to develop different technologies. Their discoveries create new innovation opportunities for all firms, so that an industry's market structure can affect its long-run direction. This direction may be socially inefficient precisely because of these spillovers. And the concentration of R&D plays a critical role by shaping innovation incentives and the market structure itself.

I provide evidence for the theory through an empirical analysis of firm patenting. Firms generally build on their own discoveries, and incumbents with greater patenting experience are slower to innovate emerging technologies. I also discuss a case study on the development of mRNA vaccines, where my theory offers a simple explanation for firm innovation decisions. This case demonstrates the advantages of decentralized R&D in promoting the exploration of new technologies, contrasting sharply with the current pattern of innovation in AI.[2]

Model. I formulate the theory in an endogenous growth model with two key ingredients. First, there are two technologies used by a large number of firms in production. Each firm can employ a team of scientists to innovate and raise its technology-specific qualities (productivities). But innovations for one technology are not useful for the other, so firms must choose *which* technology to innovate. In this sense, the technologies represent not just production processes, but *technological paradigms* that structure how firms can produce and innovate (Dosi, 1982). The concentration of R&D across firms is controlled by the number of scientists a firm can manage. With the aggregate number of scientists fixed, larger team sizes allow for faster innovation but within fewer firms.

Second, innovations produce knowledge for each technology that raises the productivity of research directed toward it. Some knowledge accumulates publicly and can be used by any firm, a key source of spillovers (P. M. Romer, 1990). But crucially I assume that knowledge also accumulates *within the firm*, embodied in its qualities. For example, firms may specialize on particular research lines within technologies, or firms may work as a means to internalize knowledge spillovers between scientists. As a result, a firm's current research productivities depend directly on its past innovation decisions, and firm boundaries matter in the growth process. This assumption is consistent with recent evidence from the auto industry (Aghion et al., 2016). I provide evidence from a broader range of industries in Section 1.5.

All firms initially innovate for an "old" technology, and incumbent firms exit at random and are replaced by new entrants. As in existing models of the firm-size distribution (e.g., Luttmer, 2007), the combination of innovation and stochastic entry and exit generates heterogeneity — incumbents with a long record of innovation have higher qualities than recent entrants. I then consider the unanticipated arrival of a "new" technology. The new technology allows faster growth than the old technology in the long run, but with little accumulated knowledge it offers slower growth in the short run. In the context of AI, the old technology might represent integration with existing products like office productivity software or internet search, while

[2] See Klinger et al. (2020), Jurowetzki et al. (2021), and Ahmed et al. (2023) for suggestive evidence that growing concentration of AI research in corporations is driving falling diversity in research topics.

the new technology might represent more radical applications like automated drug discovery or adaptive learning. Firms must choose which technology to innovate; the bulk of the paper analyzes their equilibrium choices.

Results. The model yields four main insights. First, greater concentration of R&D raises the economy's growth rate along its balanced growth paths, in which all firms innovate the same technology (Proposition 1). This result establishes a baseline role for *the firm* in economic growth even without a technology choice: Larger firms can focus more R&D resources on their research lines, developing greater expertise that helps them innovate even more quickly.[3] It also suggests an optimistic interpretation of Figure 1.1, that rising concentration of R&D might simply allow faster innovation by the largest firms.

To assess whether this holds when firms face a technology choice, I consider the introduction of the new technology. The second main result shows that innovation decisions are both *path-dependent* and *forward-looking* at the firm level (Proposition 2): Firms have incentives to continue innovating a technology in which they have developed expertise, but they also consider whether to pursue faster long-run growth by innovating the new technology. The knowledge accumulation mechanism thus generates endogenous comparative advantage *in innovation*, predicting that experienced incumbents will continue innovating the old technology while younger firms and entrants will embrace the new one. For example, the "Big Five" technology firms from Figure 1.1 have natural incentives to incorporate AI into their existing products, even if other applications hold greater long-run promise.

This explanation for heterogeneous innovation incentives differs fundamentally from most work on market structure and innovation. A large literature following Arrow (1962) studies how competition for monopoly rents can incentivize incumbents to innovate more or less than entrants, depending on whether innovation cannibalizes incumbents' existing products or allows them to escape competition from entrants.[4] A distinct literature following Henderson and Clark (1990) emphasizes that incumbent firms may have an *absolute* disadvantage at innovating new technologies because of organizational rigidities.[5] As a result, the arrival of new technologies that compete with old ones can trigger the failure of established incumbents. The knowledge accumulation mechanism I identify is independent of these effects. I show this by deliberately

[3] This result also offers a novel foundation for J. A. Schumpeter's (1942) claim that larger firms boost growth, but independent of his original justification that size enables greater surplus extraction in the goods market.

[4] The "replacement effect" is formalized by Arrow (1962) and Reinganum (1983), while the "escape-competition effect" is formalized by Gilbert and Newbery (1982), Aghion et al. (2001), and Aghion et al. (2005). Acemoglu and Cao (2015) explore the implications of the replacement effect for incremental and radical innovations, while Igami (2017a) provides an empirical analysis of these issues in the hard drive disk industry.

[5] See also Henderson (1993) and Christensen (1997).

constructing the model to rule out competition between firms or technologies in the goods market. I also assume that incumbents do not have an absolute disadvantage at innovating the new technology when it arrives. Rather, the key feature in my model is their *comparative* disadvantage at innovating the new technology.

I next turn to the aggregation of firm innovation decisions in equilibrium. The third main result shows that the economy's initial *market structure* (firm-quality distribution) can affect the aggregate technology choice in the long run (Proposition 4). As in past work on technology paradigms, spillovers through public knowledge encourage firms to innovate the same technology, so that initial innovation decisions can permanently "tip" the economy toward one technology or the other (W. B. Arthur, 1989; Acemoglu et al., 2012). But in my model firms' innovation decisions depend on their expertise, so the direction of this effect depends critically on the initial market structure. For example, an industry dominated by large incumbents may fail to exploit a new technology not just because the incumbents favor the old technology, but because they generate spillovers that induce other firms to join them. These complementarities introduce substantial technical challenges, which I resolve by defining, characterizing, and proving the existence of *monotone equilibria*. Monotone equilibria are tractable and allow for rich interactions between firms, reflecting the patterns of innovation found in the case study and the empirical analysis.

The final main result reconsiders the concentration of R&D when firms face a technology choice (Proposition 5). Greater concentration exacerbates both the path-dependent and forward-looking forces: It allows incumbents to accumulate more expertise for the old technology before the new one arrives, but it also enhances the new technology's growth advantage. The former dominates when the discount rate is sufficiently high. An increase in R&D concentration can then induce "lock-in" to the old technology, and a simple quantitative example in Section 1.6 shows that this holds for empirically reasonable parameter values. Lock-in may be inefficient because of knowledge spillovers, and in general a social planner would always transition to the new technology more quickly than in equilibrium (Proposition 6). These findings suggest an alternative interpretation to Figure 1.1: The rising concentration of R&D may drive lower diversity in research as firms follow in the footsteps of large incumbents, potentially *reducing* long-run growth as alternative directions remain unexplored.

Evidence. To provide evidence for the theory, I conduct a case study and an empirical analysis of US patent data. The case study traces the development of mRNA vaccine technology, which was critical to the recovery from the COVID-19 pandemic of 2020-2023. Historical accounts emphasize that small biotechnology firms were largely responsible for innovating this "new"

technology, while large pharmaceutical incumbents with expertise in conventional ("old") vaccines decided against it (Dolgin, 2021); these decisions are also captured in the patent record. Competition- and organization-based theories of innovation incentives have difficulty explaining this pattern, because conventional and mRNA vaccines are often not competing for applications. My theory instead offers a simple explanation based on comparative advantage in innovation, and it rationalizes why several incumbent firms have recently begun to explore mRNA technology. This case also suggests the potential stakes of R&D concentration: Had incumbents exercised greater control over R&D resources, we might have been left without a crucial tool to fight the COVID-19 pandemic.

To show that the lessons of the case study generalize, I study firm innovation decisions using a panel of US patents matched to US public firms over 1980-2021 (Arora et al., 2024). I establish three facts consistent with the theory. First, a firm's current patenting is highly correlated with its previous patenting, controlling for R&D expenditures and a variety of other determinants of innovation. This provides evidence that technological knowledge accumulates within firms, raising their research productivities. Second, after clustering patents according to the new technologies identified by Kalyani et al. (2023), I find that a firm's patenting for a technology is better predicted by previous patenting *within that technology* than by patenting in general. This suggests that knowledge is not only cumulative within firms, but technology-specific — two critical assumptions of the theory. Finally, I show that incumbents with greater patenting experience innovate substantially *less* for new technologies than less-experienced firms, though this gap shrinks as the technologies mature. This finding supports the theory's main prediction that experienced incumbents should be reluctant to embrace new technologies, and it reflects the same pattern of innovation exhibited in the case study. Thus the empirical analysis broadly supports both the assumptions and implications of the theory, and it also provides parameter estimates used to calibrate the quantitative example in Section 1.6.

Related Literature. I contribute to an expansive literature on endogenous growth (P. M. Romer, 1986, 1990; Grossman and Helpman, 1991; C. I. Jones, 1995), particularly work relating growth to market structure or firm dynamics (Aghion and Howitt, 1992; Klette and Kortum, 2004; Aghion et al., 2005; Acemoglu and Cao, 2015; Akcigit and Kerr, 2018; Akcigit and Ates, 2021, 2023). This literature primarily addresses growth within a technological paradigm and views market structure through the lens of competition, showing how the associated "replacement" and "escape-competition" effects shape the *rate* of innovation. I instead consider a choice between technological paradigms, and I highlight a distinct role for market structure to affect the *direction* of innovation through knowledge accumulation. I find that conditions favorable for growth within a technological paradigm (e.g., concentrated R&D) can reduce growth by

hindering the emergence of a new one. My framework also builds on models of the firm-productivity distribution that imply a role for *the firm* in accumulating knowledge (Luttmer, 2007; Lucas and Moll, 2014; Benhabib et al., 2021; König et al., 2022).

More closely related is a literature on paradigms and increasing returns to scale in the development and adoption of new technologies. Early work emphasizes how firms typically innovate within established paradigms, limited by bounded rationality or organizational constraints from exploring new ones (Nelson and Winter, 1982; Dosi, 1982).[6] W. B. Arthur (1989) first discusses how technological lock-in can arise from complementarities in adoption decisions. Farrell and Saloner (1986) and Katz and Shapiro (1986) develop similar ideas in the theory of network effects, and they explore how owners of the underlying technologies can strategically influence adoption. I consider network effects in the context of innovation, focusing on how heterogeneity across firms (market structure) affects the equilibrium choice of a paradigm.

A growing body of work integrates technological paradigms into models of economic growth. Acemoglu (2011) shows that creative destruction within paradigms can dissuade firms from exploring alternatives, leading to too little technological diversity in equilibrium. I show how loss of diversity can arise instead through knowledge accumulation that raises the opportunity cost of exploration, with a critical role played by market structure and the concentration of R&D. Few papers in this literature consider how different innovation incentives for incumbents and entrants impact the aggregate technology choice. Acemoglu et al. (2016) build a quantitative model of directed innovation and firm dynamics to study the speed of the clean transition under various policy regimes. They assume firms become more productive at innovating technologies based on their past experience, but this model is explicitly designed to aggregate: As in Klette and Kortum (2004), knowledge accumulates only at the product level, so that firm boundaries and market structure play no role in the innovation process. The most closely related paper to mine is a contemporaneous contribution by Aghion et al. (2024). They consider an innovation process similar to that of Acemoglu et al. (2016) and break aggregation by assuming that firms of different ages face different credit constraints. As a result, firm boundaries and market structure affect innovation only insofar as credit constraints bind, and they use their model to quantitatively explore how credit conditions can affect clean innovation.

I also contribute to the broad literature on market structure and innovation. In addition to the seminal contributions discussed above, recent work assesses how incumbents can block innovation by competitors through acquistions (Cunningham et al., 2021), defensive patenting (Argente et al., 2020), and pre-emptive hiring of inventors (Akcigit and Goldschlag, 2023). The

[6] This work was itself inspired by a literature on *scientific* paradigms following Kuhn (1970). See Brock and Durlauf (1999) for a formalization emphasizing the role of conformity effects.

knowledge accumulation mechanism I study works independently of competitive pressures, and hinges instead on how complementarities between firms lend incumbents influence over the direction of innovation. In a surprising contrast, I find that incumbents can reduce long-run growth *precisely by innovating according to their expertise*, generating knowledge spillovers on other firms that raise the opportunity cost of exploring new technologies.

Finally, my empirical analysis contributes to a large literature using panel data to estimate the determinants and effects of innovation at the firm level (e.g., Griliches, 1998; Bloom et al., 2013b; Kalyani et al., 2023; see Hall et al., 2010 for a review), particularly recent work on directed innovation and path dependence in clean and dirty technologies (Dechezleprêtre et al., 2014; Aghion et al., 2016; Dugoua and Gerarden, 2023). I provide evidence for path dependence from a broader set of industries. I also document the novel result that incumbents are slow to innovate emerging technologies, but do so after other firms make early progress.

Outline. The rest of this paper is organized as follows: Section 1.2 presents the mRNA case study. Section 1.3 sets up the baseline model of directed innovation and firm dynamics and characterizes its balanced growth paths. Section 1.4 analyzes equilibria after the introduction of the new technology. Section 1.5 describes the empirical analysis, and Section 1.6 presents a simple quantitative example. Section 1.7 concludes.

1.2 Case Study: mRNA Vaccines

In this section, I provide a brief case study of the development of mRNA vaccines. This case study is instructive in three ways. First, it provides a clear example of the economic setting described by the theory, with distinct old and new technologies that produce different goods and build on different bodies of knowledge. Second, evidence from historical accounts and the patent data shows that incumbent and entrant firms innovated different technologies. I argue that the competition- and organization-based theories of innovation incentives cannot explain this pattern, which is instead consistent with the knowledge accumulation mechanism outlined in my theory. Finally, this case also offers a sharp contrast to the concentration of R&D in AI: Pharmaceutical research is broadly decentralized across many firms, universities, and public and private research institutes. I argue that life-saving mRNA vaccines emerged *precisely because* of this decentralization, underscoring the substantial risks to concentrating R&D among a small number of firms in any industry.

1.2.1 Background

The spread of COVID-19 in early 2020 triggered one of the worst pandemics in a century and the sharpest economic contraction in the post-war period. Within a year, COVID-19 caused over half a million deaths in the United States, vaulting to third among the country's leading causes of mortality.[7] A combination of pandemic-driven uncertainty and strict public health interventions drove the unemployment rate in the United States to a high of nearly 15%, while real output fell by 7.5% in the second quarter of 2020 alone. Given the severity of these initial impacts, the subsequent recovery is all the more remarkable: In just over three years, rapid growth in productivity, employment, and business creation returned the unemployment rate to its pre-pandemic level and real output to its pre-pandemic trend (de Soyres et al., 2024). The infection rate and the risk of hospitalization or death from COVID-19 have also fallen dramatically from their peaks in 2021. A variety of policies and treatment advances supported these outcomes, but one critical innovation stands out: vaccines.

The race to develop vaccines for COVID-19 began at the onset of the pandemic, spearheaded by pharmaceutical companies with public backing through the US government's Operation Warp Speed program. The first two vaccines to receive approval were produced by biotech firm Moderna and a joint venture between pharmaceutical giant Pfizer and biotech firm BioNTech. These vaccines proved exceptionally effective, and they remain the most widely adopted COVID-19 vaccines in the United States.[8] Their underlying technology is particularly notable: Both vaccines are based on a novel technique for producing immune resistance that differs fundamentally from the mechanism used by conventional vaccines. Conventional vaccines stimulate resistance by confronting the immune system with a weakened version of a pathogen or one of its constituent proteins, both grown in a lab. The Moderna and Pfizer-BioNTech vaccines instead encode instructions for the production of a protein using a genetic material called *messenger RNA* (mRNA). *mRNA vaccines* ferry these instructions into human cells, where they are used to produce the protein en masse to stimulate immune resistance (Hedestam and Sandberg, 2023).

The contrast between conventional and mRNA vaccines offers a vivid illustration of my theory. This case features clearly identifiable old (conventional) and new (mRNA) technologies, both of which reflect the basic features of technological paradigms: Conventional and mRNA vaccine technologies are not just individual products, but sets of techniques used to *produce* and *improve* a variety of products (vaccines). Each technology also required substantial innovation before a

[7] See covid.cdc.gov/covid-data-tracker for COVID-19 statistics and Murphy et al. (2021) for causes of mortality. See fred.stlouisfed.org for all US macroeconomic statistics.

[8] See ourworldindata.org/grapher/covid-vaccine-doses-by-manufacturer?country= USA

Figure 1.2: Citation-Weighted Patents, 1999-2018

(a) mRNA Therapies

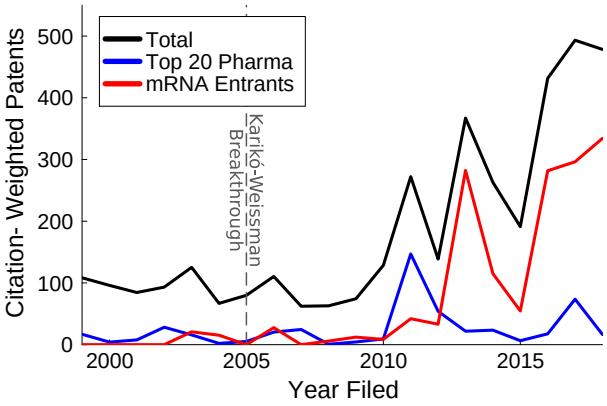

(b) Conventional Vaccines

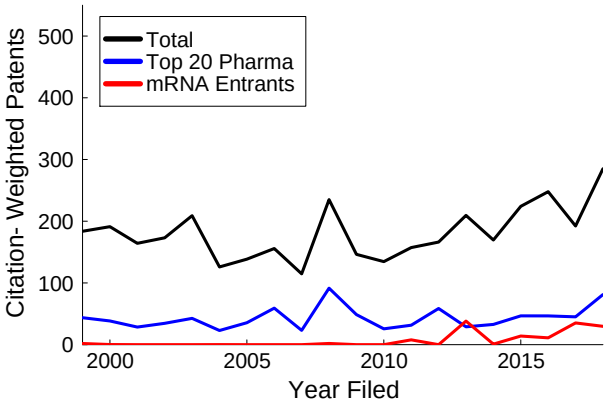

Notes: Figures 1.2(a) and 1.2(b) respectively include all mRNA therapy and conventional vaccine patents filed at the US Patent and Trademark Office (USPTO) over 1999-2018 and granted before 2023. I halt both figures in 2018 to mitigate truncation issues from the lag between patent filing and publication. The vertical dashed line in Figure 1.2(a) marks the year of the Karikó-Weissman breakthrough (Karikó et al., 2005).

viable product could be produced. For example, the practice of conventional vaccination dates back over a millenium, but the production of vaccines for many diseases was not possible before the development of microbiology in the late 19th century (Kinch, 2018). By contrast, mRNA was first discovered only in 1961, and initial experiments suggesting its therapeutic potential were conducted in the early 1990s. The final breakthrough enabling vaccine development came in experiments by Katalin Karikó and Drew Weissman in 2005, for which these researchers were awarded a Nobel Prize in 2023 (Karikó et al., 2005; Hedestam and Sandberg, 2023). As this discussion indicates, the bodies of knowledge for conventional and mRNA vaccine technologies are largely distinct.[9] A firm interested in innovating a new vaccine then faces the important choice of *which* technology to research.[10]

1.2.2 mRNA Innovation

In the case of mRNA vaccines, incumbent and entrant firms made different innovation choices. Figure 1.2 displays the number of patents related to mRNA therapies and conventional vaccines that were filed (and eventually granted) in the United States between 1999 and 2018.[11] To more accurately reflect its scientific contribution, each patent is weighted by the number of forward citations from future patents, controlling for the time horizon after publication. The black line in each figure represents the total number of patents filed each year, including patents assigned to individual researchers, firms, universities, public research organizations, and independent research institutes. The blue line represents the subset of patents assigned to the top twenty pharmaceutical companies by revenue in 2023, while the red line represents patents assigned to just four entrant firms whose founders played a key role in the development of mRNA technology: Moderna, BioNTech, the pharmaceutical firm CureVac, and the now-defunct biotech firm RNARx (Hedestam and Sandberg, 2023). All entrants were founded after 2000 specifically to commercialize mRNA therapies.

Figure 1.2 demonstrates two facts. First, a substantial share of research on mRNA therapies and conventional vaccines is performed outside of the identified incumbent and entrant firms, reflecting the broader decentralization of research in the pharmaceutical industry (Scott Morton and Kyle, 2011). Early contributions to mRNA technology were made primarily by researchers at universities and independent research institutes; Figure 1.2(a) shows that researchers outside of large pharmaceutical firms or the leading mRNA therapy firms remain influential. Second, the small set of entrant firms played a much bigger role in the innovation of mRNA therapies than the larger set of incumbent pharmaceutical firms. A historical account by Dolgin (2021) emphasizes that this was an active decision by many of these incumbents: "In the 1990s and for most of the 2000s, nearly every vaccine company that considered working on mRNA opted to invest its resources elswhere." This remained true even after the final technical barrier to mRNA therapy was resolved in 2005.[12] The entrant firms were responsible not only for much of

[9]The patent data provide additional evidence. Using the sets of mRNA therapy and conventional vaccine patents described below, I find that on average a share 0.14 of citations by an mRNA patent are made to other mRNA patents, while this "within-technology" citation share is 0.39 for conventional vaccine patents. By contrast, on average only a share 0.03 of citations by an mRNA patent are made to conventional vaccine patents, while this "cross-technology" citation share is 0.01 for conventional vaccine patents. The low cross-technology shares suggest that the knowledge produced by one technology is not generally useful for the other.

[10]In reality firms can research both, and the patent data show that they often do. But they must choose between the two technologies at least for the marginal unit of R&D expenditure, and any non-convexities in the required expenditures can make this choice more discrete.

[11]I consider a broader set of patents than just those for mRNA vaccines, because the underlying technology can also be applied to treat a variety of genetic conditions and cancers. See Appendix A.4 for a full description of the patent data, citation weights, and procedure for identifying mRNA therapy and conventional vaccine patents.

[12]In a recent interview, Weissman recalled that after their 2005 breakthrough he "told [Karikó] our phones are

the foundational basic research on mRNA technology, but also the initial therapy development: CureVac conducted the first clinical trial for an mRNA vaccine in 2017, and Moderna followed quickly to test mRNA vaccines against the Zika virus and two strains of the avian influenza virus.[13] Throughout the time period, incumbent pharmaceutical firms remained more active in conventional vaccine innovation than in mRNA therapy innovation (Figure 1.2(b)).

Why were incumbents so reluctant to invest in mRNA technology, while entrants pushed it from the laboratory to pharmacy shelves? I argue that competition-based theories of innovation incentives cannot explain this pattern, because the mRNA and conventional vaccine technologies are often *not competing* for uses. Each type of vaccine has its own advantages and disadvantages that make it suitable for different applications. For example, mRNA vaccines are often faster to develop and easier to produce than conventional vaccines, but to maintain efficacy they must be transported and stored at below-freezing temperatures. This limits their use in many contexts, where conventional vaccines may be more suitable (Gote et al., 2023). Firms developing mRNA therapies have also focused on conditions without existing vaccines and/or treatments, including infectious diseases like HIV and avian flu, genetic diseases like cystic fibrosis, and a variety of cancers.[14] These considerations limit the extent to which mRNA therapies cannibalize or "steal business" from existing therapies, the standard explanation for why incumbents might be reluctant to innovate relative to entrants (Arrow, 1962; Reinganum, 1983). Moreover, in the wake of the technology's success during the COVID-19 pandemic, several incumbent pharmaceutical firms have begun developing mRNA therapies.[15] This observation sheds doubt on organizational theories that suggests incumbents did not innovate because they *could not* innovate (Henderson, 1993).

My theory instead offers a simple explanation based on knowledge accumulation within firms: Incumbent pharmaceutical firms had already developed *expertise* in other areas of research (including conventional vaccines) by the time mRNA technology emerged in the 1990s. In line with the quote above from Dolgin (2021), these firms found it more valuable to continue innovating in those areas than fund early work on mRNA technology. With no such expertise, the entrants instead pursued research for a less developed technology with greater long-run promise. Only after learning from years' worth of innovation by entrants did the incumbent pharmaceutical firms begin seriously investing in mRNA technology. The empirical analysis in Section 1.5 shows that this pattern of innovation by incumbents and entrants generalizes

going to ring off the hook. But nothing happened. We didn't get a single call." (Yu, 2021)

[13] See Alberer et al. (2017) and trials NCT03014089, NCT03076385, and NCT03345043 at ClinicalTrials.gov.

[14] See the product pipelines for Moderna (modernatx.com/en-US/research/product-pipeline) and BioNTech (biontech.com/int/en/home/pipeline-and-products/pipeline.html).

[15] See the product pipelines for Merck (merck.com/research/product-pipeline), Pfizer (pfizer.com/science/drug-product-pipeline), and Roche (roche.com/solutions/pipeline).

to many new technologies. A key contribution of the model in Section 1.3 is to show that it naturally arises when firms accumulate expertise through innovation.

However, the model also shows that this pattern of innovation holds only if entrants have enough R&D resources to push the new technology forward. Otherwise, the entrants may be unable to make sufficient progress before it becomes more valuable for them to abandon the new technology and join the incumbents in innovating the old one. This observation highlights an advantage to decentralized R&D ecosystems like that in the pharmaceutical industry: Although incumbents almost universally decided against funding mRNA therapy research, entrants were able to develop the technology to the point that they could distribute highly effective COVID-19 vaccines within a year of the initial outbreak. This decentralization contrasts sharply with current trends in AI, where the "Big Five" technology firms from Figure 1.1 deploy the vast majority of R&D resources while hiring researchers from academia and startups at a growing rate (Ahmed et al., 2023).[16] The associated risks may be substantial: Incumbents may leave socially valuable technologies undeveloped, and entrants may simply choose to join them instead of independently developing other technologies. The model developed in the next two sections describes why and when this might occur.

1.3 Model

This section describes the baseline model of directed innovation and firm dynamics. I set up the model and define an equilibrium in Section 1.3.1, I solve the "static block" of the model in Section 1.3.2, and I define and characterize balanced growth in Section 1.3.3.

1.3.1 Setup

Consumption. The economy is deterministic and exists in continuous time, populated by a mass $L > 0$ of *workers*, a mass $S > 0$ of *scientists*, and a mass $N \in (0, 1]$ of *entrepreneurs*. To focus attention on firm innovation decisions, I keep the demand side of the economy as simple as possible: All agents have linear preferences over a unique consumption good, with common discount rate $\rho > 0$. The economy admits a representative consumer who evaluates consumption streams $[C(t)]_t$ with the utility function

$$\int_0^\infty \exp(-\rho t) C(t) dt. \tag{1.1}$$

[16] See also https://www.nytimes.com/2024/06/13/opinion/big-tech-ftc-ai.html.

Workers inelastically supply one unit of labor at wage $w_L(t)$, while entrepreneurs and scientists own all firms and earn profits as described below. All agents can risklessly save and borrow against the total value of firms $\mathcal{A}(t)$ at the equilibrium interest rate $r(t)$. Given the path of interest rates $[r(t)]_t$, the representative consumer solves a standard consumption-savings problem to maximize her utility (1.1) subject to her budget constraint and no-Ponzi condition:

$$\dot{\mathcal{A}}(t) \leq w_L(t)L + r(t)\mathcal{A}(t) - C(t), \tag{1.2}$$

$$0 \leq \lim_{t\to\infty} \mathcal{A}(t)\exp\left(-\int_0^t r(s)ds\right). \tag{1.3}$$

I take the final (consumption) good as the numeraire. The main simplifying assumption here is that preferences are linear, so that in equilibrium the interest rate is fixed at $r(t) = \rho$. With less elastic intertemporal preferences, the interest rate would vary along a technological transition and introduce additional dynamics into firm innovation decisions. This general equilibrium effect is potentially interesting but obscures the more fundamental innovation incentives at the heart of my analysis.

Production. A competitive firm produces the final good by combining labor supplied by workers with intermediates. Intermediates come in two types $\theta \in \{A, B\}$ and are supplied by a unit measure of monopolistic *incumbent* firms; I refer to each type θ as a *technology*. Each incumbent owns one intermediate for each technology θ with an endogenous and firm-specific *quality* $q_\theta(t)$. An incumbent is fully characterized by its vector of qualities $q \equiv (q_A, q_B)$ at each time, and I let $F(q,t)$ denote the distribution of qualities across incumbents at time t. In the context of the mRNA case study, the technologies A and B correspond to conventional and mRNA vaccines, and the qualities q represent a firm's ability to produce effective vaccines with each technology.

Given labor input $L(t)$ and intermediate inputs $x_A(q,t)$ and $x_B(q,t)$ from each firm with qualities q, final output is

$$Y(t) \equiv \frac{1}{1-\beta}\left(\int \sum_{\theta\in\{A,B\}} q_\theta^\beta x_\theta(q,t)^{1-\beta} dF(q,t)\right) L(t)^\beta. \tag{1.4}$$

The final producer chooses the inputs $L(t)$ and $[x_\theta(q,t)]_{\theta,q}$ at each time to maximize its profits, taking as given the wage for workers $w_L(t)$ and the intermediate prices $[p_\theta(q,t)]_{\theta,q}$. The wage

$w_L(t)$ is set competitively to clear the market for production labor:

$$L(t) = L. \tag{1.5}$$

The intermediate prices $[p_\theta(q,t)]_{\theta,q}$ are chosen by firms to maximize profits as described below. Each unit of an intermediate is produced using $\gamma > 0$ units of final output. All remaining output is used for consumption, yielding the market-clearing condition

$$Y(t) = C(t) + \gamma \int \sum_{\theta \in \{A,B\}} x_\theta(q,t) dF(q,t). \tag{1.6}$$

This production structure is used frequently in models of endogenous growth with quality upgrading, though I introduce the distinction between two sets of intermediates $\theta \in \{A,B\}$.[17] Its key features are additive separability across intermediates in the final production function (1.4) and the absence of any labor reallocation between intermediates. As a result, the final producer's demand curve for each intermediate is independent of the qualities and prices of all others, $x_\theta(q,t) = q_\theta p_\theta(q,t)^{-\frac{1}{\beta}} L$. The flow profits for firms inherit these properties, as each firm sets the prices of its intermediates $p_A(q,t)$ and $p_B(q,t)$ to maximize profits π given its qualities q:[18]

$$\pi(q) \equiv (q_A + q_B)\,\bar{\pi}, \quad \text{where} \quad \bar{\pi} \equiv \max_p (p-\gamma) p^{-\frac{1}{\beta}} L. \tag{1.7}$$

This expression demonstrates that the economy features no competition between firms or substitution between technologies that could affect innovation incentives. Competition is ruled out because each firm's profits do not depend on other firms' qualities, so firms will not innovate to steal business or preempt competition from rival firms. Similarly, the profits earned by technology A intermediates do not depend on the qualities of technology B intermediates. Such dependence would be natural if technologies A and B instead produced distinct goods that were imperfectly substitutable in demand or production, as in the existing literature on directed innovation. Improvement in one technology would then trigger relative price and market size adjustments that affect the profits for the other.[19] I exclude these competition and "demand-pull" forces to focus instead on how the innovation process itself shapes firms' incentives. The mRNA

[17] For example, see Howitt (1999) and Acemoglu et al. (2006).

[18] Flow profits are also linear and symmetric in the qualities q. Linearity is essentially a normalization from the definition of q_θ, while symmetry ensures that firms have no reason to favor either technology based on differences between the goods markets.

[19] These price and market size effects are detailed in Acemoglu (1998, 2002), and they are applied to study the direction of innovation between clean and dirty technologies in Acemoglu et al. (2012).

case study provides just one example where these assumptions are realistic, but the forces I study arise even with competition in the goods market.

Innovation and Firm Dynamics. Intermediate qualities are determined endogenously through innovation, entry, and exit. Each firm can raise the quality of its intermediates by employing scientists to conduct research. If a firm with qualities $q(t)$ employs $s_\theta(q(t),t)$ scientists to research technology θ at t, the quality of its intermediate evolves according to

$$\dot{q}_\theta(t) = [\lambda q_\theta(t) + \sigma_I K_\theta(t)]\,\eta_\theta s_\theta(q(t),t). \tag{1.8}$$

Here $\eta_\theta > 0$ denotes the basic productivity of research for θ. This productivity is augmented by the accumulation of knowledge that raises research productivity, represented by the term in brackets. A central feature of my theory is that knowledge accumulates through two channels. First, following an extensive literature on endogenous growth initiated by P. M. Romer (1990), I suppose that each firm can learn from innovations made by all others. This public knowledge is technology-specific, and it is embodied in each technology's *knowledge stock* $K_\theta(t) > 0$. The knowledge stock reflects, for example, all information about technology θ found in the patent data and scientific publications, or shared among scientists at conferences. Its initial value $K_\theta(0) > 0$ is exogenous, and it increases as firms innovate for θ:

$$\dot{K}_\theta(t) = \int \underbrace{[\lambda q_\theta(t) + \sigma_I K_\theta(t)]\,\eta_\theta s_\theta(q(t),t)}_{\dot{q}_\theta(t)}\, dF(q(t),t). \tag{1.9}$$

Let $K(t) \equiv (K_A(t), K_B(t))$ denote the vector of knowledge stocks at time t. Second, I make the novel assumption that knowledge also accumulates *within each firm*, embodied in its quality $q_\theta(t)$. This internal knowledge is again technology-specific, and it represents any information produced by the firm that disproportionately improves its own future research efforts. The parameters $\lambda \geq 0$ and $\sigma_I \geq 0$ in (1.9) control the extent to which firms draw on internal and public knowledge when innovating, respectively.

Several mechanisms could induce knowledge accumulation within firms. For example, intermediates for a technology θ could represent distinct research lines, and the knowledge generated within a research line may be more useful for future innovation than knowledge generated by others. A firm may work as a mechanism to coordinate scientists on a single research line and accelerate innovation. The mRNA case study reflects this interpretation: mRNA technology can be used to treat or prevent a variety of different diseases, including genetic conditions, cancer, and viral infections like COVID-19. Each application builds on and

contributes to our general understanding of mRNA technology. But the technology must also be tailored to each case, and the leading mRNA firms have indeed specialized within particular applications.[20] Alternatively, work dating back to Marshall (1890) documents that innovation is spatially concentrated, suggesting that knowledge spillovers between researchers happen through direct communication (e.g., Jaffe et al., 1993; Kalyani et al., 2023). A firm may employ researchers to facilitate communication or ensure proximity in a common office or lab, again catalyzing "internal" knowledge spillovers. Both mechanisms imply a role for the firm *per se* in the innovation process, rooted in the insights of Coase (1937) and Alchian and Demsetz (1972) that firms form to internalize externalities and coordinate complementary activities. This role for the firm is essential to the model, and I discuss its implications throughout the analysis below.

Consistent with these microfoundations, I assume that scientists are organized into teams managed by entrepreneurs. A firm innovates by employing one entrepreneur and her team to research its intermediates. For simplicity, I assume that each entrepreneur manages the same number $s \equiv S/N$ of scientists, so all innovating firms have equally sized "R&D departments." However, a fraction $1 - N$ of firms are unable to innovate at each time because entrepreneurs are in scarce supply. The team size s is a natural measure of the *concentration of scientists* across firms: Holding the total mass of scientists S fixed, an increase in s reduces the number of firms N that can innovate while raising the R&D resources available to any firm that still can. To reduce notation without any essential changes to the model, I assume that any firm unable to hire an entrepreneur must exit. The distribution $F(q,t)$ describes the qualities for active firms and is scaled by their total measure $N > 0$.

In this baseline model, firm entry and exit are exogenous: Each incumbent receives an independent exit shock at rate $\delta > 0$, at which time it ceases production and is replaced by an entrant with initial qualities

$$q_\theta^E(t) \equiv \sigma_E K_\theta(t). \tag{1.10}$$

Here $\sigma_E \geq 0$ determines the strength of knowledge spillovers to entrants. The entrant immediately employs the entrepreneur and scientists from the exiting firm.

Before exit, I assume that all profits generated by an innovating firm accrue to its entrepreneur and scientists. The scientists are then allocated at each time to maximize the firm's value $V(q,t)$, taking as given its initial qualities q and the trajectories of the knowledge stocks $[K(t)]_t$ and

[20] Moderna and BioNTech have focused on mRNA vaccines for contageous diseases and cancer, respectively.

the interest rate $[r(t)]_t$:

$$V(q,t) = \max_{[s_\theta(q(\tau),\tau)]_{\theta,\tau}} \int_t^\infty \exp\left[-\int_0^\tau (r(t')+\delta)dt'\right]\pi(q(\tau))d\tau, \tag{1.11}$$

where maximization is subject to the resource constraint $s_A(q,\tau)+s_B(q,\tau)\leq s$ and the quality evolution equation (1.8). Asset market clearing requires that the total demand for assets from the consumer equal the total value of incumbents at each time:

$$\mathcal{A}(t) = \int V(q,t)dF(q,t). \tag{1.12}$$

This specification of innovation, entry, and exit implies that the distribution $F(q,t)$ evolves according to the Kolmogorov forward equation (KFE)

$$\frac{\partial F(q,t)}{\partial t} = -\int \dot{q}_A(t)F\left(q_A, dq'_B, t\right) - \int \dot{q}_B(t)F\left(dq'_A, q_B, t\right) + \delta N \mathbb{1}_{q\geq q^E(t)} - \delta F(q,t). \tag{1.13}$$

The mass of incumbents with qualities below q declines as incumbents with technology A qualities q_A improve their A intermediates. The first term on the right side of (1.13) captures the corresponding loss of mass per unit of time, or *flux*, through the boundary $\{q' : q'_A = q_A\}$. The second term similarly captures the flux through the boundary $\{q' : q'_B = q_B\}$. The third term reflects the increase in mass from entry, while the last term gives the fall in mass from exit.

Definition 1. An *equilibrium* is a set of trajectories for total output $[Y(t)]_t$, consumption $[C(t)]_t$, assets $[\mathcal{A}(t)]_t$, labor demand $[L(t)]_t$, intermediate quantities $[x_\theta(q,t)]_{\theta,q,t}$, wages $[w_L(t)]_t$, intermediate prices $[p_\theta(q,t)]_{\theta,q,t}$, knowledge stocks $[K(t)]_t$, allocations of scientists $[s_\theta(q,t)]_{\theta,q,t}$, incumbent values $[V(q,t)]_{q,t}$, and the quality distribution $[F(q,t)]_{q,t}$ such that

1. the representative consumer chooses her consumption and asset holdings to maximize her utility (1.1) subject to her budget constraint (1.2) and no-Ponzi condition (1.3);
2. labor demand and intermediate quantities are chosen by the final producer to maximize profits, given input prices;
3. intermediate prices are chosen to maximize incumbent flow profits (1.7);
4. the markets for labor, goods, and assets clear (1.5, 1.6, 1.12);
5. the knowledge stocks satisfy the evolution equation (1.9);

6. the incumbent value function satisfies (1.11), while the scientist allocation solves the corresponding maximization problem; and

7. the quality distribution satisfies the KFE (1.13).

I maintain several parameter restrictions throughout. I assume $\lambda + \sigma_I > 0$ so that innovation is possible, though either of the parameters λ or σ_I may equal zero. I also assume that spillovers to entrants are positive $\sigma_E > 0$, which ensures that the economy generates long-run growth. Finally, I assume that the exit rate δ is sufficiently large that an initial cohort of incumbents cannot generate aggregate growth in the long run:

$$\delta > (\lambda + \sigma_I N) \max\{\eta_A, \eta_B\} s. \tag{1.14}$$

1.3.2 Equilibrium: Static Block

To simplify the characterization of equilibrium, I note that the quality distribution F and the knowledge stocks $K \equiv (K_A, K_B)$ are the state variables in this economy. With the exception of the incumbent value function V and the allocation of scientists $[s_\theta(q,t)]_{\theta,q,t}$, all remaining equilibrium variables are statically determined as a function of the state (F, K). In fact, only the *aggregate qualities* $Q \equiv (Q_A, Q_B)$ are needed to determine this "static block" of the economy, where the aggregate quality of technology θ is

$$Q_\theta(t) \equiv \int q_\theta dF(q,t).$$

Integration by parts reveals that the aggregate quality of θ increases with incumbent innovation but potentially declines as incumbents exit:[21]

$$\dot{Q}_\theta(t) = \dot{K}_\theta(t) + \delta \left[N q_\theta^E(t) - Q_\theta(t) \right]. \tag{1.15}$$

The following lemma characterizes the static block of the economy along with the equilibrium interest rate $r(t)$:

Lemma 1. *In equilibrium,*

1. *the interest rate is $r(t) = \rho$;*

[21] The convention in the endogenous growth literature is to identify the knowledge stock K_θ with the aggregate quality Q_θ. This literature generally does not consider entrants that may arrive with lower qualities than the average incumbent ($\sigma_E < 1$), which directly reduces aggregate quality. I introduce the distinction between K_θ and Q_θ so that public knowledge is not artificially reduced by firm entry and exit.

2. *intermediate prices and the wage for workers are*

$$p_\theta(q,t) = \frac{\gamma}{1-\beta} \quad \textit{and} \quad w_L(t) = [Q_A(t) + Q_B(t)]\,\frac{\beta \bar{x}^{1-\beta} L^{-(1-\beta)}}{1-\beta};$$

3. *production labor demand is* $L(t) = L$*;*

4. *intermediate quantities and the corresponding flow profis are*

$$x_\theta(q,t) = q_\theta \bar{x}, \quad \textit{where} \quad \bar{x} \equiv L\left(\frac{\gamma}{1-\beta}\right)^{-\frac{1}{\beta}},$$
$$\pi(q_\theta) = q_\theta \bar{\pi}, \quad \textit{where} \quad \bar{\pi} \equiv \beta L\left(\frac{\gamma}{1-\beta}\right)^{-\frac{1-\beta}{\beta}};$$

5. *total output and consumption are*

$$Y(t) = [Q_A(t) + Q_B(t)]\,\bar{Y}, \quad \textit{where} \quad \bar{Y} \equiv \frac{\bar{x}^{1-\beta} L^{\beta}}{1-\beta},$$
$$C(t) = [Q_A(t) + Q_B(t)]\,\bar{C}, \quad \textit{where} \quad \bar{C} \equiv \bar{Y} - \gamma \bar{x}.$$

All proofs are found in Appendix A.1. The wage $w_L(t)$, total output $Y(t)$, and consumption $C(t)$ are all linear in the sum of the aggregate qualities $Q_A(t) + Q_B(t)$, again reflecting the absence of substitution between technologies A and B. As noted above, the fixed interest rate $r(t) = \rho$ aids tractability. In the remaining analysis, I focus on the "dynamic block" of the economy: firm innovation decisions and the resulting dynamics of the state (F, K).

1.3.3 Balanced Growth

The main goal of the analysis is to characterize innovation decisions and transitional dynamics after a "new technology" B is introduced into the economy. To lay a foundation for this exercise, I first study the economy's steady states, or *balanced growth paths* (BGPs). The characterization of BGPs in Proposition 1 yields the first core result of the theory: When firms accumulate knowledge internally ($\lambda > 0$), increasing the concentration of scientists s always raises the economy's BGP growth rate. It also produces changes in the firm-quality distribution F that play a critical role in the analysis of technology choice in Section 1.4.

Toward a definition of balanced growth, note that the distribution of qualities F is never stationary in equilibrium: As firms innovate and contribute to the knowledge stocks K_θ, they improve the initial qualities for entrants $q_\theta^E(t) = \sigma_E K_\theta(t)$. This process repeats as entrants

begin innovating, and it produces an upward shift in the distribution of qualities over time. We can attempt to stabilize the distribution by normalizing the qualities q_θ by the quality of an entrant $q_\theta^E(t)$. Define the *relative quality* of an intermediate by

$$z_\theta(t) \equiv \frac{q_\theta(t)}{q_\theta^E(t)}, \tag{1.16}$$

and denote a firm's vector of relative qualities by $z = (z_A, z_B)$. Let $H(z,t) \equiv F(zq^E(t), t)$ denote the corresponding distribution, where $zq^E(t)$ denotes the pointwise product $(z_A q_A^E(t), z_B q_B^E(t))$.

Definition 2. A *balanced growth path (BGP)* is an equilibrium in which

1. K_θ and Q_θ grow at a constant rate $g_\theta^* \geq 0$ for each technology $\theta \in \{A, B\}$; and
2. the relative quality distribution is stationary: $H(z,t) = H^*(z)$ for all $t \geq 0$.

We can observe immediately that the economy admits at most three BGPs, only two of which are locally stable. For each technology θ, there exists a stable BGP in which all firms exclusively innovate for θ. As a result, technology θ grows asymptotically while the other technology θ' permanently stagnates. These two very different BGPs arise because innovation for each technology features dynamic increasing returns: Innovation for θ produces knowledge externally (through K_θ) and within the firm (through q_θ) that raises the productivity of future research for θ. With positive spillovers to entrants $\sigma_E > 0$, any technology that develops a large enough lead in accumulated knowledge attracts all innovation by entrants, reinforcing its advantage. Note that without substitution between technologies A and B in demand or production, there are no relative price adjustments that could redirect innovation toward the lagging technology and ensure a unique "interior" BGP.

Whenever the two stable BGPs exist, there also exists a third BGP that features equal and positive growth rates for both technologies. Scientists must be balanced across the technologies to sustain equal growth rates. However, this BGP is fragile to perturbations in which a greater share of scientists is temporarily directed to just one of the two technologies, which can push the economy toward one of the stable "corner" BGPs. Finally, we can rule out BGPs in which the growth rates g_A^* and g_B^* are both positive but unequal: Eventually all entrants would choose to research the technology with the faster growth rate, and the growth rate of the other would be driven to zero.

I restrict attention to the economy's two stable BGPs. To characterize them, suppose all firms permanently direct their scientists toward technology θ. The evolution of intermediate qualities

is then completely mechanical:

$$\dot{q}_\theta(t) = [\lambda q_\theta(t) + \sigma_I K_\theta(t)]\,\eta_\theta s \quad \text{and} \quad \dot{q}_{\theta'}(t) = 0.$$

With no innovation directed toward θ', all relative qualities $z_{\theta'}$ remain fixed at one. The BGP is then summarized by the growth rate g^*_θ and the marginal distribution $H^*_\theta(z_\theta)$ of relative qualities for θ. We can use the evolution equation (1.9) for the knowledge stock K_θ to write the growth rate g^*_θ as a function of the distribution H^*_θ:

$$g^*_\theta = \frac{\dot{K}_\theta}{K_\theta} = \int [\lambda \sigma_E z_\theta + \sigma_I]\,\eta_\theta s dH^*_\theta(z_\theta). \tag{1.17}$$

The distribution H^*_θ is determined by the evolution of qualities through the KFE (1.13). To see this, let $F_\theta(q_\theta, t)$ denote the non-stationary marginal distribution of qualities for technology θ. This distribution satisfies the following one-dimensional version of the KFE:

$$\frac{\partial F_\theta(q_\theta, t))}{\partial t} = -[\lambda q_\theta + \sigma_I K_\theta(t)]\,\eta_\theta s f_\theta(q_\theta, t) + \delta\left[N \mathbb{1}_{q_\theta \geq q^E_\theta(t)} - F_\theta(q_\theta, t)\right].$$

Here $f_\theta(q_\theta, t) \equiv \partial F_\theta(q_\theta, t)/\partial q_\theta$ denotes the density of F_θ. The first term denotes the loss of mass as firms with the quality q_θ innovate, while the second term denotes the net change in mass due to entry and exit. Since the equation $H^*_\theta(z_\theta) = F_\theta(z_\theta q^E_\theta(t), t)$ must hold for all t along a BGP, we can differentiate to find a time-invariant differential equation for H^*_θ:

$$0 = -\left[\left(\lambda - \frac{g^*_\theta}{\eta_\theta s}\right) z_\theta + \frac{\sigma_I}{\sigma_E}\right] \eta_\theta s h^*_\theta(z_\theta) + \delta N - H^*_\theta(z_\theta). \tag{1.18}$$

The solution to this differential equation gives the distribution H^*_θ as a function of the growth rate g^*_θ. Candidates for the BGP growth rate g^*_θ and stationary distribution H^*_θ must solve the system (1.17, 1.18).

To ensure that a solution exists and delivers finite values for firms and the consumer, I maintain the following parameter restrictions:

$$s > \sigma_E S, \tag{1.19}$$

$$\rho > -\frac{\delta - \lambda \eta_\theta s - \sigma_I \eta_\theta S}{2} + \sqrt{\left(\frac{\delta - \lambda \eta_\theta s - \sigma_I \eta_\theta S}{2}\right)^2 + (\lambda \sigma_E + \sigma_I)\,\delta \eta_\theta S}. \tag{1.20}$$

The first condition (1.19) ensures that spillovers to entrants are sufficiently small that an

incumbent's quality $q_\theta(t)$ grows faster than the entrant quality $q_\theta^E(t)$.[22] This assumption implies that relative qualities $z_\theta(t)$ are weakly above one, which is essential to obtain a non-degenerate stationary distribution H_θ^*. The second condition (1.20) ensures that the discount rate ρ is large enough that the consumer's discounted utility is finite in equilibrium. Together with the lower bound on the exit rate (1.14), it is also sufficient to ensure that firm values are finite.

Given these assumptions, the first proposition characterizes the economy's stable BGPs:

Proposition 1. *The economy has two locally stable BGPs, one for each technology θ. In the BGP for technology θ:*

1. *All scientists research technology θ.*

2. *The knowledge stock K_θ and aggregate quality Q_θ grow at rate*

$$g_\theta^* = -\frac{\delta - \lambda\eta_\theta s - \sigma_I \eta_\theta S}{2} + \sqrt{\left(\frac{\delta - \lambda\eta_\theta s - \sigma_I \eta_\theta S}{2}\right)^2 + (\lambda\sigma_E + \sigma_I)\,\delta\eta_\theta S}.$$

3. *The stationary distribution H_θ^* is a generalized Pareto distribution with location parameter 1, shape parameter $\varphi_\theta^* > 0$, and tail parameter $\xi_\theta^* \in (-\infty, 1)$:*

$$H_\theta^*(z_\theta) = \begin{cases} N\left[1 - \left(1 + \xi_\theta^* \frac{z_\theta - 1}{\varphi_\theta^*}\right)^{-\frac{1}{\xi_\theta^*}}\right] & \xi_\theta^* \neq 0, \\ N\left[1 - \exp\left(-\frac{z_\theta - 1}{\varphi_\theta^*}\right)\right] & \xi_\theta^* = 0. \end{cases}$$

   The parameters satisfy

$$\varphi_\theta^* = \frac{\left(\lambda + \frac{\sigma_I}{\sigma_E}\right)\eta_\theta s - g_\theta^*}{\delta} \quad \text{and} \quad \xi_\theta^* = \frac{\lambda\eta_\theta s - g_\theta^*}{\delta}.$$

Along the BGP for technology θ, all aggregates including the knowledge stock K_θ, aggregate quality Q_θ, total output Y, and consumption C grow at the common rate $g_\theta^* > 0$. This growth rate is naturally increasing in the extent of knowledge spillovers (σ_E, σ_I) and the productivity of research (λ, η_θ). With fully endogenous growth, the model also features a scale effect whereby the growth rate is increasing in the total mass of scientists S. The growth rate is declining in the exit rate δ because all innovation is undertaken by incumbents.

The first core result of the theory concerns the concentration of scientists s. Provided that

[22] A condition is needed because $q_\theta(t)$ grows with innovation by a single firm while $q_\theta^E(t) = \sigma_E K_\theta(t)$ grows with spillovers from all firms.

incumbents build to some extent on their own past advances ($\lambda > 0$), the growth rate strictly increases with the concentration of scientists. This observation clarifies a role for *the firm* per se in the growth process: If scientists build on the knowledge generated by others, and if these spillovers are particularly intense between scientists at the same firm, then larger firms can better catalyze spillovers and accelerate innovation. Any of the microfoundations discussed in Section 1.3.1 could account for this "theory of the firm," and Proposition 1 draws a sharp implication for balanced growth.

The second part of Proposition 1 shows that the BGP relative quality distribution falls into the familiar generalized Pareto class. The shape and tail parameters of the distribution depend endogenously on the growth rate g_θ^* and the parameters of the innovation process. To understand this relationship, note that an intermediate's relative quality $z_\theta(t)$ increases as the growth rate of its quality outpaces the aggregate growth rate g_θ^*:

$$\frac{\dot{z}_\theta(t)}{z_\theta(t)} = \frac{\dot{q}_\theta(t)}{q_\theta(t)} - \frac{\dot{K}_\theta(t)}{K_\theta(t)} = \left[\lambda + \frac{\sigma_I}{\sigma_E} z_\theta(t)^{-1}\right] \eta_\theta s - g_\theta^*.$$

This calculation follows directly from the evolution equation (1.8) for $q_\theta(t)$. Proposition 1 shows that the numerator of the shape parameter φ_θ^* is equal to the growth rate of relative quality for an entrant ($z_\theta(t) = 1$). The numerator of the tail parameter ξ_θ^* is instead the limiting growth rate for a long-lived incumbent ($z_\theta(t) \to \infty$). An increase in either parameter reallocates mass toward larger relative qualities; the shape parameter primarily affects the "body" of the distribution (low z_θ), while the tail parameter naturally affects the tail (high z_θ). Hence any change in the model primitives that raises the growth rate of a firm relative to the growth rate of the economy will produce a more skewed distribution. Because it plays a key role in the analysis of technology choice in Section 1.4, the following corollary formally shows that this holds for the concentration of scientists s:

Corollary 1. *The shape parameter φ_θ^* is strictly increasing in s. The tail parameter ξ_θ^* is strictly increasing in s if and only if $\lambda > 0$.*

This result is not immediate, because an increase in s raises both the growth rate of a firm $\frac{\dot{q}_\theta(t)}{q_\theta(t)}$ and the economy's growth rate g_θ^* when $\lambda > 0$. The first effect is larger because aggregate growth is generated by firm quality growth *in combination with* entry and exit, which attenuates the link between aggregate and firm-level growth. I prove additional comparative statics for φ_θ^* and ξ_θ^* in Appendix A.1.

1.4 Equilibrium with Technology Choice

In this section, I analyze the economy's equilibrium when firms can choose to innovate for technology A or technology B. I suppose that initially only A is available, and that incumbents have innovated for A before B arrives at $t = 0$. These technologies pose a simple trade-off: B has the higher basic research productivity $\eta_B > \eta_A$, so it can support faster growth in the long run. But its initial knowledge stock $K_B(0) > 0$ may be lower than that for A, limiting growth in the short run. Initial incumbents at $t = 0$ are endowed with the entrant quality $q_B^E(0)$, so that no firm has any absolute advantage in innovating B to start. To study the effect of market structure on the aggregate technology choice, I suppose an arbitrary initial distribution $H_{A0}(z_A)$ of relative qualities for A. I later endogenize it using the BGP distribution H_A^* from Proposition 1 to explore how the concentration of scientists s affects the equilibrium.

I begin by characterizing the solution to the firm's problem (1.11) in Section 1.4.1. Firm innovation decisions are generally *path-dependent* and *forward-looking*, and knowledge spillovers generate two kinds of strategic complementarities that critically shape (and complicate) these decisions. To resolve the resulting technical challenges, I study the class of *monotone equilibria* in Section 1.4.2. These equilibria exist under weak conditions, and they display the pattern of innovation observed in the mRNA case study (Section 1.2) and the empirical analysis (Section 1.5): Initial incumbents are reluctant to innovate for the new technology B, but they may begin doing so after entrants have made substantial progress. The initial market structure H_{A0} plays a critical role, as incumbents with greater expertise for A are more reluctant to innovate for B. In Section 1.4.3 I discuss the benchmark case with no knowledge spillovers across incumbents ($\sigma_I = 0$), which features a unique, closed-form equilibrium. When the economy is initially following the BGP for A ($H_{A0} = H_A^*$), I show that the concentration of scientists s has an ambiguous effect on the equilibrium direction of innovation. But it tends to slow or prevent a transition to B provided the discount rate ρ is sufficiently high. The social planner may not always transition to B *more often* than in equilibrium, but when it is optimal, the social planner always transitions *more quickly*.

1.4.1 The Firm's Problem

Consider a firm with qualities q at time $t \geq 0$. Given a trajectory for the knowledge stocks $[K(\tau)]_\tau$, the firm's problem (1.11) is to allocate its scientists at all times $\tau \geq t$ to maximize the present value of its profits. To understand the underlying incentives, let

$$k_\theta(t) \equiv \lambda q_\theta + \sigma_I K_\theta(t). \tag{1.21}$$

denote the total knowledge of technology θ available to the firm. The firm's marginal value of research (scientists) for technology θ at t then satisfies

$$\frac{dV(q,t)}{ds_\theta(q,t)} = \frac{\partial V(q,t)}{\partial q_\theta} k_\theta(t)\eta_\theta, \tag{1.22}$$

To interpret, note that an additional scientist at t raises the quality q_θ by $[\lambda q_\theta + \sigma_I K_\theta(t)]\eta_\theta$, increasing the firm's value in proportion to $\partial V(q,t)/\partial q_\theta$. The marginal value of research $dV(q,t)/ds_\theta(q,t)$ plays a key role in the firm's problem, because any solution must allocate scientists only to the technology with the larger marginal value at each t. The next proposition characterizes this marginal value by showing how the firm's value $V(q,t)$ varies with its qualities q. This result clarifies the path-dependent and forward-looking forces as well as the impact of spillovers on the firm's innovation decisions.

Proposition 2. *The incumbent's marginal value of quality for technology θ is*

$$\frac{\partial V(q,t)}{\partial q_\theta} = \bar{\pi}\Psi_\theta(t),$$

$$\textit{where} \qquad \Psi_\theta(t) \equiv \int_t^\infty \exp\left(-\int_t^{t'} [\rho + \delta - \lambda\eta_\theta s_\theta(q(\tau),\tau)]\, d\tau\right) dt'.$$

The marginal value of research for θ is then

$$\frac{dV(q,t)}{ds_\theta(q,t)} = \bar{\pi}\Psi_\theta(t)[\lambda q_\theta + \sigma_I K_\theta(t)]\eta_\theta.$$

When $\lambda > 0$, this marginal value is strictly increasing in the current quality q_θ and the mass of scientists $s_\theta(q(\tau),\tau)$ researching θ at future times $\tau > t$. When $\sigma_I > 0$, the marginal value is strictly increasing in the current knowledge stock $K_\theta(t)$. It does not depend on technology $\theta' \neq \theta$.

The expression for the marginal value of quality $\partial V(q,t)/\partial q_\theta$ follows by integrating the evolution equation (1.8) for $q_\theta(t)$ and substituting into the firm's objective (1.11). Given an increase in the quality q_θ, the function $\Psi_\theta(t)$ aggregates the additional profits at each time $t' \geq t$, discounted back to t. The discount rate at $\tau \geq t$ is $\rho + \delta - \lambda\eta_\theta s_\theta(q(\tau),\tau)$, where the last term reflects the additional profits generated because the innovation raises the productivity of future research for technology θ. The comparative statics for the marginal value of research $dV(q,t)/ds_\theta(q,t)$ in Proposition 2 are immediate from the corresponding expression.

Proposition 2 yields the second core result of the theory: When knowledge accumulates within firms ($\lambda > 0$), innovation decisions are both *path-dependent* and *forward-looking* at the firm level. Path dependence arises because a firm's past innovation for technology θ raises its

current research productivity for θ, incentivizing continued innovation in that direction. This is immediate from Proposition 2, because the marginal value of research $dV(q,t)/ds_\theta(q,t)$ is strictly increasing in the current quality q_θ when $\lambda > 0$. But history is not always destiny, because the firm also takes into account how its current innovation affects its future research productivities. This is also apparent from Proposition 2, because with $\lambda > 0$ the marginal value of research $dV(q,t)/ds_\theta(q,t)$ increases with future research for θ via the aggregation factor $\Psi_\theta(t)$. Holding fixed any effect of spillovers on the firm's direction, this forward-looking force pushes the firm to innovate for B: The future allocation of scientists is also chosen by the firm, and B allows for faster quality *growth* through cumulative innovation ($\eta_B > \eta_A$). A key insight from the model is that the path-dependent and forward-looking forces arise from the same assumptions that knowledge is technology-specific and accumulates within the firm.

With stochastic entry and exit, these features of the innovation process naturally generate heterogeneity in innovation incentives across firms. Through their past innovation, initial incumbents have higher initial equalities $q_A(0)$ for technology A than entrants, but they share the same initial qualities $q_B^E(0)$ for technology B. This gives initial incumbents an endogenous comparative advantage *in innovating A*, incentivizing them to continue with A after B arrives. In this way, the theory readily explains the initial innovation decisions made by firms in the mRNA case study: Entrants opted to explore the potentially promising mRNA technology, while incumbent pharmaceutical firms with existing expertise "opted to invest [their] resources elsewhere" (Dolgin, 2021).

Knowledge accumulation within firms is crucial for these results. When firms instead draw only on aggregate knowledge ($\lambda = 0$), the marginal value of research simplifies to

$$\frac{dV(q,t)}{ds_\theta(q,t)} = \frac{\bar{\pi}}{\rho+\delta}\eta_\theta \sigma_I K_\theta(t).$$

Each firm allocates its scientists to the technology with the larger marginal value at each t, so in this case firm heterogeneity has no impact on innovation decisions. The assumption of purely public knowledge accumulation is pervasive in the endogenous growth and directed innovation literatures, precluding any meaningful role for the firm (or the concentration of scientists s) in the growth process. See Appendix A.3 for a detailed discussion of this case.

Proposition 2 is also suggestive of the third core result of the model, developed in greater detail in Sections 1.4.2 and 1.4.3: A firm's innovation decisions are influenced by spillovers from other firms, providing a mechanism by which the decisions made by initial incumbents can persistently affect the aggregate direction of innovation. The model features two sources of spillovers. First, the initial qualities of an entrant are determined by the aggregate knowledge

stocks, $q_\theta^E(t) = \sigma_E K_\theta(t)$. When firm innovation decisions are path-dependent ($\lambda > 0$), past innovation for technology θ then raises the entrant's marginal value of research for θ. These "backward-looking" spillovers generate a strategic complementarity between past and current firms. A second source of spillovers arises when $\sigma_I > 0$, so that incumbents draw on public knowledge when innovating. In this case, a firm has incentives to innovate for the technology with the larger current knowledge stock $K_\theta(t)$, generating a strategic complementarity between past and current firms and across all current firms. If in addition $\lambda > 0$, so that the firm's choice of direction is forward-looking, then the firm must account for the future path of the knowledge stocks: The expectation of high growth in $K_\theta(t)$ may induce the firm to innovate for technology θ, even if its current research productivity is lower.

The analysis above provides insight into firm innovation incentives, but it does not yield a sharp characterization of innovation decisions. Two difficulties remain. First, when $\lambda > 0$, the firm's marginal value of research for a given technology θ is increasing in its past and future research for θ. These complementarities render the firm's problem non-convex, so first-order conditions are not sufficient to characterize the solution.[23] Second, knowledge spillovers across incumbents can complicate innovation decisions and introduce multiple equilibria. For example, an incumbent innovating for θ may switch to θ' if the other knowledge stock $K_{\theta'}(t)$ is growing (or expected to grow) quickly enough. Moreover, with knowledge spillovers across incumbents ($\sigma_I > 0$) and knowledge accumulation within incumbents ($\lambda > 0$), incumbents are more likely to innovate for a technology θ if they expect its knowledge stock $K_\theta(t)$ to grow in the future. Since the knowledge stock grows more rapidly as more firms innovate for θ, these spillovers can produce multiple equilibria.

To address the first difficulty, I reformulate the firm's problem (1.11) as an optimal stopping problem. As noted above, at each time the firm optimally allocates all scientists to the technology θ with the larger marginal value of research. As a result, I can cast the firm's problem as the choice of an initial innovation direction $\theta_0 \in \{A, B\}$ and a sequence of stopping times $t < T_1 \leq T_2 \leq \ldots$ at which the firm completely reverses its direction. The next lemma records this observation and provides a first-order necessary condition – the *smooth-pasting condition* – for interior stopping times.[24]

Lemma 2. *Given initial qualities q and a trajectory for the knowledge stocks $[K(\tau)]_\tau$, for any solution to the firm's problem (1.11) there exists an initial innovation direction $\theta_0 \in \{A, B\}$ and a sequence of stopping times $t < T_1 \leq T_2 \leq \ldots$ such that the firm exclusively innovates for θ_0 at*

[23] For example, it can be optimal for a firm to innovate for A at t assuming it exclusively innovates for A in the future, but it might raise value even further to switch all current and future innovation to B.

[24] I adopt the terminology of Dixit (1993) for this optimality condition, though I do not emphasize its connection to the differentiability of the firm's value function at the stopping times.

$t' \in [t, T_1)$ *and reverses its innovation direction at each stopping time* T_j. *Every interior stopping time* $T_j \in (t, \infty)$ *must satisfy the smooth-pasting condition*

$$\Psi_B(T_j)k_B(T_j)\eta_B = \Psi_A(T_j)k_A(T_j)\eta_A. \tag{1.23}$$

The smooth-pasting condition (1.23) simply requires that the marginal value of research must be equalized across technologies at each stopping time T_k. Lemma 2 provides a much more tractable description of the firm's problem than the infinite-dimensional optimal control problem (1.11). But it does not address the second difficulty discussed above: With spillovers across incumbents $\sigma_I > 0$, the firm may reverse its research direction repeatedly, and the economy may feature multiple equilibria. In the next section, I define and characterize *monotone equilibria* that rule out repeated reversals of a firm's innovation direction over time.

1.4.2 Monotone Equilibria

When innovation decisions are both inherently forward-looking ($\lambda > 0$) and subject to knowledge spillovers across incumbents ($\sigma_I > 0$), the model features a dynamic strategic complementarity that renders a full characterization of all equilibria intractable. To simplify while maintaining the key economics, I make two restrictions. First, I consider only equilibria that converge to one of the economy's stable BGPs. This restriction excludes cyclical equilibria as well as equilibria that happen to converge to the economy's unstable BGP. Second, I consider only equilibria in which innovation decisions are appropriately monotone over time:

Definition 3. An equilibrium is *monotone* if the equilibrium allocation of scientists $s_\theta(q(t), t)$ is monotone in $t \geq 0$ for every possible starting time $t_0 \geq 0$ for a firm. Here $q(t)$ evolves according to (1.8) for $t > t_0$. If $t_0 > 0$, then $q(t) = q^E(t)$ for $t \leq t_0$.

To understand the definition, first set $t_0 = 0$. Monotonicity then implies that an initial incumbent reverses its innovation direction at most once. By Lemma 2, this reduces the incumbent's problem to a choice of an initial research direction and a *single* stopping time T. Instead setting $t_0 > 0$, monotonicity implies the same restriction on entrants, with one additional requirement: If at any time $t > 0$ the initial innovation direction for entrants reverses from θ to θ', all subsequent entrants permanently innovate for θ'. In this sense, monotonicity requires that innovation is monotone both within firms and across entrants over time.

I focus on monotone equilibria to maintain analytical tractability while allowing for knowledge spillovers across incumbents ($\sigma_I > 0$), which can generate realistic innovation dynamics. For example, we will see in Section 1.4.3 that in the absence of these spillovers firms never reverse

their initial innovation directions. Spillovers to entrants still generate linkages across firms, but they cannot rationalize why, for example, incumbent pharmaceutical firms are now researching and developing new mRNA therapies. By contrast, spillovers across incumbents can induce a firm to alter its direction of innovation. Monotonicity simply ensures that this occurs at most once for each firm, and so works as a joint restriction on σ_I and the trajectories of the knowledge stocks $[K(t)]_t$. The empirical analysis in Section 1.5 also suggests that these spillovers are important to explain firm patenting behavior, and it shows that incumbent innovation for many new technologies is generally monotone in the sense of Definition 3.

To understand the structure of monotone equilibria, consider one that converges to the BGP for technology B. Each initial incumbent chooses an innovation direction θ_0 and a stopping time $T \in (0, \infty]$ at which to *permanently* reverse it. Since the equilibrium converges to the BGP for B, it can be shown that any reversal must be from A to B. An initial incumbent then either permanently innovates for B or begins innovating for A before switching permanently to B (if at all). These firms differ only by their initial relative qualities $z_A(0)$ for A, so path dependence implies a cutoff value $z_{A0} \geq 1$ such that an incumbent starts innovating for A if and only if $z_A(0) > z_{A0}$.

For entrants, there exists a time $T_E \in [0, \infty]$ after which all entrants innovate permanently for one of the technologies; since the equilibrium converges to the BGP for B, this must be technology B. Any entrant that arrives before T_E then begins by innovating for A, but it also chooses a stopping time $T \in (T_E, \infty]$ at which to permanently switch to B. To simplify, I note that an equilibrium with $T_E = 0$ exists whenever an equilibrium with $T_E > 0$ exists — the latter is essentially a "translation" of a monotone equilibrium with $T_E = 0$, with all firms innovating for A until $T_E > 0$. I restrict to $T_E = 0$ in what follows.

To characterize the stopping times T, note that each such time must satisfy the smooth-pasting condition (1.23), which in this case can be written

$$\frac{k_B(T)\eta_B}{\rho + \delta - \lambda\eta_B s} = \frac{k_A(T)\eta_A}{\rho + \delta}. \tag{1.24}$$

Intuitively, at the stopping time T the value of the firm's initial research for B must be equal to the value of its final research for A. Given the non-convexities in the firm's problem described in Section 1.4.1, the smooth-pasting condition (1.24) alone does not fully characterize a firm's stopping time T. But it does provided that we include the associated second-order condition, that the left side of (1.24) be weakly increasing relative to the right side just before T.

Thus a monotone equilibrium converging to B is determined by the initial cutoff z_{A0}, with all

stopping times T characterized by the smooth-pasting condition (1.24) and the associated second-order condition. The next proposition summarizes this description. It additionally demonstrates that innovation decisions can be written in terms of a *cutoff function* $\chi(t)$ for the *firm knowledge ratio* $\frac{k_B(t)}{k_A(t)}$, such that a firm innovates for B at t if and only if $\frac{k_B(t)}{k_A(t)} > \chi(t)$. The initial value $\chi(0)$ is related to the relative quality cutoff z_{A0} by the identity

$$\chi(0) = \frac{\lambda\sigma_E + \sigma_I}{\lambda\sigma_E z_{A0} + \sigma_I}\kappa(0),$$

where I define the *aggregate knowledge ratio*

$$\kappa(t) \equiv \frac{K_B(t)}{K_A(t)}. \tag{1.25}$$

This ratio measures the "gap" between the two technologies, and it plays a key role in the analysis below. The proposition also provides a monotone comparative static for the cutoff function $\chi(t)$ with respect to the initial relative quality distribution H_{A0}. I state the analogous result for monotone equilibria converging to the BGP for technology A in Appendix A.2.

Proposition 3. *In any monotone equilibrium converging to the BGP for technology B, there exists a cutoff function $\chi(t)$ such that a firm innovates for B if and only if*

$$\frac{k_B(t)}{k_A(t)} > \chi(t).$$

The cutoff function satisfies the differential equation (A.8), and the monotone equilibrium is unique up to the initial value $\chi(0)$. The knowledge stocks $K(t)$ solve the dynamical system (A.13).

The initial value $\chi(0)$ is increasing in H_{A0} in the sense of first-order stochastic dominance, and strictly so whenever $\chi(0) \in (0, \kappa(0))$ and $\sigma_I > 0$.

Figure 1.3 depicts equilibrium dynamics for initial incumbents, both when σ_I is low and when σ_I is high. These cases differ because incumbents may reverse their initial innovation directions if and only if σ_I is sufficiently high.[25] In both cases, equilibrium is characterized by a simple cutoff function $\chi(t)$ for the firm knowledge ratio $\frac{k_B(t)}{k_A(t)}$, reflecting the dynamic sorting of firms to technologies based on comparative advantage in innovation. The cutoff $\chi(t)$ decreases with the knowledge ratio $\frac{k_B(t)}{k_A(t)}$ of the least-experienced incumbent to innovate for A. It remains constant when it reaches the level identified by the smooth pasting condition (1.23) as the value of the knowledge ratio $\frac{k_B(t)}{k_A(t)}$ at which a firm innovating for A reverses to B.

[25] For an equilibrium converging to the BGP for B, an incumbent researching technology A must eventually reverse to B if and only if $g_B^* > \lambda\eta_A s$. This inequality is satisfied if and only if σ_I is sufficiently high.

Figure 1.3: Equilibrium Converging to B

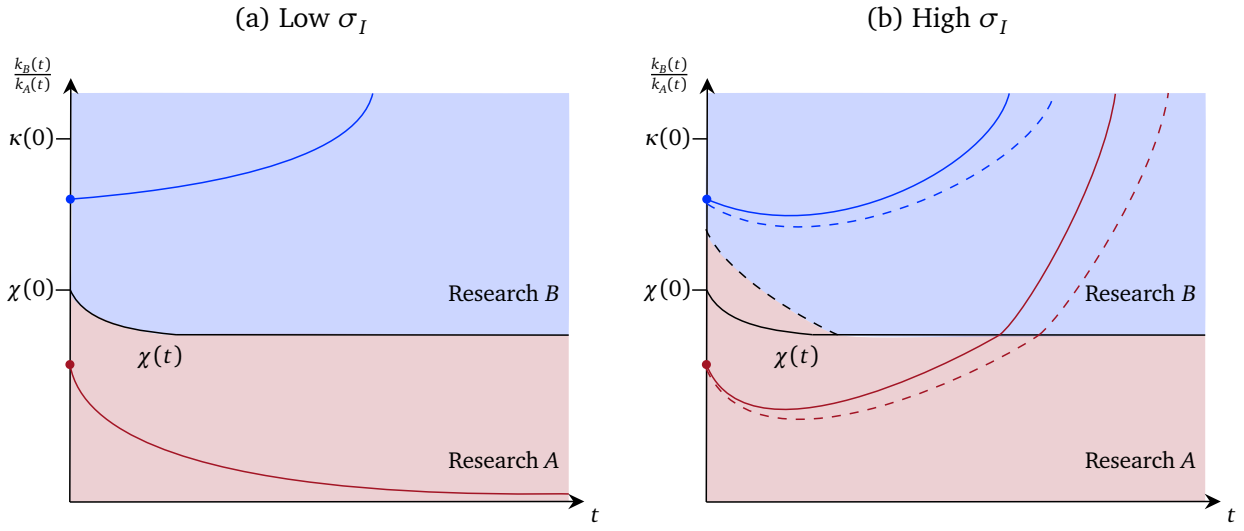

Notes: In each figure, the red and blue regions indicate the sets of knowledge ratios at which a firm would research technology A and B, respectively. Similarly, the solid red and blue lines are the relative knowledge stock trajectories for incumbents that initially research technologies A and B, respectively. The black line is the trajectory of the cutoff $\chi(t)$, which always limits to the value $\frac{\eta_A}{\eta_B}\frac{\rho+\delta-\lambda\eta_B s}{\rho+\delta}$. The dashed lines in Figure 1.3(b) are counterfactual trajectories after an "increase" in the initial distribution H_{A0} in the sense of first-order stochastic dominance.

Proposition 3 also demonstrates that both the cutoff $\chi(t)$ and the knowledge stocks $K(t)$ can be recovered by integrating a second-order dynamical system, dramatically simplifying the computation of equilibria. Moreover, monotone equilibria are unique up to the initial condition $\chi(0)$. The key substantive implication of the Proposition concerns the comparative static with respect to H_{A0}: When $\sigma_I > 0$, the entire cutoff function $\chi(0)$ is increasing in the initial relative quality distribution H_{A0}, so that greater initial experience for technology A always slows the transition to technology B. This provides a first indication of how the economy's initial market structure H_{A0} can persistently affect aggregate innovation.

As the discussion before Proposition 3 clarifies, monotonicity disciplines equilibrium by ensuring that firms do not reverse their research directions more than once. This may not hold for an arbitrary equilibrium converging to the BGP for technology B, because the knowledge stock $K_B(t)$ for technology B may initially decline relative to the knowledge stock $K_A(t)$ for technology A. With strong enough spillovers across incumbents, these dynamics could induce an incumbent to start innovating for B, reverse to A as $\kappa(t)$ falls, and reverse back to B as $\kappa(t)$ diverges asymptotically. To rule out this effect, in Appendix A.2 I derive a condition on primitives and the trajectory of the knowledge stocks $[K(t)]_t$ to ensure that any firm innovating for B can never satisfy the optimality conditions necessary for a reversal to A. The following proposition provides a weaker sufficient condition, exclusively involving primitives, under which monotone

Figure 1.4: Equilibrium Thresholds

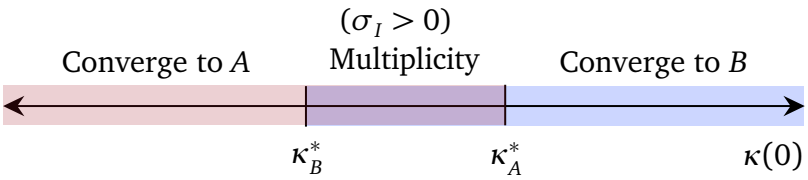

equilibria converging to each technology's BGP exist.

Proposition 4. *Suppose*

$$\lambda \eta_B s \geq \frac{\sigma_I}{\lambda \sigma_E + \sigma_I} g_A^*.$$

Then there exist thresholds $\kappa_B^* \leq \kappa_A^*$ such that:

1. *A monotone equilibrium converging to B exists if and only if $\kappa(0) \geq \kappa_B^*$.*
2. *A monotone equilibrium converging to A exists if and only if $\kappa(0) \leq \kappa_A^*$.*

The thresholds satisfy $\kappa_B^* < \kappa_A^*$ if and only if $\lambda > 0$ and $\sigma_I > 0$. When $\lambda > 0$, the thresholds are strictly increasing in H_{A0}.

This result shows that an equilibrium converging to the BGP for technology θ can only exist provided that the initial advantage for θ' is not too large. Moreover, multiple equilibria can arise for intermediate values of the aggregate knowledge ratio $\kappa(t)$ when the firm's problem (1.11) is dynamic ($\lambda > 0$) and there are knowledge spillovers across incumbents ($\sigma_I > 0$). As the discussion above indicates, these multiple equilibria are supported by different expectations about other firms' future innovation decisions. Figure 1.4 depicts the economy's equilibrium set given the initial knowledge stock ratio $\kappa(0)$.

This proposition also establishes the third core result of the theory: The economy's initial market structure H_{A0} can have a permanent impact on the economy's aggregate technology choice. Greater initial experience makes it easier to sustain an equilibrium converging to A but harder to sustain an equilibrium converging to B. This holds because firm-level path dependence implies that a greater fraction of initial incumbents will choose to innovate for technology A initially, reducing the knowledge ratio $\kappa(t)$ and raising the incentives for later entrants to innovate technology A too.

1.4.3 Benchmark: $\sigma_I = 0$

To develop additional intuition, I consider the benchmark case with no spillovers across incumbents ($\sigma_I = 0$). This case provides a simple illustration of equilibrium dynamics, and it provides several new insights about how market structure and the concentration of R&D affect the aggregate direction of innovation. I sketch the equilibrium derivations below, describing firm innovation decisions, their aggregation, comparative dynamics for the concentration of scientists s, and potential inefficiencies. I provide full details in Appendix A.3.

Firm Innovation. The case with $\sigma_I = 0$ is particularly tractable because all spillovers between firms are "backward-looking," so each firm is unaffected by other firms' current and future innovation decisions. This conveniently implies that each firm chooses its initial innovation direction with no reversals.

To characterize innovation decisions, let $V^\theta(q,t)$ denote the value from innovating permanently for technology θ given initial qualities q at time t. Integrating the quality evolution equation (1.8) yields the explicit formula

$$V^\theta(q,t) = \frac{q_\theta + q_{\theta'}}{\rho+\delta} + \frac{1}{\rho+\delta}\frac{\lambda\eta_\theta s}{\rho+\delta-\lambda\eta_\theta s}q_\theta.$$

The first term is the value of the firm's initial qualities q, while the second term is the value of subsequent innovation for θ. The firm innovates for B if an only if B yields a higher value, $V^B(q,t) \geq V^A(q,t)$. Substituting the expression above, this holds if and only if

$$\frac{\eta_B}{\rho+\delta-\lambda\eta_B s}q_B \geq \frac{\eta_A}{\rho+\delta-\lambda\eta_A s}q_A. \tag{1.26}$$

We can apply this observation to conveniently characterize entrant and incumbent innovation decisions. First consider an entrant at $t \geq 0$. Substituting the entrant qualities $q^E_\theta(t) = \sigma_E K_\theta(t)$ into the inequality (1.26), we can rearrange to find that it innovates for technology B when the aggregate knowledge ratio $\kappa(t)$ is above the *entry threshold*

$$\kappa^E \equiv \frac{\eta_A}{\rho+\delta-\lambda\eta_A s}\left(\frac{\eta_B}{\rho+\delta-\lambda\eta_B s}\right)^{-1} \in (0,1).$$

Now consider an initial incumbent at $t=0$ with qualities $q_A(0) = z_A(0)q^E_A(0)$ and $q_B(0) = q^E_B(0)$. Substituting these qualities into (1.26), we similarly find that the incumbent permanently

innovates for technology B when its relative quality $z_A(0)$ is below the cutoff

$$z^*_{A0} \equiv \frac{\kappa(0)}{\kappa^E}$$

This cutoff is inversely related to the initial value of the cutoff function $\chi(t)$ described by Proposition 3, which in this case satisfies $\chi(0) = \kappa^E$. That is, the equilibrium in this case is monotone, and with $\chi(0)$ pinned down it must be unique. Since firms never reverse their innovation directions in equilibrium, there is no need to track the full path of the cutoff function $\chi(t)$, and I can focus on z^*_{A0} in the analysis below.

Aggregation. Using this characterization of innovation decisions, we can directly derive the dynamical system described by Proposition 3 for the evolution of the knowledge stocks K. Suppose $\kappa(0) > \kappa^E$, and consider technology A. For t near zero, all innovation for A is conducted by initial incumbents that chose not to innovate for B. The quality q_A for any such incumbent grows at rate $\lambda\eta_A s$, but these incumbents exit at rate δ. Differentiating the evolution equation (1.9) for K_A then yields

$$\ddot{K}_A(t) = -(\delta - \lambda\eta_A s)\dot{K}_A(t). \tag{1.27}$$

With the assumed lower bound on the exit rate (1.14), this equation implies that the knowledge stock K_A increases more slowly over time as initial incumbents exit and entrants instead choose to innovate for technology B. The initial conditions for this differential equation are the initial knowledge stock $K_A(0)$ and the initial growth rate $\dot{K}_A(0)/K_A(0)$, which is decreasing in the cutoff z^*_{A0}:

$$\frac{\dot{K}_A(0)}{K_A(0)} = \lambda\eta_A s\sigma_E \int_{z^*_{A0}}^{\infty} z_A dH_{A0}(z_A). \tag{1.28}$$

Now consider technology B. For t near zero, entrants as well as initial incumbents with $z_A(0) \leq z^*_{A0}$ innovate for technology B. The quality q_B for any such firm grows at rate $\lambda\eta_B s$, but firms exit at rate δ. Exiting firms are replaced by a mass δN of entrants with initial quality $\sigma_E K_B(t)$, so differentiating the evolution equation (1.9) for K_B yields

$$\ddot{K}_B(t) = -(\delta - \lambda\eta_B s)\dot{K}_B(t) + \lambda\eta_B s\delta N\sigma_E K_B(t). \tag{1.29}$$

The initial conditions for this differential equation are the initial knowledge stock $K_B(0)$ and

the initial growth rate $\dot{K}_B(0)$, which is increasing in the cutoff z^*_{A0}:

$$\frac{\dot{K}_B(0)}{K_B(0)} = \lambda \eta_B s \sigma_E H_{A0}(z^*_{A0}). \tag{1.30}$$

The equations (1.27, 1.29) form a second-order, autonomous, linear system of differential equations for the evolution of the key state variables K (while entrants innovate for B). This is a natural kind of dynamical system to describe the evolution of an economy undergoing an endogenous technological transition: As C. I. Jones (1995) discusses in detail, all models of endogenous *balanced* growth rely on a *first*-order, autonomous, linear differential equation of the form $\dot{X}(t) = gX(t)$ to generate exponential growth, where X is an appropriately-defined knowledge stock or productivity variable.[26] But along an endogenous technological *transition*, the growth rates of the old and new technologies must adjust to accommodate the rise of the new technology. A *second*-order linear system is perhaps the simplest way to describe the evolution of state variables with smoothly changing growth rates.

I show in Appendix A.3 that the system (1.27, 1.29) can be integrated in closed form. The solution has two key implications: The aggregate knowledge ratio $\kappa(t)$ depends on initial conditions only through its initial value $\kappa(0)$, in which it is strictly increasing; and $\kappa(t)$ is "U-shaped" over time.[27] These properties are useful because the system (1.27, 1.29) only describes the dynamics of the knowledge stocks $K(t)$ while entrants continue to research technology B, $\kappa(t) \geq \kappa^E$. If this condition is ever violated, the economy fails to transition in aggregate to technology B and instead converges back to the BGP for technology A. Thus there exists a threshold κ^* such that the economy converges to the BGP for technology B if and only if $\kappa(0) \geq \kappa^*$. Note that this threshold characterization of the aggregate direction of innovation reflects Proposition 4 for the case in which the equilibrium is unique.

Equilibrium Transition. The threshold κ^* determines the economy's propensity to transition to technology B in equilibrium, and it depends richly on model primitives. To gain intuition for the key forces, let $H_{A0} = H^*_A$, and note that the analysis above implies that a simple sufficient condition for the economy to transition is for the initial growth rate of K_B to dominate the initial growth rate of K_A. By (1.28, 1.30), this holds if and only if

$$\eta_B H^*_A\left(z^*_{A0}\right) \geq \eta_A \int_{z^*_{A0}}^{\infty} z_A dH^*_A(z_A). \tag{1.31}$$

[26] For example, the present model satisfies this relation for $X = K_\theta$ along each stable BGP.

[27] The latter holds because exit and entry are sluggish, so that technology A can continue improving relative to technology B until enough entrants have begun innovating for B.

Figure 1.5: Concentration Comparative Statics

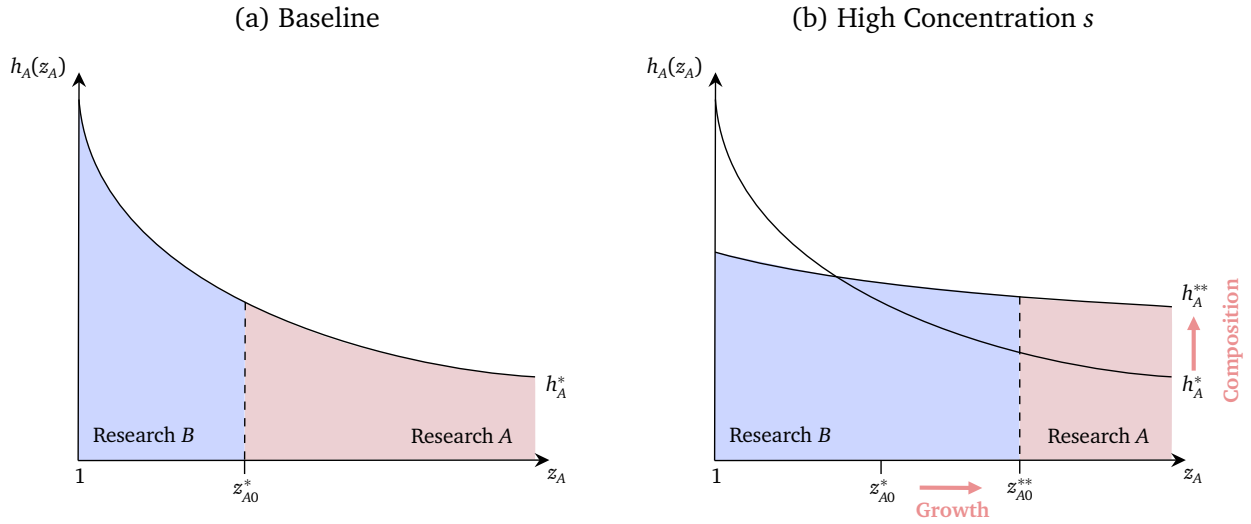

Notes: Figure 1.5(a) shows the BGP relative quality density h_A^* when technology B is introduced. The cutoff z_{A0}^* separates initial incumbents that research technology A (red area) from initial incumbents that research technology B (blue area). Figure 1.5(b) how the composition and growth effects alter the masses of initial incumbents that research each technology.

This inequality depends on the initial aggregate knowledge ratio $\kappa(0)$ only through the cutoff z_{A0}^*, which is strictly increasing in $\kappa(0)$. As $\kappa(0)$ rises, the left side increases as a larger mass of initial incumbents innovate for B, raising the initial growth rate of K_B. The right side instead decreases as fewer initial incumbents innovate for A, lowering the initial growth rate of K_A. There exists a unique value $\bar{\kappa}$ at which the inequality (1.31) binds:

$$1 = \frac{\eta_A}{\eta_B} \frac{1}{1-\xi_A^*} \frac{\frac{\bar{\kappa}}{\kappa^E}}{\left(\frac{\bar{\kappa}}{\kappa^E}\right)^{1/\xi_A^*} - 1}. \tag{1.32}$$

Here I have substituted $z_{A0}^* = \frac{\kappa(0)}{\kappa^E}$ and the BGP relative quality distribution H_A^* from Proposition 1, which is a standard Pareto distribution when $\sigma_I = 0$.

The right side of equation (1.32) is strictly increasing in ξ_A^* and κ^E and strictly decreasing in $\bar{\kappa}$, and it delivers a simple but powerful intuition about the drivers of a technological transition: Any change that thickens the tail of the old technology's firm-quality distribution slows the transition, because it raises both the relative mass of incumbents who choose not to transition and their initial innovation rates. Both effects increase their collective influence over the aggregate direction of innovation, which may be decisive if it induces entrants to switch back to innovating for technology A. However, any change that raises incentives for new firms to

innovate for the new technology instead accelerates the transition by raising the relative mass of incumbents who choose to transition. The tail parameter ξ_A^* and the entry threshold κ^E respectively capture these "composition" and "growth" forces, depicted in Figure 1.5. But they depend on many of the same model primitives. The following proposition provides an explicit comparative static for $\bar{\kappa}$ with respect to the concentration of scientists s:

Proposition 5. *There exists a discount rate $\bar{\rho} \geq 0$ that depends on model primitives such that $\bar{\kappa}$ is strictly increasing (decreasing) in s locally if and only if ρ is larger (smaller) than $\bar{\rho}$.*

The trade-off between the composition and growth effects hinges on the discount rate ρ. When ρ is relatively high, firms are relatively myopic, and the increase in the growth rate of technology B with s has little effect on firm incentives. The composition effect then dominates, and the increase in the concentration of scientists can delay or prevent a transition to technology B. In Section 1.6, I provide a simple calibrated example to show that an increase in the concentration of scientists can raise the thresholds $\bar{\kappa}$ and κ^* for reasonable values of ρ.

Efficiency. To show that this may be inefficient, I suppose that a social planner instead chooses the initial cutoff z_{A0} to maximize the representative consumer's utility (1.1) given the initial knowledge stocks $K(0)$, with all innovation decisions by entrants as in equilibrium. The following proposition characterizes properties of the solution $\hat{z}_{A0}$:

Proposition 6. *A solution $\hat{z}_{A0}$ to the social planner's problem exists and depends on $K(0)$ only through $\kappa(0)$. There exists a threshold $\hat{\kappa}$ such that*

1. *the solution $\hat{z}_{A0}$ yields a transition to technology B if and only if $\kappa(0) \geq \hat{\kappa}$;*
2. *$\hat{z}_{A0} > z_{A0}^*$ if $\kappa(0) \geq \hat{\kappa}$; and*
3. *$\hat{z}_{A0} \leq z_{A0}^*$ if $\kappa(0) < \hat{\kappa}$, with equality only if $z_{A0}^* = 1$.*

In general, the transition thresholds for the social planner $\hat{\kappa}$ and the equilibrium κ^* cannot be ranked. This holds because the social planner internalizes knowledge spillovers on future entrants when choosing the long-run direction of innovation, but these spillovers are not necessarily always larger for a given technology: Technology B spillovers are larger in the long-run given $\eta_B > \eta_A$, but technology A spillovers may be larger in the short-run given incumbents' initial expertise for technology A (i.e., the initial distribution H_A^*). However, Proposition 6 shows that for a given long-run innovation direction, the social planner always prefers to direct greater initial innovation in that direction than in equilibrium.

1.5 Empirical Analysis

In this section, I present an empirical analysis of firm innovation decisions using patent publications. I document three facts consistent with the theory: First, a firm's current patenting is highly correlated with its previous patenting, controlling for other determinants of innovation outcomes (Section 1.5.2). Second, for the collection of new technologies identified by Kalyani et al. (2023), a firm's current patenting for a given technology is better predicted by past patenting *within that technology* than patenting in general (Section 1.5.3). Third, incumbents with greater patenting experience patent *less* for a new technology than less-experienced firms, though this gap shrinks as the technology matures (Section 1.5.4). The first two facts support the theory's key assumptions that knowledge is both cumulative within firms and technology-specific, generalizing existing evidence from the auto industry (Aghion et al., 2016). The third supports the key prediction that experienced incumbents are reluctant to innovate for new technologies given their expertise in old ones. I discuss these results in Section 1.5.5, with additional results and robustness checks in Appendix A.5.

1.5.1 Data

The main dataset for the analysis is a panel of US public firms over 1980-2021, which includes measures of (i) each firm's patenting, both overall and within specific technologies; (ii) the aggregate stock of knowledge (patents) available to each firm, both overall and within specific technologies; and (iii) each firm's R&D expenditures. Below I summarize the main data sources, with additional details in Appendix A.4.

Patents. I use data on the set of all utility patents filed at the US Patent and Trademark Office (USPTO) after 1980 and granted through 2023. This dataset covers almost seven million patents and is made available through the USPTO's PatentsView platform. It includes each patent's title, abstract, assignees (initial owners), inventors, technology area, and citations made to other USPTO patents.

Patent-Firm Matching. Accurately grouping patents by firm is difficult, because assignee names are not standardized and are not adjusted to reflect changes in firm ownership through mergers and acquisitions.[28] I address these issues using the DISCERN 2.0 dataset, which employs an extensive matching process to identify all patents granted to US public firms from

[28] PatentsView provides disambiguated assignee and inventor names meant to resolve the first issue, but my experience matching patents to pharmaceutical companies for the case study in Section 1.2 suggests that the PatentsView disambiguation protocol is still a work-in-progress.

1980 to 2021 (Arora et al., 2024). The resulting 1,865,633 patents are matched to 5,680 firms, and for tractability I restrict the analysis to the top 10% of firms by total number of patents. The final subset includes 1,659,998 patents matched to 568 firms. Restricting to public firms reduces the scope of the analysis, but it ensures an accurate matching of patents to firms and is common throughout the empirical literature on firm innovation (Hall et al., 2010).

R&D Expenditures. I use data on yearly nominal R&D expenditures from Compustat North America. I normalize by US GDP per capita each year to measure inflation-adjusted R&D expenditures in "scientist equivalent" units.

New Technologies. To group patents into technologies in a theory-consistent way, I start with the list of new technologies produced by Kalyani et al. (2023). They identify new technologies by extracting all two-word combinations that appear in the text of US patents over 1976-2014, excluding any combination found in a representative sample of pre-1970 American English text. Of the remaining "novel" combinations, the authors retain only the ones that appear in a sufficiently large number of citation-weighted patents. Finally, to ensure that these combinations refer to technologies instead of scientific concepts or problems to solve, the authors search for a Wikipedia page corresponding to each combination and verify that it describes a technology. Each of the resulting 1,148 two-word combinations is a *new technology*.

Kalyani et al. (2023) also define an "emergence year" for each new technology. This is meant to capture the time just before the new technology became prominent, but after it became available for firms to innovate. It is defined as the first year in which the number of citation-weighted patents that mention the technology (i) reaches a minimal threshold and (ii) grows by ten percent over each of the next five years. I use this emergence year in my analysis to mark the arrival of each new technology.

I identify a set of patents related to each new technology by searching for the two-word combination in the title and abstract of each patent. I include any patent that mentions the combination, as well as any patent that cites such a patent. Table 1.1 lists the top ten new technologies by total number of patents, along with the number of citation-weighted patents and the emergence year. The list primarily includes information and communication technologies ("data store," "code-division multiple access") along with biomedical technologies ("thermometer," "ion channel") and industrial technologies ("heat treating"), representative of the broader set of 1,148 technologies. Across this broader set, on average each technology has 13204.1 raw patents, 27804.1 citation-weighted patents, and emergence year 1991.3.

How well do these groups of patents capture the notion of a technology in the theory? The latter

Table 1.1: Top New Technologies

Technology	Patents (Raw)	Patents (Citation-Weighted)	Emergence Year
data store	133992	241300	1990
codedivision multiple access	107638	144829	1986
memory address	107358	163577	1992
thermometer	106336	214837	1991
heat treating	104587	154475	1983
microsoft access	100097	203024	1987
error detection correction	96819	146012	1986
holographic optical element	95231	138734	1992
textbased user interface	94379	197112	1994
ion channel	94097	144825	1992

corresponds to (i) a shared set of production techniques and (ii) a shared base of knowledge used to improve those techniques. As Table 1.1 suggests, the new technologies identified in the patent data generally meet the first condition: Each new technology spans patents with similar applications, from data storage ("data store") to biological targets for drug treatment ("ion channel"). They also meet the second condition essentially by design: Any patent associated with a new technology must either mention it *or cite* one that does, so that it contributes directly to or builds on the technology's base of knowledge. To see that patents for a new technology are generally focused on similar topics, we can analyze how they are classified into technology areas by patent examiners. Each patent is assigned to potentially several of 132 "classes" and 672 "subclasses" defined in the Cooperative Patent Classification (CPC). These groupings are meant to identify all patents with similar technological content, with greater specificity for subclasses than classes. On average, I find that a share 0.58 of patents in a new technology belong to a *single class*, while a share 0.45 belong to a *single subclass*, so that patents associated with a new technology are highly concentrated within technology areas.[29]

This method to group patents into technologies has several limitations. First, identifying an initial set of patents by keywords may include patents that simply mention the keywords but actually address other applications. I mitigate this concern by searching only in the title and abstract of each patent, which are more tightly focused on the patent's core applications. Second, Kalyani et al. (2023) explicitly identify *new* technologies that arrive over 1976-2014, but my

[29]I do not use the CPC hierarchy to directly define technologies for two reasons. First, classes and subclasses are still too broad to provide a compelling definition: Examples include class A61 "Medical or Veterinary Science; Hygiene" and subclass A61K "Preparations for Medical, Dental, or Toiletry Purposes", which include patents with a variety of applications. Second, many reasonably-defined technologies include patents assigned to many different classes and subclasses. For example, the two most common subclasses among mRNA therapy patents from Section 1.2 are A61K and C12N "Microorganisms or Enzymes; Compositions Thereof; Propogating, Preserving, or Maintaining Microorganisms; Mutation or Genetic Engineering; Culture Media."

theory focuses on the contrast between new and old technologies. In my empirical analysis, I proxy for old technologies with the set of all patents not associated with a new technology that were filed before its emergence year. Finally, some of the new technologies identified by Kalyani et al. (2023) have substantial overlap, such as "data store" and "microsoft access" (a database management system). This will attenuate my estimates of the technology specificity of knowledge, to the extent that I misclassify patents as unrelated to a given new technology. However, it may also artificially increase precision by inflating the number of new technologies above its true value.

1.5.2 Fact 1: Knowledge Accumulation within Firms

I first provide evidence for knowledge accumulation within firms: A firm's current patenting is highly correlated with its previous patenting, holding fixed R&D expenditures, previous patenting by all firms, and the firm's latent propensity to patent. The regression is

$$Pat_{it} = \exp\left(\beta_1 \log\left(K_{it}^{\text{Firm}}\right) + \beta_2 \log\left(K_{it}^{\text{Agg}}\right) + \beta_3 \log\left(s_{it}\right) + \alpha' X_{it} + \varepsilon_{it}\right). \tag{1.33}$$

Here Pat_{it} denotes the number of patents filed by firm i in year t. I weight each patent by the number of forward citations it receives to better reflect the value of the underlying innovations.[30] The firm-level knowledge stock K_{it}^{Firm} measures firm i's accumulated internal knowledge at time t. I construct this stock from past patent flows by the perpetual inventory method, setting the depreciation rate ν to the standard value of 0.15 (e.g., Hall et al., 2005):

$$K_{it}^{\text{Firm}} = (1-\nu)K_{it-1}^{\text{Firm}} + Pat_{it-1}.$$

The aggregate knowledge stock K_{it}^{Agg} measures the aggregate knowledge available to firm i in year t. Since each firm draws on different knowledge depending on its area of focus, I construct K_{it}^{Agg} in two steps. I first compute the knowledge stocks for each CPC subclass by the perpetual inventory method. I then compute K_{it}^{Agg} as a weighted average of these stocks in each year t, with weights given by the distribution of firm i's total patents across all CPC subclasses. Finally, I control for R&D expenditures s_{it} and a vector of controls X_{it} that includes separate fixed effects by year, firm, and the number of years for which the firm has been publicly listed.

The regression (1.33) is an empirical analogue to the quality evolution equation (1.8). I use the patent flow Pat_{it} in place of the change in quality $\dot{q}_\theta(t)$, the firm-level knowledge stock K_{it}^{Firm} in place of the quality $q_\theta(t)$, and K_{it}^{Agg} in place of the aggregate knowledge stock $K_\theta(t)$. However,

[30] See Appendix A.4 for additional details about the variable constructions.

in this section the regression (1.33) is not technology-specific; it instead captures determinants of overall innovation. I adopt a log-linear functional form so that the coefficients of interest β_1, β_2, and β_3 may be interpreted as unit-free elasticities. The hypothesis that knowledge accumulates within firms implies $\beta_1 > 0$, so that a firm's past patenting has a larger impact on its current patenting than predicted by its effect on the aggregate knowledge stock K_{it}^{Agg}.

I estimate the regression (1.33) on the panel of public firms described in Section 1.5.1. To mitigate truncation issues in the patent data and in the unbalanced panel of firms, I restrict to the years 1985-2016, and I drop any observations for the first year when a firm becomes publicly listed.[31] I estimate directly by negative binomial regression and by linear regression (OLS) with dependent variable $\log(1 + Pat_{it})$. The corresponding estimates are found in columns (1) and (2) of Table 1.2.

In each specification, I find that a firm's patenting is positively correlated with its past patenting (K_{it}^{Firm}) conditional on the other covariates. The estimates of the corresponding elasticity suggest that a 10% increase in a firm's internal knowledge stock K_{it}^{Firm} is associated with a 7% increase in the firm's contemporaneous patenting Pat_{it}. This is consistent with the assumption of the theory that knowledge accumulates within firms ($\lambda > 0$), so that the productivity of a firm's R&D expenditures s_{it} is increasing in its own past patenting K_{it}^{Firm}. The regression controls for many alternative explanations: For example, firm fixed effects ensure that this relationship is not driven by variation in firms' latent propensity to patent at all times, which would yield a positive correlation between past and current patenting. Controlling for R&D expenditures similarly ensures that the relationship is not driven by variation in the scale of each firm's R&D program. To the extent that R&D expenditures are elastic to *persistent* innovation opportunities, this control also rules out the explanation that firms simply patent more when they discover a valuable line of research.

Table 1.2 also shows that a firm's patenting is positively correlated with the measure of the aggregate knowledge K_{it}^{Agg} available to the firm, which reflects positive spillovers across incumbents in the theory ($\sigma_I > 0$). The estimates of the corresponding elasticity imply that a 10% increase in the firm's external knowledge K_{it}^{Agg} is associated with a 2.5-3% increase in contemporaneous patenting. I also find a positive elasticity of patenting to R&D expenditures of 8-9%, in line with previous estimates in the literature (e.g., Bloom et al., 2013b).

[31] Truncation issues arise in the patent data because I only observe *granted* patents, and there is often a multi-year lag between when a patent is filed and when it is granted. A similar problem arises for observed citations, which I address by normalizing forward citations both within year and across CPC subclasses (see Appendix A.4). A left truncation issue appears for firms because I only observe them after they become publicly listed, biasing my initial estimates of their internal knowledge stocks K_{it}^{Firm} downward if they filed for patents before then.

Table 1.2: Firm Patenting Regressions

	Overall Patenting		**Technology Patenting**	
	(1) Neg. Bin.	(2) OLS	(3) Neg. Bin.	(4) OLS
$\log(K_{it}^{\text{Firm}})$	0.6962***	0.7161***	0.3557***	0.0714***
	(0.0330)	(0.0335)	(0.0605)	(0.0145)
$\log(K_{it}^{\text{Agg}})$	0.2973**	0.2481**	-0.4386*	-0.1386***
	(0.1249)	(0.1187)	(0.2641)	(0.0530)
$\log(s_{it})$	0.0778***	0.0951***	0.0877**	0.0211***
	(0.0186)	(0.0203)	(0.0431)	(0.0079)
$\log(K_{i\theta t}^{\text{Firm}})$			0.5156***	0.3338***
			(0.0140)	(0.0228)
$\log(K_{\theta t}^{\text{Agg}})$			0.2093***	0.0416***
			(0.0330)	(0.0046)
Observations	12,237	12,237	562,080	564,062
Firms	555	555	435	453

*Significance: ***: 0.01, **: 0.05, *: 0.1*

Notes: The overall patenting regressions include fixed effects by year, firm, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero internal knowledge stock K_{it}^{Firm}, zero aggregate knowledge stock K_{it}^{Agg}, and zero R&D expenditures s_{it}. The technology-specific patenting regressions also include fixed effects by technology and the time since emergence and dummy variables for zero internal and aggregate technology-specific knowledge stocks $K_{i\theta t}^{\text{Firm}}$ and $K_{\theta t}^{\text{Agg}}$. All standard errors are clustered at the firm level.

1.5.3 Fact 2: Technology-Specific Knowledge

The evidence above supports the claim that knowledge accumulates within firms, but they say nothing about the extent to which this knowledge is *technology-specific*. To assess this second assumption of the theory, I consider a related regression for a firm's patenting in a given technology θ:

$$\begin{aligned} Pat_{i\theta t} = \exp(&\beta_1 \log\left(K_{it}^{\text{Firm}}\right) + \beta_2 \log\left(K_{it}^{\text{Agg}}\right) + \beta_3 \log\left(s_{it}\right) \\ &+ \beta_4 \log\left(K_{i\theta t}^{\text{Firm}}\right) + \beta_5 \log\left(K_{\theta t}^{\text{Agg}}\right) + \alpha' X_{i\theta t} + \varepsilon_{i\theta t}). \end{aligned} \tag{1.34}$$

Here the dependent variable $Pat_{i\theta t}$ is the number of citation-weighted patents filed by firm i for technology θ in year t. Regressors again include the firm's internal knowledge stock K_{it}^{Firm}, aggregate knowledge stock K_{it}^{Agg}, and R&D expenditures s_{it}, none of which are technology-

specific. But I also add the technology-specific knowledge stocks $K^{\text{Firm}}_{i\theta t}$ and $K^{\text{Agg}}_{\theta t}$. Here $K^{\text{Firm}}_{i\theta t}$ is defined analogously to K^{Firm}_{it} as the discounted sum of past citation-weighted patents by firm i for technology θ. The aggregate knowledge stock $K^{\text{Agg}}_{\theta t}$ is simply the discounted sum of all past patents for technology θ and does not vary across firms. The vector of controls $X_{i\theta t}$ includes the same fixed effects as in the first regression (1.33), in addition to fixed effects by technology θ and the number of years after the emergence of the technology. The hypothesis that knowledge is technology-specific implies $\beta_4 > \beta_1$ and $\beta_5 > \beta_2$, so that knowledge related to a given technology θ has a larger impact on a firm's current patenting for θ than generic knowledge.

I consider technologies θ with emergence years after 1990. I make a final restriction to ensure that the estimates of the regression (1.34) apply to firm-technology pairs such that the firm could plausibly innovate the technology: As described above for firms, for each technology θ I compute the distribution of its associated patents across CPC subclasses. I then exclude any firm-technology pairs for which the cosine similarity between the firm and technology patent distributions falls below 0.6. This excludes, for example, the pair of Microsoft and "microsatellite" (cosine similarity 0.44) while including the pair of Microsoft and "flash memory" (cosine similarity 0.61). Including more firm-technology pairs attenuates the coefficient estimates, as firms patent little for technologies outside their areas of focus.

The negative binomial and OLS estimates for the technology-specific regression (1.34) are found in columns (3) and (4) of Table 1.2. These estimates broadly demonstrate that knowledge is technology-specific: The estimated elasticities corresponding to the technology-specific knowledge stocks $K^{\text{Firm}}_{i\theta t}$ and $K^{\text{Agg}}_{\theta t}$ are significantly positive and generally larger than the elasticities corresponding to the generic knowledge stocks K^{Firm}_{it} and K^{Agg}_{it}. For example, a 10% increase in a firm's past patenting for a technology θ correlates with a 4-5% increase in the firm's contemporaneous patenting for θ, while the corresponding elasticity for the firm's generic past patenting is only 1.3-3%. Finally, note that the elasticity of patenting for technology θ with respect to R&D expenditures naturally attenuates relative to the firm-level specification (1.33): s_{it} is at best an imprecise proxy for the R&D expenditures specifically allocated by the firm to technology θ, which I cannot observe.

1.5.4 Fact 3: Incumbents and New Technologies

When knowledge accumulates within firms and is technology-specific, the theory predicts that incumbent firms with extensive experience in existing technologies should be reluctant to innovate new ones, relative to less-experienced incumbents or entrants (Proposition 2). To test this prediction, I estimate several regressions that relate a firm's patenting for a new technology

θ to the characteristics of the firm at the emergence time T_θ:

$$\log(1+Pat_{i\theta}^{10}) = \beta_1 \log\left(K_{iT_\theta}^{\text{Firm}}\right) + \beta_2 \log\left(K_{iT_\theta}^{\text{Agg}}\right) + \beta_3 \log\left(s_{iT_\theta}\right) + \beta_4 \log\left(K_{i\theta T_\theta}^{\text{Firm}}\right) + \beta_5 \log\left(K_{\theta T_\theta}^{\text{Agg}}\right) + \alpha' X_{i\theta} + \varepsilon_{i\theta}. \tag{1.35}$$

Here $Pat_{i\theta}^{10}$ denotes the total number of citation-weighted patents filed by firm i for technology θ in the ten years after the technology's emergence. The regressors include all firm-specific and aggregate knowledge stocks at the emergence time T_θ, along with the firm's R&D expenditures s_{iT_θ}. The vector of controls $X_{i\theta}$ includes fixed effects by firm i, emergence time T_θ, and the number of years for which the firm has been publicly listed at the emergence time.[32] I also consider the alternative specification in which the outcome variable is the *share* of technology θ in all patents filed by firm i in the ten years after emergence:

$$TechShare_{i\theta}^{10} \equiv \frac{Pat_{i\theta}^{10}}{\sum_{t=T_\theta+1}^{T_\theta+10} Pat_{it}}.$$

The regression (1.35) is similar to the previous one with technology-specific patenting (1.34), but with no panel dimension. It simply relates a firm's medium-run innovation for a new technology $Pat_{i\theta}^{10}$ to the firm's characteristics when that technology emerges. The coefficient of interest is β_1, which measures how the firm's initial knowledge stock at the time of emergence $K_{iT_\theta}^{\text{Firm}}$ predicts its subsequent patenting in the new technology θ. When the dependent variable is $TechShare_{i\theta}^{10}$, the theory predicts that this coefficient should be negative: The empirical results documented thus far suggest that firm innovation decisions should be to some extent path-dependent. As a result, incumbents with greater internal knowledge stocks when a new technology θ emerges should devote a smaller share of their R&D resources to technology θ, yielding in a smaller share for technology θ in the firm's total patents $TechShare_{i\theta}^{10}$. This need not imply that these firms patent *absolutely* less for technology θ, because their greater research productivity may allow them to patent more in many different technologies.[33] Only when the path dependence force is particularly strong does the theory predict that β_1 should be negative when the dependent variable in the regression (1.35) is the transformed patent count $\log(1+Pat_{i\theta}^{10})$.

I estimate the regression (1.35) by OLS for both dependent variables. For brevity, I display

[32] In this baseline specification, I exclude technology fixed effects so as to estimate the effect of the initial aggregate knowledge stock $K_{\theta T_\theta}^{\text{Agg}}$. The results in Table 1.3 are robust to the inclusion of technology fixed effects (see Appendix A.5).

[33] This force is absent in the theoretical model because all knowledge is technology-specific and firms cannot deploy more R&D resources as their research becomes more productive.

Table 1.3: Technology Patenting Level Difference

	Technology Patents		Technology Patent Share	
	(1) Full Sample	(2) No Early Patents	(3) Full Sample	(4) No Early Patents
$\log(K^{\text{Firm}}_{iT_\theta})$	-0.0694**	-0.0848***	-0.0023***	-0.0016**
	(0.0345)	(0.0273)	(0.0007)	(0.0007)

*Signifance: ***: 0.01, **: 0.05, *: 0.1*

Notes: All regressions include fixed effects by firm, emergence year, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. All standard errors are clustered at the firm level.

the estimates for β_1 in Table 1.3 and present the estimates for the remaining coefficients in Appendix A.5. Columns (1) and (3) provide the estimate $\hat{\beta}_1$ for the full sample of firms, while columns (2) and (4) exclude firms with any patenting for technology θ before the emergence year T_θ. Consistent with firm-level path dependence, the estimates suggest a sizable negative effect of a firm's initial patenting experience on its subsequent patenting for new technologies: A 10% increase in a firm's initial internal knowledge stock $K^{\text{Firm}}_{iT_\theta}$ is associated with a 0.8% decrease in the firm's subsequent patents for the new technology. It also corresponds to a decrease of 0.01-0.02 percentage points in the technology's share of the firm's subsequent patents. To provide a sense of the effect size, note that the 20th and 80th percentiles of the distribution of the initial internal knowledge stock $K^{\text{Firm}}_{iT_\theta}$ are 2.75 and 6.47, respectively. Moving from the 20th percentile to the 80th percentile of the initial knowledge distribution reduces the firm's subsequent patents for a new technology by 23%. It also reduces the technology's share of the firm's subsequent patents by 0.0086, or approximately 59% of the mean share in the sample (0.0145). These effects suggest substantial path dependence in incumbent innovation, consistent with a fundamental prediction of the theory.

To assess how path dependence varies as a new technology matures, I estimate a dynamic version of the regression (1.35), allowing the effect of the initial internal knowledge stock $K^{\text{Firm}}_{iT_\theta}$ to vary over time:

$$\begin{aligned}\log(1+Pat_{i\theta w}) = &\sum_{\tilde{w}=-2}^{6} \beta_{1\tilde{w}} \log\left(K^{\text{Firm}}_{iT_\theta}\right) \mathbb{1}_{w=\tilde{w}} + \beta_2 \log\left(K^{\text{Agg}}_{iT_\theta}\right) + \beta_3 \log\left(s_{iT_\theta}\right) \\ &+ \beta_4 \log\left(K^{\text{Firm}}_{i\theta T_\theta}\right) + \beta_5 \log\left(K^{\text{Agg}}_{\theta T_\theta}\right) + \alpha' X_{i\theta w} + \varepsilon_{i\theta w}.\end{aligned} \tag{1.36}$$

Here I group all years from 1980-2022 into three-year windows, and for each technology I

Figure 1.6: Technology Patenting Level Difference, Dynamic

(a) Patents: $\log(1 + Pat_{i\theta w})$ (b) Patent Share: $TechShare_{i\theta w}$

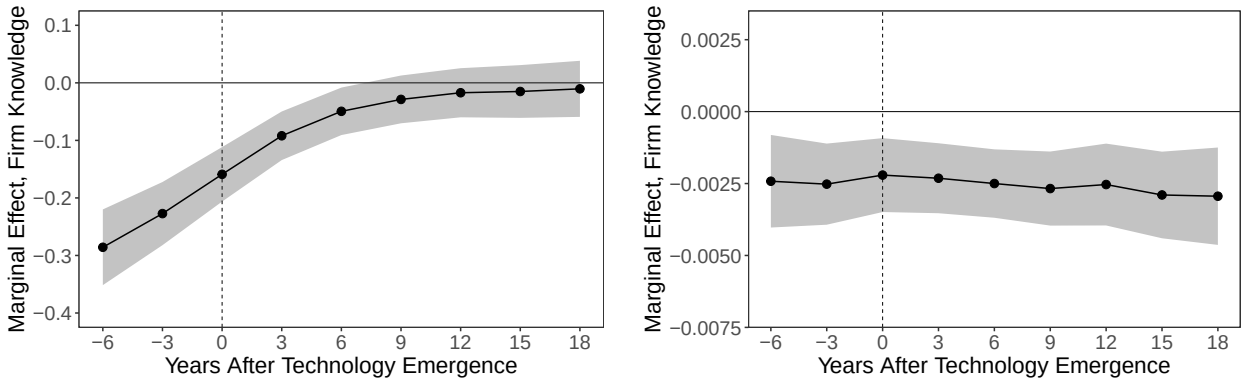

Notes: All regressions include fixed effects by firm-year, technology-window, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. The gray area denotes the 95% confidence interval, with all standard errors are clustered at the firm level.

denote the window containing the emergence time T_θ by $w = 0$. I let $Pat_{i\theta w}$ denote the firm's total citation-weighted patents for technology θ in window w. I group years into windows so that I can consider the alternative specification in which the outcome variable is the share of technology θ in all patents filed by firm i in window w:

$$TechShare_{i\theta w} \equiv \frac{Pat_{i\theta w}}{Pat_{iw}}.$$

The timing of patent filings is noisy from year to year, so grouping years into windows reduces the number of observations dropped when the denominator of $TechShare_{i\theta w}$ is zero.

The regression (1.36) relates the firm's patenting after the emergence of a technology to its various knowledge stocks at the time of emergence, allowing for heterogeneous effects across time for the firm's generic internal knowledge stock $K^{\text{Firm}}_{iT_\theta}$. The vector of controls $X_{i\theta w}$ now includes fixed effects for each firm-year it, each technology-window θw, and the number of years for which the firm has been publicly listed at the time of emergence T_θ. Note that I do *not* include firm-technology fixed effects, so the regression (1.36) does *not* correspond to an event study. Rather, the coefficients of interest β_{1w} capture the *level differences* in patenting across firms with different internal knowledge stocks $K^{\text{Firm}}_{iT_\theta}$ at the emergence time. The key prediction of the theory is that these coefficients should be negative $\beta_{1w} < 0$ for windows w near zero (i.e., near emergence).

I estimate the regression (1.36) by OLS for the dependent variables $\log(1 + Pat_{i\theta w})$ and $TechShare_{i\theta w}$. The corresponding estimates ("marginal effects") $\hat{\beta}_{1w}$ are plotted in Figure 1.6. These estimates provide clear evidence for firm-level path dependence: A firm with greater initial patenting experience patents strictly less for an emerging technology, both absolutely and relative to its total flow of patents.[34] Figure 1.6(a) shows that this negative effect on total patenting attenuates over time. Experienced incumbents are particularly unlikely to innovate for an emerging technology, but after 10 years a firm's initial experience has no effect on its total patenting for the new technology. Figure 1.6(b) replicates the initial negative effect of the firm's internal knowledge stock for the share of the new technology in the firm's total patents. This negative effect is more persistent. On average, an experienced incumbent always has a lower share of patents for a new technology than less-experienced firms.

1.5.5 Discussion

The empirical results presented above support the key assumptions and implications of my theory. The results in Section 1.5.2 provide evidence that firms accumulate knowledge through innovation that raises the productivity of their future R&D. Most microeconomic evidence on the return to R&D at the firm level ignores this dynamic channel (Hall et al., 2010; Bloom et al., 2013b), which my results suggest is quantitatively important. Klette (1996) provides an early discussion of this issue, which I extend with an alternative regression framework for firm-level innovation and more comprehensive evidence about the knowledge accumulation channel.

The results in Section 1.5.3 additionally demonstrate that knowledge produced through innovation is highly technology-specific, indicating that firm innovation decisions are subject to the path-dependent and forward-looking forces highlighted in the theory (Lemma 2). These findings generalize existing work in environmental economics that show knowledge from dirty and clean innovation is technology-specific (Dechezleprêtre et al., 2014) and may generate firm-level path dependence (Aghion et al., 2016). I show that these properties are pervasive, extending to a variety of technologies in many industries. I also explicitly control for other firm-level determinants of innovation outcomes (like R&D expenditures) that could offer an alternative explanation for these findings.

Finally, the analysis in Section 1.5.4 tests the theory's core implication that firms with substantial expertise in existing technologies should be reluctant to innovate for new ones. The findings are consistent with the theory, and they provide novel evidence that experienced incumbents

[34] As expected, the coefficient magnitudes are similar to those from the cross-sectional regressions in Table 1.3.

generally play a limited role in innovating emerging technologies. The pattern of innovation uncovered by this analysis is also consistent with the history of innovation in mRNA vaccines described in Section 1.2: Incumbents initially innovate less for a new technology, but their efforts intensify after younger firms develop the technology over several years. Lessons from the case study appear to generalize to the emergence of a variety of new technologies.

1.6 Quantitative Example

In this section, I provide a simple calibrated example of the model to assess how an increase in the concentration of scientists might affect equilibrium innovation and social welfare. I describe the calibration in Section 1.6.1, and I provide simulation results in Section 1.6.2.

1.6.1 Calibration

I calibrate all parameters of the model necessary to compute the equilibrium path of the knowledge stocks K and the aggregate qualities Q. I first normalize the total mass of scientists to $S = 1$. I set the exit rate to $\delta = 0.075$, in line with recent estimates of the aggregate exit rate for US firms reported by Hopenhayn et al. (2022). I choose the discount rate $\rho = 0.075$ to ensure that the firms' total discount rate $\rho + \delta = 0.15$ falls in the middle of the range reported by Gormsen and Huber (2023).

To calibrate the parameters of the innovation process, I note that the estimates of the firm patenting equation (1.33) reported in Table 1.2 indicate that firms build substantially on their own past knowledge when innovating, with a smaller contribution from public knowledge. To keep the equilibrium dynamics simple, I consider the limiting case described in Section 1.4.3 and set the contribution of spillovers across incumbents to zero, $\sigma_I = 0$. I then normalize $\lambda = 1$. I choose the parameters $\sigma_E = 0.4$ and $\eta_A = 0.026$ jointly to ensure that the initial BGP growth rate is $g_A^* = 0.02$ and that the tail parameter of the BGP firm-quality distribution $\xi_A^* = 0.472$ matches the corresponding value for the US firm-productivity distribution as in Benhabib et al. (2021). I similarly choose technology B's productivity parameter $\eta_B = 0.03$ to ensure that this technology allows for 0.5% faster growth than technology A along its BGP, $g_B^* = 0.025$. Finally, I set the concentration of scientists $s = 2.1$ to deliver a firm-quality growth rate of $\lambda\eta_A s = 0.055$, within the range of net employment growth rates for continuing US firms reported by Haltiwanger et al. (2013). At $t = 0$ the economy is following the BGP for technology A, with corresponding relative quality distribution H_A^*.

Figure 1.7: Knowledge Ratio and Aggregate Quality Trajectories

(a) Knowledge Ratio κ (b) Aggregate Qualities Q_θ

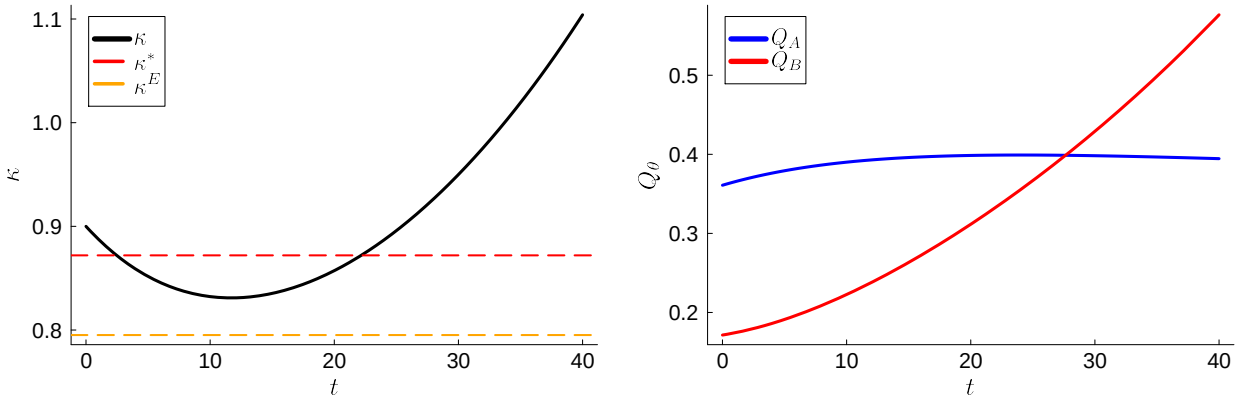

Notes: The equilibrium was simulated with initial knowledge stock ratio $\kappa(0) = 0.9 > \kappa^*$.

1.6.2 Simulation Results

The calibrated model has a unique equilibrium, and it features a transition from technology A to technology B if and only if the aggregate knowledge ratio $\kappa(t)$ is initially above the threshold $\kappa^* \approx 0.872$. Figure 1.7 displays the trajectories of the aggregate knowledge ratio $\kappa(t)$ and the aggregate qualities $Q(t)$ for an equilibrium with a transition. The aggregate knowledge ratio $\kappa(t)$ is "U-shaped," initially decreasing before increasing asymptotically. But it remains above the threshold κ^E at which entrants would instead innovate for technology A, ensuring that the economy converges to the BGP for technology B in the long run. The aggregate quality Q_B for technology B increases with innovation, and its growth rate initially increases as incumbents exit and entrants reallocate their scientists to technology B. The aggregate quality Q_A for technology A initially increases before declining to a positive limiting value due to firm reallocation.

To explore how the concentration of scientists s can affect the equilibrium's propensity to transition, in Figure 1.8 I plot the equilibrium threshold κ^* for different values of s and ρ, keeping all other parameters of the model fixed. The figure indicates that an increase in the concentration of scientists increases the threshold κ^* under the baseline calibration of the model, implying a *lower* propensity to develop the high-growth technology B in equilibrium. The threshold κ^* increases more rapidly with s for higher values of the discount rate ρ, while it can instead decline with s for lower values of ρ. These observations are exactly consistent with the trade-off between the composition and growth effects discussed in Section 1.4.3. When the discount rate ρ is high, firms have short planning horizons, so any increase in future growth from greater concentration of scientists s has a negligible impact on firm innovation decisions.

Figure 1.8: Equilibrium Threshold

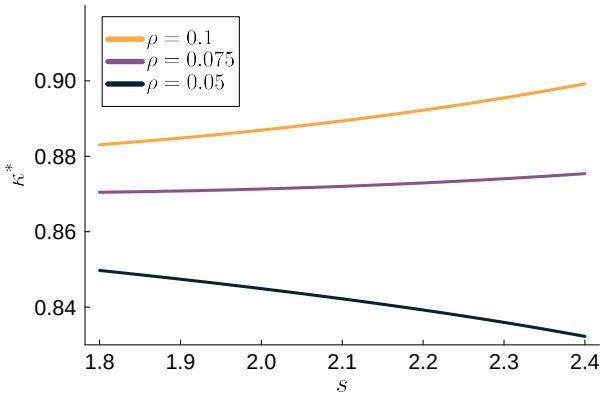

Figure 1.9: Equilibrium and Efficient Cutoffs z_{A0}

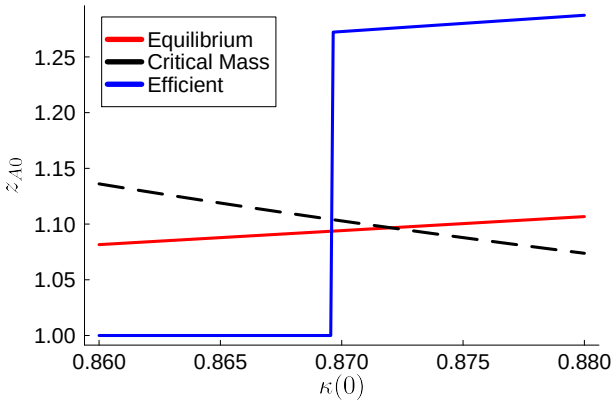

However, it also produces a more skewed firm-quality distribution for technology A, leading a larger mass of initial incumbents to continue innovating for technology A.

Finally, it is straightforward to see that an increase in the concentration of scientists s can actually reduce social welfare. In Figure 1.9, I plot the initial incumbent cutoff z_{A0} chosen in equilibrium (red line) and by the social planner (blue line) as a function of the initial aggregate knowledge stock ratio $\kappa(0)$. This cutoff determines which initial incumbents innovate for technology B ($z_A(0) \leq z_{A0}$), and it must be greater than the dashed "critical mass" line to ensure a transition to the BGP for technology B. Under the baseline calibration, the social planner always prefers to transition more often than in equilibrium. Consistent with Proposition 6 the efficient cutoff $\hat{z}_{A0}$ is always greater than the equilibrium cutoff z^*_{A0} when a transition to technology B is efficient.

An increase in s raises the equilibrium transition threshold κ^*, represented in Figure 1.9 by the intersection between the red and black lines. If $\kappa(0)$ is found to the right of this intersection, the increase in s can prevent a transition to technology B, strictly reducing the economy's long-run growth rate and social welfare. This holds despite the uniformly positive effects of an increase in s on growth and welfare along the economy's BGPs.

1.7 Conclusion

In this paper, I presented a new model of innovation and firm dynamics to clarify a novel connection between market structure, the direction of innovation, and economic growth. The accumulation of knowledge within firms generates an endogenous comparative advantage for incumbents *in innovating* old technologies relative to entrants; the accumulation of public

knowledge instead generates complementarities across firms. These features of the innovation process jointly imply that an industry's initial market structure affects its propensity to explore new technologies and hence its long-run growth. This mechanism has immediate relevance for the growing concentration of R&D in AI: Increasing concentration allows for faster progress for existing applications, but it also risks leaving valuable alternative innovation directions unexplored, as large incumbent firms develop AI according to their existing expertise. Empirical evidence from patenting for emerging technologies and a case study of mRNA vaccines supports the theory's mechanism and its key implication that incumbent firms are slow to innovate new technologies.

The broad implication of this paper is that focusing on complementarities between firms, not competition, provides a new perspective on the role different firms play in the innovation process. This basic idea touches on many aspects of innovation that would be interesting to explore in future research. For example, this paper considered only firms that are small relative to the size of the industry and labor market for scientists. But a "large" (granular) firm may internalize its effect on the innovation incentives for outside researchers, utilizing its size in the labor market to affect the direction of innovation in the industry. Similarly, complementarities in innovation may also explain (and suggest harm from) the investments large technology companies have made in small but innovative start-ups. More generally, these complementarities suggest a new theory of harm from industry consolidation even when competition is not threatened. Innovation or antitrust policies that promote diverse and vibrant innovation ecosystems may be essential for long-run growth.

Chapter 2

Combining Complements: Theory and Evidence from Cancer Treatment Innovation

This chapter was co-authored with Rebekah Dix.

2.1 Introduction

Many of the technologies that shape our world are inherently combinatorial, employing multiple innovative components together to achieve outcomes superior to what each component could accomplish individually. For example, smartphones enable wireless communication and information processing by integrating a radio transceiver, a microprocessor, and a lithium-ion battery, technologies that were individually recognized with Nobel prizes: Marconi and Braun (1909), Alferov et al. (2000), and Goodenough et al. (2019), respectively. Electric vehicles similarly combine many innovations used in gas-powered vehicles, such as suspension systems, aerodynamic designs, and safety features, with more recent breakthroughs in battery technology and electric motors. And leading treatments for medical conditions ranging from ADHD to HIV/AIDS to COVID-19 consist of multiple drugs that are more effective when used together.[1] While combination innovations like these are thought to play a critical role in long-run technological progress (J. Schumpeter, 1934; P. Romer, 1992; Weitzman, 1998), firms' incentives to invent new combinations remain less well examined.

[1]Adderall, approved in 1996, treats ADHD using amphetamine and dextroamphetamine, discovered in the 1920s and 1930s. Pfizer's COVID-19 drug, Paxlovid, combines nirmatrelvir and ritonavir.

In this paper, we study incentives to innovate combinations. Innovation, in general, results in numerous externalities that can distort the efficiency of equilibrium investment (Arrow, 1962). Innovating firms may steal business from existing substitutes, generate positive spillovers for consumers through better products or lower prices, or benefit other firms through information spillovers that improve their innovation productivities.[2] We argue that combination innovation introduces a new externality on other firms, a *market expansion effect* that arises when it increases profits for underlying component products not owned by the innovator.[3] This positive externality reduces the private value of innovation below its social value, potentially leading to underinvestment in combinations. Moreover, component owners may have an incentive to free ride off others' combination innovation, leading to delayed innovation.

To what extent does the market expansion externality drive a wedge between equilibrium and socially optimal innovation of new combinations, and how should policy respond? We develop an empirical framework to answer this question in a healthcare context: the market for cancer drug combination therapies, where combination innovation occurs through running clinical trials for new combinations of cancer drugs.

Cancer combination therapies offer a policy-relevant setting in which to study combination innovation. Cancer is the second leading cause of death in the United States (CDC, 2024), and innovations in chemotherapy have been important contributors to reductions in cancer mortality (Cutler, 2008; Lichtenberg, 2010; Siegel et al., 2023), with combination therapies emerging as some of the most potent weapons in the war on cancer (Mokhtari et al., 2017).[4] Innovating a new combination therapy by demonstrating safety and efficacy in clinical trials is costly, and pharmaceutical companies face substantial legal, logistical, and strategic barriers that prevent them from internalizing externalities through Coasian bargaining (Humphrey et al., 2011; Institute of Medicine et al., 2012; Deng, 2015; Boshuizen and Peeper, 2020; Podkonjak et al., 2021; Sanofi, 2022).[5] Government intervention is common in this market through publicly-funded combination trials (Holbeck et al., 2017; Meric-Bernstam et al., 2023), opening questions about designing cost-effective policies to support combination innovation

[2] Early literature on innovation focused on a single innovation at a time, while the more recent "Schumpeterian" literature (Aghion and Howitt, 1992) emphasizes dynamic models of repeated innovation. Quality ladder models feature business stealing, where each innovation replaces demand for previous generations and enables the development of the next.

[3] Similar effects can arise when an innovation increases profits for complementary products.

[4] Recently, the promise of more effective treatments and the sense that new combination therapies have been slow to arrive have prompted extensive discussion in the medical research community about the causes of delay (Boshuizen and Peeper, 2020).

[5] For example, uncertainty over the success and potential qualities of combination therapies may contribute to imperfect or private information across different parties. Dynamically, drug owners may have an incentive to free ride off of the combination innovation of others. We discuss these forces in detail in Section 2.2.1 and Appendix B.7.

that balance private and public funding (Schilsky, 2013).

We use data from publicly recorded investments in clinical trials to observe combination innovation decisions for both successful and failed innovations, and data on drug demand and prices to estimate the value of successful combinations to innovators and the spillovers on patients and firms. An advantage in studying drug combination innovation is that we can readily identify the risk set of potential innovations (i.e. all possible combinations of existing drugs), something that is typically more challenging to do in empirical studies of innovation.[6] This greatly enhances our ability to characterize the direction of innovation and the nature of "missing innovations."

We formalize externalities arising from combination innovation using a stylized model of cancer drug combination innovation. Firms own one or more individual drugs and can run a clinical trial to, if successful, introduce a treatment regimen that combines multiple drugs. The introduction of a new combination regimen may impose pecuniary externalities on consumers and other firms, leading an innovator investing in combination innovation to not fully internalize the change in industry profits and consumer surplus.

An important case of such externalities is when there are *missing property rights* from combination innovation. This issue is especially stark in our setting of cancer treatment. Unlike many oral medications (e.g., Adderall and Paxlovid[7]), cancer combination therapies are typically not packaged together since component drugs may be taken over the course of many days or weeks in a clinical outpatient setting and in dosages that depend on patient characteristics. Because of this, an innovating firm cannot prevent patients (or doctors) from adopting the new regimen using its underlying drugs. Thus, the innovator of a combination can profit from the innovation only by selling more of the component drugs it owns at their non-discriminatory prices.[8] Similar missing property rights arise in other settings, such as the innovation of software and hardware.[9]

In our setting, the externalities that result from introducing a new drug combination include a market expansion externality distinctive to combination innovation, as well as two externalities

[6] A recent exception is Kim (2023), which studies innovation in structural biology and uses a database of all known proteins to study the direction of innovation.

[7] These drugs are examples of Fixed Dose Combinations, combinations of two or more drugs in a single dosage form. We discuss the differences from our setting in more detail in Section 2.2.1.

[8] The fact that the combination cannot be sold as a physically combined drug also precludes selling at a bundled discount since the component drugs could be resold.

[9] New computer software or smartphone applications increase the demand for manufacturers of the associated hardware (e.g., processors, memory, screens, ...). However, developers may struggle to appropriate all of the surplus because of the competition that they face in their product markets. Conversely, improvements in hardware quality increases the demand for software but the associated positive externality on developers may not be fully internalized.

that are common to many kinds of innovation, consumer surplus spillovers and business stealing from existing regimens. This market expansion externality arises whenever a firm introduces a regimen that uses drugs patented by another firm, raising demand for those complementary drugs and hence their owners' profits. Consumer surplus spillovers are also exacerbated by missing property rights, because a firm cannot price discriminate based on the value of a drug used in a particular regimen. These new features tend to reduce the private value of innovation below its social value.

We use our stylized model to formulate three predictions about private and public innovation decisions that we test using comprehensive data on clinical trials for cancer drugs between 1990 and 2022. First, all else equal (particularly, regimen prices and qualities), market expansion externalities imply that a firm's research portfolio will consist of a smaller share of combination trials than will a social planner's. The data support this prediction: while 58% of cancer drug clinical trials are for combinations rather than single-drug therapies, this percentage is significantly lower for trials sponsored solely by firms (49%) compared to trials sponsored by publicly-funded (i.e., likely more social-welfare minded than firms) innovators (63%).

Second, the missing property right problem implies that firms will trial combinations of their own drugs more often than combinations involving drugs they do not own, therefore directing innovation towards own-drug combinations. Focusing on two-drug combinations, firm trials are biased towards combinations consisting of their own on-patent drugs—11 times more likely than if combinations were drawn uniformly at random from the set of possible combinations—instead of combinations with other firms' on-patent drugs or generic drugs. One might propose an alternative explanation for this result, arguing that firms prefer trialing their own drugs due to greater familiarity or technological compatibility. We present two pieces of evidence against these alternative explanations. First, we show that the probability of a firm trialing its own on-patent drug and a generic drug is largely unaffected by whether the generic drug was initially patent by the firm. Second, we examine data from laboratory tests (Holbeck et al., 2017) of pre-clinical measures of efficacy from all possible two-drug combinations from a set of marketed cancer drugs. These "bench test" measures show that combinations that consist of drugs owned by the same firm or two different firms have similar measures of efficacy across a wide range of cancer cells. Together, these results suggest that financial incentives created by property rights for drugs drive trialing decisions.

Third, the missing property right problem implies that the owner of a drug will run fewer combination trials including that drug after it faces generic competition, while other innovators will run more. Price declines after generic entry may increase demand and thus the value of the combination to other innovators. To examine this prediction, we compare innovation

before and after generic entry. We find that a drug is used in more combination trials on average after generic entry (11.0% increase), with the original owner running fewer (-59.4% decrease) and publicly-funded innovators and other firms running more (7.9% and 41.7% increase, respectively). The increase in total trialing is driven by an increase in the probability of a newly generic drug being trialed with an on-patent drug.

Motivated by these empirical results, which are consistent with inefficient underinvestment in combination therapies because of the market expansion externality and missing property right problem, we develop an empirical framework to quantify the externalities from combination innovation and design innovation funding policies to correct these externalities. This framework consists of two steps. First, we model cancer drug demand and price setting, which allows us to estimate the welfare change from combination innovation and decompose this change into externalities on patients and firms. Second, we develop a dynamic combination innovation model that captures the incentive to free ride off others' combination innovation and the possibility of public crowd-out. We use this model to recover primitives of the combination innovation decision and evaluate counterfactual innovation funding policies.

In this first step, we estimate a model of patient demand for cancer drugs and drug price setting. Each patient makes a joint decision with her doctor over what treatment regimen to take, either a single drug or combination therapy. The demand model captures complementarities between drugs by modeling demand over bundles (regimens) of drugs following Gentzkow (2007), using insights from the medical literature to define the set of recommended regimens (Chu and DeVita, 2019). Substitution patterns depend flexibly on patient demographics and type of insurance (public or private), which is important for accurately measuring the value of new combinations. We estimate the model using claims data for Medicare and privately insured patients, from the Marketscan database, between 1998 and 2019. These microdata allow us to develop a novel method to measure market shares and prices of combination regimens, which cannot be recovered from aggregate drug sales alone. We model price setting as simultaneous Nash bargaining between drug owners and private insurers, which captures the role of insurers as intermediaries in determining drug prices for their relatively price inelastic beneficiaries. We recover marginal costs of drug production and parameters of the bargaining problem. These estimates are important inputs into firm profit functions, including predicting how prices respond to innovation.

We use these models to quantify the externalities from combination therapies that have been successfully trialed and introduced to the US market between 1999 and 2019. This analysis is selected on combination therapies that, despite the market expansion externality, were likely privately profitable to develop. We find that new combination therapies with at least two

branded drugs owned by different firms have large positive spillovers on other firms that own component drugs in the regimen but did not sponsor the trial, so that the sum of business stealing and market expansion is large and positive on average: one year after introduction it averages approximately $27 million, and extrapolated over the average patent length of the component drugs implies positive profit externalities upwards of $200 million over the life-cycle of each new combination therapy. New combination therapies also have positive spillovers on patients ($31 million per year per new combination therapy on average), but there are also large negative externalities from combination innovation on insurers that increase their costs ($24 million per year per new combination therapy on average). These results suggest that firms are often under-incentivized to conduct trials for combination therapies because of large, positive spillovers to other firms and patients.

Finally, in the second step of the empirical framework, we study how these externalities affect the path of combination innovation by developing and estimating a dynamic model of combination therapy innovation decisions. For each combination regimen that is trialed in the data, we model the timing of when that regimen is trialed, and which innovator trials the regimen, as a dynamic discrete-choice game. The game is finite horizon, capturing the fundamental non-stationarity of the setting that occurs through new drug introduction and intellectual property protections. The agents in the game are innovators that have interests in trialing the regimen (i.e., firms that have at least one on-patent drug in the regimen and a publicly-funded innovator). Each innovator maximizes the discounted sum of variable surplus from successfully trialed regimens net of trialing costs. Market expansion externalities create an incentive for firms to free ride off others' combination innovation. And innovation by the public innovators similarly may crowd out private combination innovation.

We estimate the model using a full-solution likelihood-based approach, recovering the fixed cost of innovation and parameters of the public innovator's objective function. To make this estimation procedure computationally tractable, and facilitate the computation of counterfactuals, we reduce the size of the state space using a method similar to partially oblivious equilibrium (Weintraub et al., 2008, Benkard et al., 2015) and moment-based Markov equilibrium (Ifrach and Weintraub, 2016) that creates separability across the games for different regimens. We also apply sieve value function approximation (Arcidiacono et al., 2013), which allows us to approximate each game's solution.

We then use the estimated model to design cost-effective policy solutions to support combination innovation. We focus on combination regimens that were trialed in the data and three types of policies that could affect the speed with which these new regimens were trialed. Specifically, we study the effects of (i) research subsidies, (ii) varying the amount of public innovation, and

(iii) varying the direction of public innovation. These three policies are motivated by existing interventions in this market which take the form of research grants and publicly-funded trials, and our counterfactuals are informative about the design of these interventions. They are also simple to implement compared to the relatively information intensive Pigouvian subsidy.

Increasing public innovation is a cost-effective policy for increasing the rate of combination innovation, even though the model predicts private firms will to some extent free ride off public innovation. Focusing on colorectal cancer, increasing the public trialing budget by approximately $416 million can increase consumer surplus and profits by as much as $840 million, with a total welfare increase of $616 per patient-year. Despite being an untargeted policy, research subsidies can also be a cost-effective way to advance combination innovation as they increase private trialing probabilities while reducing public-crowd out, increasing consumer surplus and profits by $750 million at a cost of $307 million, with a total welfare increase of $642 per patient-year. Finally, redirecting public innovation towards combinations that firms are particularly under-incentivized to trial, that is, combinations with many generic drugs or combinations with high market expansion potential, minimizes public crowd out of private investments and can increase the rate of combination innovation and total welfare while remaining budget neutral. These policies increase total welfare for colorectal cancer regimens by as much as $367 million ($533 per patient-year), giving similar gains to uniformly increasing public innovation or research subsidies but at much lower implementation cost.

Related Literature: Our work contributes first to a vast empirical literature on the efficiency of private innovation decisions (including Griliches, 1979, 1991; Jaffe, 1986; Klette, 1996; C. Jones and Williams, 1998, 2000; Hall et al., 2010; Bloom et al., 2013a; Lucking et al., 2019; Zacchia, 2020; B. Jones and Summers, 2021). Papers in this literature use a variety of methods to measure innovation externalities, and typically estimate that social returns to R&D are higher than private returns. Our study of combination innovation externalities highlights a new force leading to underinvestment. Recent work including Bloom et al. (2013a) focuses in particular on quantifying the knowledge spillovers and business stealing externalities emphasized in the Schumpeterian growth literature (Aghion and Howitt, 1992). We emphasize the distinctive pecuniary externalities that arise under combination innovation, including market expansion and severe consumer surplus spillovers from missing property rights, and we develop a structural model to estimate them in the cancer drug market. To our knowledge, this paper is the first to provide an empirical analysis of combination innovation in a market setting, complementing a primarily theoretical literature on its role in economic growth (J. Schumpeter, 1934; Weitzman, 1998; B. Arthur, 2009; Clancy, 2018; C. Jones, 2023). Relative to these papers, our paper also highlights how market expansion externalities can arise in models with imperfect competition.

Since combination innovations necessarily build on existing technologies, our work is also related to the literature on follow-on innovation (e.g., Kitch, 1977; Green and Scotchmer, 1995; Heller and Eisenberg, 1998; Lerner and Tirole, 2004; Scotchmer, 2004; Williams, 2013; Sampat and Williams, 2019), which typically has not focused on quantifying the externalities that result from follow-on innovation.[10] Our discussion of the missing property rights problem for combination therapies is particularly related to the missing property rights problem for "new uses" for generic drugs discussed in Roin (2013) and Conti et al. (2020), but we emphasize that it also shapes incentives to innovate combination therapies when the underlying drugs are still on-patent.

We also contribute to a literature on the economics of the pharmaceutical industry, including work on pharmaceutical demand and pricing (papers with closely related models include Dubois and Lasio, 2018; Dubois et al., 2022; Maini and Pammolli, 2020; Cuddy, 2021; Dafny et al., 2023; see Scott Morton and Kyle, 2011 for a broader review) and innovation (Budish et al., 2015; Dubois et al., 2015; Gilchrist, 2016; Rao, 2020; Krieger et al., 2022; Agha et al., 2022; Aryal et al., 2022; Frankel et al., 2023). Most related are several papers that study demand for combination therapies and the pricing of the underlying drugs: Song et al. (2017) estimates a model of demand and pricing for combination therapies for colon cancer, and uses it to explore the potential price-reducing effects of mergers between firms with complementary drugs. Cao and Chatterjee (2023) similarly studies the pricing of *fixed-dose* combinations, in which the constituent drugs are bundled in a single dosage form. We instead focus on combination therapies that use separately marketed drugs, which are much more common in cancer treatment and limit firms' ability to price discriminate. We focus on the incentives to innovate combination therapies rather than pricing, and we develop a structural model of combination innovation. Our characterization of the pecuniary externalities from combination innovation is similar to that of Brekke et al. (2023), that develops a model of duopoly pricing in the presence of a combinatorial good. We use our empirical model to quantify these externalities and explore the implications for innovation, and we stress the role of intellectual property protections for individual drugs in shaping innovation incentives. Finally, Wang (2022, 2023) studies the effects of the Medicines Patent Pool on drug diffusion and innovation, focusing on generic firms' production of HIV drug cocktails. Such patent pools have not been widely used for cancer drugs, and we focus on quantifying the incentives to innovate these cocktails that contain combinations of patented and generic drugs and the resulting externalities.

Finally, our combination innovation model draws on existing work in the structural modeling of dynamic games, especially related to entry and innovation decisions (e.g., Pakes, 1986; Bajari et

[10] A follow-on innovator may have positive externalities on a previous innovator who has IP protection.

al., 2007; Pakes et al., 2007; Goettler and Gordon, 2011; Sweeting, 2013; Igami, 2017b; Igami and Uetake, 2019; Bodéré, 2023; Hodgson, 2024). We develop a computationally tractable model of combinatorial product innovation by applying approximation methods (Arcidiacono et al., 2013) and reducing the dimension of the state space with a method similar to oblivious equilibrium (Weintraub et al., 2008, Benkard et al., 2015) or moment-based Markov equilibrium (Ifrach and Weintraub, 2016).

Outline: The remainder of the paper is organized as follows. Section 2.2 describes the empirical setting and data. Section 2.3 presents a stylized model of combination innovation and facts consistent with predictions of this stylized model using data on investments in clinical trials, suggestive of inefficiencies in the innovation of cancer combination therapies. Section 2.4 develops our model of cancer drug demand and price setting, and Section 2.5 uses this model to quantify the externalities from combination therapy introduction. Section 2.6 develops a dynamic model of combination innovation to examine how externalities shape combination innovation decisions, and Section 2.7 designs cost-effective policies to support combination innovation. Section 2.8 concludes.

Table 2.1: Example Cancer Drug Combination Regimens

Cancer	Drugs	Firm	Dosage	Cycle	Trial Sponsor (Paper)	Trial Year
Colorectal	5-Fluorouracil Leucovorin Oxaliplatin	(Generic) (Generic) Sanofi	3,000 mg/m^2 IV day 1, 2 200 mg/m^2 IV day 1 100 mg/m^2 IV day 1	2 weeks	Public (de Gramont et al., 2000)	1999
CLL	Chlorambucil Obinutuzumab	(Generic) Roche	.5 mg/kg PO on day 1, 15 100 mg IV day 1; 900 mg IV day 2; 1000 mg IV day 8, 15	4 weeks	Roche (Goede et al., 2014)	2013
Breast	Gemcitabine Trastuzumab	Eli Lilly Genentech	1,200 mg/m^2 IV day 1, 8 2 mg/kg IV day 1	3 weeks	Eli Lilly (O'Shaughnessy et al., 2004)	2002
Non-Small Cell Lung	Carboplatin Paclitaxel Bevacizumab	Bristol-Myers Squibb Bristol-Myers Squibb Genentech	AUC of 6, IV day 1 200 mg/m^2 IV day 1 15 mg/kg IV day 1	3 weeks	Public (Sandler et al., 2006)	2001

Notes: Table shows four example cancer combination regimens from Chu and DeVita (2019), for colorectal, chronic lymphocytic leukemia (CLL), breast, and non-small cell lung cancer, respectively. For each regimen, the second column gives the component drugs, the third column gives the drug owner (original patent holder) at the year of trial submission, the fourth column gives the dosage of each drug, the fifth column gives the cycle length, the sixth shows who ran the pivotal clinical trial, and the final column is the submission year of that trial to ClinicalTrials.gov. IV means intravenous, and these drugs are administered into a vein using a needle or tube. PO means taken orally. AUC stands for area under the curve, and measures the exposure of a drug.

2.2 Cancer Combination Therapies: Setting and Data

2.2.1 Setting

Combination therapies, treatment regimens consisting of two or more drugs, are widely used in the treatment of most cancers.[11] Biological justifications for using combination therapies include reduced drug resistance,[12] reduced toxicity or side effects, and chemical synergies (Chu and DeVita, 2019). Most combination therapies consist of injectable (IV) drugs delivered in an outpatient clinical setting, while some regimens also contain prescription drugs taken orally. Example combination regimens are shown in Table 2.1, which displays the component drugs, dosage information, and firms with patented drugs in the regimen at the time of the first trial. These example regimens highlight that combinations are often taken over the course of many days or weeks, and the exact dosage can depend on patient characteristics (e.g., size of the tumor). Combinations contain a mix of on-patent drugs, potentially owned by different firms, and generic drugs.

[11] One of the first, and most influential, cancer combination therapies was discovered in the 1960s by Emil Frei (Frei et al., 1965) and was used to treat pediatric patients with acute lymphoblastic leukemia, resulting in dramatic reductions in mortality. It is known as the "VAMP" regimen and consists of 4 drugs: vincristine, amethopterin (methotrexate), 6-mercaptopurine, and prednisone.

[12] Approximately 90% of cancer-related deaths are associated with drug resistance (Bukowski et al., 2020; Dhanyamraju, 2024).

Combination therapies are evaluated in clinical trials to prove safety and efficacy.[13] Every recommended regimen in authoritative treatment guidelines such as Chu and DeVita (2023), UpToDate (Connor, 2024), and National Comprehensive Cancer Network (2024) cites a scientific paper reporting results from pivotal clinical trials demonstrating the regimen's safety and efficacy. While off-label use in untested combinations is theoretically feasible, it is uncommon: in addition to concerns about safety, efficacy, and liability, there are potential issues with insurance coverage (NCI, 2014).[14] Combination therapies often undergo combined Phase I-II or II-III trials rather than separate Phase I, II, and III trials.[15] In our sample of combination clinical trials (defined below), we estimate a success rate of 1.5% across cancers, with a range from .1% to 5% by individual cancer.[16] These clinical trials can be extremely costly, with average per-patient-phase costs in cancer trials exceeding $\$100,000$ (Sertkaya et al., 2014, Moore et al., 2020),[17] with combination therapy trials enrolling 116.2 patients on average.[18] Various entities may conduct combination trials, including private firms, universities, government agencies (including the National Cancer Institute, NCI), and other research institutions.[19] These trials can occur at different stages of a drug's life cycle. While most drugs are initially approved for solo use, some receive their first approval as part of a combination therapy.

In the spirit of Coase (1960), firms may attempt to internalize the market expansion externality by collaborating on trials (splitting trial costs) or contracting over the resulting revenues. However, such negotiations are limited in practice by a variety of frictions, including uncertainty about commercial potential of regimens or newly generated IP, concerns about violating the IP

[13] Combinations do not go through the same approval process as individual drugs, as the FDA does not typically explicitly approve combinations (unless the results from a combination trial are used in support of an individual drug's new drug application). Combinations are instead evaluated in clinical trials, and the results are reviewed by experts, such as oncologists, before being included in treatment guidelines.

[14] For example, CMS and the private insurer UnitedHealthcare use the NCCN Compendium (National Comprehensive Cancer Network, 2024) as their reference for oncology coverage policy. UnitedHealthcare requires prior authorization for these treatments, which necessitates the regimen be included in the NCCN Compendium.

[15] Phase I trials focus on safety and determining dose range. Phase II trials evaluate efficacy, and Phase III trials evaluate efficacy relative to the current standard of care.

[16] This success rate is calculated as the fraction of regimens trialed (in any phase) that appear in cancer treatment guidelines (Chu and DeVita, 2019) and are taken by patients in our drug usage microdata. We discuss these guidelines and drug usage microdata in more detail in Sections 2.2.2 and 2.4.3. Wong et al. (2019) estimates a success rate of 3.4% for oncology drugs, where the success rate is measured as the fraction of trialed drugs (single-agents) that receive FDA approval. The success rate we measure differs for two reasons. First, it is a success rate of combination therapies rather than individual drugs. Second, our definition of success is more restrictive in that we require the combination therapy to appear in treatment guidelines and be taken by patients.

[17] A breakdown of the cost components of clinical trials using data from industry reports is included in Appendix B.7.

[18] Unger et al. (2019) estimates costs for a subset Phase 3 trials run through the National Cancer Institute Clinical Trial Network, finding an average cost to the NCI of $16.6 million per trial.

[19] The NCI allocated approximately $857 million to research projects conducting clinical trials in 2022 (NCI, 2022b) out of its budget of $6.8 billion. The NCI also allocates significant funds to projects such as Cancer Centers and Clinical Cooperative Groups, which may also fund clinical trials (NCI, 2022a).

of component drugs, potential liabilities associated with the combination and negative effects on component drugs, and potential antitrust enforcement (Humphrey et al., 2011; Institute of Medicine et al., 2012; Deng, 2015; Boshuizen and Peeper, 2020; Podkonjak et al., 2021).[20]

The price of a combination therapy is constrained to be the sum of the prices of component drugs, precluding price discrimination based on whether a particular drug is used in a regimen. This feature limits a firm's ability to extract surplus from innovating combination therapies, and occurs for two key reasons. First, the drugs in combination therapies are typically not packaged together, as drugs are taken over the course of many days or weeks and in dosages that depend specifically on patient characteristics (e.g., measurements of the tumor), as shown in the example regimens in Table 2.1. This feature distinguishes cancer combination therapies from from fixed-dose combinations (e.g., Adderall) in which multiple drugs are packaged in a single physical dosage form (see, e.g., Cao and Chatterjee, 2023), and patenting is possible. Second, doctors' right to prescribe drugs for off-label use allows them to use any treatment regimen that has successfully completed clinical trials by simply prescribing and combining the individual drugs for their patients, making "self-production" of the regimen easy. These features reduce the value of patents for combinations, and pharmaceutical firms have generally not obtained exclusivity extensions for new combinations.[21]

2.2.2 Data

Making progress on characterizing incentives to innovate combinations requires data on cancer drug characteristics, clinical trials, treatment guidelines, drug usage and prices, and patient characteristics. We summarize each component in this section, while a detailed description of the data and sample construction procedure is given in Appendix B.1.

Cancer Drugs and Characteristics: We combine data from GlobalData, Drugs@FDA, and the Surveillance, Epidemiology, and End Results (SEER) Program to construct a comprehensive list

[20] For example, uncertainty regarding the effectiveness and attributes of combination therapies can lead to imperfect or asymmetric information among stakeholders. Moreover, drug owners may have dynamic incentives to free ride on the combination innovations developed by others. These frictions, and others, are discussed in detail in Appendix B.7.

[21] In general, firms may apply for New Clinical Investigation (NCI) exclusivity for demonstrating that a drug (or drug in a combination) is effective for some new indication not covered in the initial NDA. However, the exclusivity term that may be granted from successfully trialing a drug in a combination for a new indication will only apply to that new indication, rather than the original indications the drug was approved for. Enforcement of this exclusivity is difficult, especially for combinations, as the FDA may approve ANDAs for the original indications during this NCI exclusivity period. The potential to arbitrage across regimens and indications limits the usefulness of this type of exclusivity extension. We describe these regulatory details in more detail in Appendix B.7. We also we use data from the FDA Orange Book (Durvasula et al., 2023) to confirm that exclusivity extensions due to successful combination therapy clinical trials are relatively rare in practice, which is consistent with the small benefits from this type of NCI exclusivity.

of cancer drugs and their characteristics. Our primary source is GlobalData, which provides information on approximately 27,000 pipeline and marketed oncology drugs.[22] This database includes details on original drug manufacturers (i.e., patent holders), patent and exclusivity information, and other drug characteristics. The earliest drugs in our sample were approved in the 1950s, and we have information on drugs through 2023, including investigational drugs. For drugs marketed in the US, we supplement this data with information on generic competition from Drugs@FDA, quantified by the number of approved Abbreviated New Drug Applications (ANDAs). Finally, we use data from the SEER CanMED Healthcare Common Procedure Coding System (HCPCS) database to compile current and historical HCPCS codes for each drug, which we use to identify patient usage of drugs delivered in an inpatient or outpatient setting.

Clinical Trials: We use a registry of privately and publicly-funded clinical trials from ClinicalTrials.gov, focusing on trials run between 1990 and 2022. Submission of trials run in the US has been mandatory since 2007,[23] though many trials were documented prior to this date due to factors such as patient recruitment needs, funding requirements, or journal publication mandates. This data includes information on the particular treatment tested (including what drugs), sponsor, collaborators, indication, start and end dates, and trial phase.[24] We classify sponsors as either private or public, where public includes the NIH, universities, and other non-profit research groups.

We subset to oncology clinical trials using the list of oncology Medical Subject Headings (MeSH) terms and free text keywords in Califf et al. (2012). We use a large language model, OpenAI's GPT-4o, to extract information on drugs used in the control and treatment arm(s) of each trial and remove oncology trials that use only non-pharmaceutical treatment methods like surgery or radiation.[25] Our sample of clinical trials includes trials that test at least one drug with at

[22] GlobalData acquires its information from myriad sources, including scientific publications and conferences, company annual reports, regulatory filings, and direct contact with companies. Their comprehensive coverage of pipeline drugs begins in 2000, and marketed drugs in 1983.

[23] We will use clinical trials run from 1990 to 2022 as our main sample, but we show robustness of our analysis to focusing on clinical trials run after 2007 given the beginning of mandatory reporting in that year.

[24] ClinicalTrials.gov defines a "sponsor" as the entity that "initiates the study and who has authority and control over the study", while a "collaborator" is any organization other than the sponsor that provides support (funding, design, implementation, data analysis, or reporting). The Food and Drug Administration Amendments Act of 2007 requires that clinical trial information submissions include the sponsor, but not collaborators (42 CFR 11). It is not clear how extensive reporting of collaborators is. Manual comparison of a small set of clinical trials and their associated publications (author affiliations and funding source disclosure) suggests that the collaborators field for firms involved in the trial is relatively complete, though additional work could be done to further complete the field by examining funding disclosures in research papers associated with trials.

[25] We test the performance of this approach by constructing a random sample of 100 oncology clinical trials from 2007. We manually extract the drugs used in the treatment arm(s) of each trial, if applicable. The large language model produces the correct result for 98 of the 100 trials. In one case, the model omitted a component drug of combination, and in the other failed case it incorrectly classified a surgical implement as a drug.

least one location in the United States,[26] leaving approximately 28,000 trials. We merge this panel of clinical trials with the drug information constructed above for each drug used in a treatment arm of the trial. This requires harmonizing sponsor names and drug names over the life-cycle (investigational, branded, generic)—details about our merging procedure are given in Appendix B.1. Approximately 85% of trialed drugs merge to GlobalData, and drugs that do not merge are typically pipeline drugs that are trialed in at most Phase I. We define a combination trial as a trial that tests at least two drugs together in a treatment arm.

Table 2.2a presents summary statistics for the clinical trials in our sample, which includes both single-agent and combination therapy trials. 58.1% of trials are for combination therapies. The median number of drugs tested in a combination trial is 2, with a median of 1 generic drug and 1 on-patent drug as a component. The table additionally shows the fraction of trials run by different lead sponsor–collaborator pairs. Firms with no collaborators run a higher fraction of single-agent trials (38.7%) than combination trials (24.2%). There are relatively few trials with multiple firms collaborating (3.9% of single-agent trials, 4.8% of combination trials), and there is a higher fraction of firm-public collaboration for combination trials (21.7%) compared to single-agent trials (13.6%). Combination therapy trials enroll more patients than single-agent trials on average, with a mean enrollment of 116.2 patients compared to 104.0 patients. Table 2.2b presents summary statistics for drugs that are trialed in cancer clinical trials in our sample. Trials may include drugs that are approved or that are still investigational (i.e., not yet approved), whether for single-agent therapies or combination therapies.

Cancer Treatment Guidelines: While the clinical trial data provide a comprehensive list of trialed single-agent and combination therapies, it is difficult to parse the results of those trials to determine what regimens were "successful."[27] We use treatment guidelines from the medical literature (Chu and DeVita, 2019) to define successful regimens for each cancer, and this subset of successful regimens then forms the set of products for our demand system, developed in Section 2.4.[28] These treatment guidelines are a reference for oncologists when determining treatment for a patient, and are updated regularly to incorporate new treatments. These guidelines contain similar regimens to compendia that are used by insurers to determine coverage of cancer drugs, such as the National Comprehensive Cancer Network (2024) compendium. Chu and DeVita (2019) contains information about approximately 700 regimens for 40 different

[26] We focus on trials with at least one location in the United States as these trials are the most likely to be included in treatment guidelines published in the US and thus taken in the US, which our demand data covers.

[27] Furthermore, many cancer clinical trials do not provide timely public reporting of results (Kao et al., 2023).

[28] While we could use the set of all trialed regimens as the products in the demand system, it is computationally convenient to focus on the (significantly) smaller set of successful regimens as defined by Chu and DeVita (2019). We show that these regimens account for most of patient drug usage, and "unsuccessful" or regimens that have not been trialed typically see little use.

Table 2.2: Cancer Clinical Trial and Drug Summary Statistics

a Cancer Drug Clinical Trials

	Single			**Combination**		
	Mean	Median	Std Dev	Mean	Median	Std Dev
Number of Drugs	1.0	1.0	0.0	3.03	2.0	1.64
Number of Generic Drugs	0.07	0.0	0.25	1.0	1.0	1.37
Number of Patented Drug Owners	1.0	1.0	0.0	1.5	1.0	0.87
Sponsor Firm Solo	0.39			0.24		
Sponsor Firm + Firm	0.04			0.05		
Sponsor Firm + Public	0.14			0.22		
Sponsor Public Solo	0.44			0.49		
Enrollment	104.01	38.0	255.15	116.16	40.0	276.59
Total Trials	11,959			16,596		

b Cancer Drugs in Trials

	Approved			**Not Approved**		
	Mean	Median	Std Dev	Mean	Median	Std Dev
Number of Drugs	399			3,281		
Total Trials	107.9	25.0	223.61	3.3	1.0	7.32
Fraction Combination	0.68	0.75	0.28	0.47	0.5	0.44
Biologic Indicator	0.09			0.07		

Notes: Panel 2.2a shows summary statistics of characteristics of clinical trials in our sample, separately for single (columns 1-3) and combination (columns 4-7) trials. Number of Drugs is the number of drugs included in treatment arms of the trial. Number of Generic drugs is the number of generic drugs (determined at the start of the trial) included in the treatment arms of the trial. Number of Patented Drug Owners is the number of unique firms that own on-patent drugs in the regimen (calculated at the time of trial). Sponsor Firm Solo is the fraction of trials run with a firm as the lead sponsor and no collaborators. Sponsor Firm + Firm is the fraction of trials run with a firm as the lead sponsor with a firm (and no other agents) as a collaborator. Sponsor Firm + Public is the fraction of trials with a Firm and Public innovator collaborating. Sponsor Public Solo is the fraction of trials run with a public innovator as the lead sponsor and no firms (but potentially other public innovators) as collaborators. Enrollment is the number of patients enrolled in the trial. Panel 2.2b shows summary statistics of drugs that are trialed in cancer clinical trials, separately for drugs that were eventually approved (i.e., approved at some point during our sample) (columns 1-2) versus not (or not yet) approved during our sample period (columns 3-4). Number of Drugs is the number of unique drugs in each group. Total Trials is the total number of cancer drug clinical trials a drug is involved in. Fraction combination is the fraction of a drug's trials that are for combinations. Biologic indicator is an indicator of whether a drug is a biologic drug.

cancers. The average number of drugs included in a regimen is 2.05, with a range between 1 and 11. Each recommended regimen cites a scientific paper that presents results from a pivotal clinical trial showing that regimen's efficacy. When new combination therapies demonstrate efficacy in clinical trials, they may be incorporated into existing guidelines, either as additional options, modifications to current protocols, or in some cases, replacements for older treatments, depending on their comparative benefits and safety profiles for certain patients.

Cancer Drug Usage, Prices, and Patients: To determine regimen usage, prices, and patient demographics, we analyze claims and enrollment data from two sources: the Center for Medicare and Medicaid Services (CMS) for traditional Medicare beneficiaries, and Marketscan Commercial Claims and Encounters database for a sample of privately insured individuals. The Medicare data is from 1998 to 2019, and includes information about a 20% random sample of Medicare beneficiaries, of which we observe claims information for traditional Medicare beneficiaries.[29] The MarketScan data is from 1996 to 2013, and is a sample of individuals under age 65 with employer-provided health insurance and their dependents. The sample of patients included in Marketscan expands considerably throughout the sample period, covering approximately 3 million individuals in 1996 and over 50 million by 2013. We discuss this data, including how we measure regimen usage and prices, in detail in Section 2.4.3.

2.3 Stylized Model and Descriptive Evidence

This section introduces a stylized model of combination innovation in the pharmaceutical industry and presents evidence supporting its predictions regarding combination innovation decisions. We first use the model to characterize privately and socially optimal innovation decisions, highlighting the existence of market expansion externalities and missing property rights for combinations. The model also reveals the economic primitives underlying the innovation externalities that we estimate in Sections 2.4 and 2.5.

We then document three facts consistent with predictions of the stylized model with market expansion externalities: Private firms are significantly less likely than public researchers to trial combinations relative to single-drug therapies (Fact 1). When private firms do trial combinations, they are biased toward combinations consisting of their own branded drugs (Fact 2). Finally, after (and in anticipation of) generic entry, a drug is less likely to be trialed in a combination

[29] Physician-administered drug claims from Medicare Advantage (Part C) patients are excluded from our CMS data. These Part C plans were first introduced in 1985, and enrollment has dramatically increased since the beginning of our sample. At the end of our sample, 2019, approximately 39% of Medicare beneficiaries were enrolled in a Medicare Advantage plan, with the remaining 61% enrolled in traditional Medicare, as reported by the Kaiser Kaiser Family Foundation (Ochieng et al., 2023).

by its original owner, but it is more likely to be trialed by other firms, suggesting potential delays in combination innovation due to intellectual property protections (Fact 3). These facts demonstrate that private innovation decisions are driven by the ownership structure of drugs, and thus that market expansion and missing property rights decisively shape the magnitude, direction, and timing of combination innovation. They also suggest inefficient underinvestment in combination therapies by firms, which we explore in greater detail in the remainder of the paper.

2.3.1 A Stylized Model of Externalities Arising from Combination Innovation

The key features of combination innovation can be best understood using a static model of the market for drug regimens used to treat a given cancer.[30] A regimen is a set of drugs $r = \{d_{r1}, \ldots, d_{rn}\}$ taken in fixed proportions by a patient, and for simplicity we assume that a patient who takes regimen r consumes one unit of each drug $d \in r$. The set of all drugs that can be combined to form regimens is $\mathcal{D}$, and the set of all regimens that have been introduced and can be taken by patients is $\mathcal{R} \subseteq 2^{\mathcal{D}}$, where the power set $2^{\mathcal{D}}$ is the set of all possible regimens. Each drug d is either generic or owned by a firm, and we identify each firm with the set of drugs it owns, $f = \{d_{f1}, \ldots, d_{fn}\}$. The set of all firms is $\mathcal{F} \subseteq 2^{\mathcal{D}}$.

Drug d is produced at marginal cost $mc_d \geq 0$ and sold at price $p_d \geq mc_d$.[31] As discussed in Section 2.2, firms cannot price discriminate between patients taking different regimens, so the price of any regimen r is simply the sum of the underlying drug prices, $p_r = \sum_{d \in r} p_d$. The share of patients who take regimen r is $s_r(p)$, and we normalize the total mass of patients to one for convenience. The profits earned by drug d through regimen r are

$$\pi_{dr}(p) \equiv (p_d - mc_d) s_r(p) \mathbb{1}[d \in r].$$

The total profits earned by drug d are then $\pi_d(p) \equiv \sum_{r \in \mathcal{R}} \pi_{dr}(p)$, while we assume profits for generic drugs are zero. We let $CS(p)$ denote aggregate consumer surplus, and we denote total

[30] We incorporate dynamics into our structural model in Section 2.6.

[31] In this stylized model, we make a simplifying assumption for exposition that each drug has a single price that all patients pay, ignoring the potential difference between list prices and out-of-pocket prices and price dispersion across patients. We model these features in our empirical framework in Section 2.4.

welfare (Marshallian surplus) by[32]

$$W(p) \equiv \text{CS}(p) + \sum_{d \in \mathcal{D}} \pi_d(p).$$

A pharmaceutical firm can pay a fixed cost $\kappa > 0$ to run a clinical trial for a new regimen r^+. With probability $\chi > 0$ the trial is successful and the regimen can be taken by patients; with complementary probability $1 - \chi$ the trial fails. When successful, the addition of the "new product" r^+ leads to new market shares $\tilde{s}_r(p)$ and a new set of equilibrium drug prices $\tilde{p}$. The change in shares reflects how patients adjust their regimen choices at each price vector, while the change in equilibrium prices reflects pricing conduct in the market. We use tildes ($\sim$) to distinguish post-introduction values from pre-introduction values, and we use Δ to denote the equilibrium change in a value. For example, $\tilde{\pi}_d(p)$ denotes the post-introduction profit function for drug d, while $\Delta\pi_d \equiv \tilde{\pi}_d(\tilde{p}) - \pi_d(p)$ denotes the change in equilibrium profits.

Private versus Socially Optimal Innovation Incentives: In this static model, a private firm $f \in \mathcal{F}$ finds it profitable to trial the new regimen r^+ given existing regimens $\mathcal{R}$ if and only if the expected change in its profits exceeds the trial cost, $\chi \sum_{d \in f} \Delta\pi_d > \kappa$. We can decompose the change in profits $\Delta\pi_d$ for each drug d into three terms:

$$\begin{aligned}\Delta\pi_d = &\underbrace{\sum_{r \in \mathcal{R}} (p_d - mc_d)[\tilde{s}_r(p) - s_r(p)]\mathbb{1}[d \in r]}_{\text{business stealing}} + \underbrace{(p_d - mc_d)\tilde{s}_{r^+}(p)\mathbb{1}\left[d \in r^+\right]}_{\text{market expansion}} \\ &+ \underbrace{\sum_{r \in \mathcal{R} \cup \{r^+\}} [(\tilde{p}_d - mc_d)\tilde{s}_r(\tilde{p}) - (p_d - mc_d)\tilde{s}_r(p)]\mathbb{1}[d \in r]}_{\text{price adjustment}}. \end{aligned} \tag{2.1}$$

The first term captures the change in the profits drug d derives from existing regimens $r \in \mathcal{R}$, holding initial prices p fixed. This *business stealing* effect arises as patients substitute away from existing regimens and toward to the new regimen r^+, and it is weakly negative provided that patient preferences satisfy the Weak Axiom of Revealed Preference (WARP). The second component is instead weakly positive and captures the profits earned from the new regimen r^+ at initial prices p, which we refer to as the *market expansion* effect. At fixed prices, any drug not used in the new regimen $d \notin r^+$ earns lower profits because of the business stealing effect, while a drug used in the new regimen $d \in r^+$ earns higher profits as the market expansion effect

[32] In this stylized model, we exclude insurer costs from the total welfare function for simplicity. Patients are often not responsible for the full cost of drugs and the total welfare function must capture insurer welfare. We defer analysis of insurers (both public and private) to the empirical model in Section 2.4.

dominates the business stealing effect.[33] Without property rights over regimens, a firm can only profit from the introduction of r^+ if it owns a drug used in this regimen, and even then only if the market expansion effect dominates the business stealing effect on average across its drugs. The balance of these effects depends critically on patient substitution patterns. The final term reflects all changes to profits because of the equilibrium price adjustment from p to $\tilde{p}$. This adjustment depends on pricing conduct, drug ownership, and the change in the demand system. It affects both the observed market shares $\tilde{s}_r(\tilde{p})$ and the margins $\tilde{p}_d - mc_d$.

To derive the *net externality* of firm f's introduction of regimen r^+, we compare the change in surplus internalized by the firm to the change in total welfare. The change in total welfare includes both the change in consumer surplus and the change in profits:

$$\Delta W \;=\; \Delta \text{CS} \;+\; \sum_{d \in \mathcal{D}} \Delta \pi_d.$$

The *net externality* of firm f's introduction of regimen r^+ is then:[34]

$$\text{Net Externality}_{r^+} \;\equiv\; \Delta \text{CS} \;+\; \sum_{d \notin f} \Delta \pi_d. \tag{2.2}$$

The first term is the *consumer surplus externality* (i.e., the standard non-appropriability of consumer surplus that arises in innovation), and WARP implies that it is weakly positive when prices are fixed ($\tilde{p} = p$).[35] This externality is not specific to combination innovation and arises whenever an innovating firm cannot appropriate all consumer surplus, though it may be particularly large in this setting since firms cannot price discriminate by regimen. The second term is the *net profit externality*, which captures the pecuniary externality on other firms $f' \neq f$. As the decomposition (2.1) demonstrates, the net profit externality includes business stealing, market expansion, and price adjustment effects. Even at fixed prices, the presence of business stealing implies that the net profit externality, and thus the overall net externality, may be positive or negative.[36] But if the positive effects on consumers through increased surplus and

[33] More generally, the business stealing and market expansion effects arise whenever a firm innovates a new product that is respectively substitutable or complementary with an existing product. A distinctive feature of combination innovation is that these effects appear simultaneously.

[34] The externality from the firm's decision to *trial* regimen r^+ is simply the net externality multiplied by the success probability χ.

[35] When quantifying these externalities in Section 2.5, we additionally include a term for the fiscal externality on insurers that pay remaining drug costs after patient deductibles, co-payments, or co-insurance. The change in this net surplus (consumer surplus minus insurer costs) may be positive or negative.

[36] The endogenous growth literature identifies another externality, knowledge spillovers, which arises when one firm's innovation builds on (and thus benefits from) another firm's innovation (P. M. Romer, 1990). This externality plays a critical role in combination innovation, because new combinations often use components invented by

other firms through market expansion dominate business stealing, then firms will underinvest in combination innovation relative to the social planner.

Measuring Innovation Externalities: The stylized model clarifies the economic primitives that we must measure in order to quantify the net externality from combination innovation and any corresponding inefficiency. The definition of the net externality (2.2) from firm f's introduction of regimen r^+ requires that we measure the corresponding change in consumer surplus and the change in profits earned by other firms' drugs, neither of which are directly observable. In Section 2.4, we estimate a discrete-choice random coefficients demand system for regimens that allows us to compute ΔCS given prices p and $\tilde{p}$. Similarly, to recover marginal costs of drug production and compute the profit change $\Delta\pi_d$, we estimate a model of Nash bargaining over drug prices between manufacturers and insurers. Both the demand system and the pricing model are crucial for decomposing the profit change $\Delta\pi_d$ into the business stealing, market expansion, and price adjustment terms in (2.1). To assess the extent to which the net externality yields inefficient innovation decisions by firms, we must additionally recover the innovation fixed costs κ and the success probability χ. In Section 2.6 we estimate a dynamic innovation model that allows us to recover these parameters, making use of the estimated demand system and pricing model as well as firms' observed clinical trial decisions.

Before recovering these primitives, we first document three facts related, respectively, to the amount, direction and timing of combination innovation. These facts are consistent with predictions of the stylized model having positive market expansion externalities and missing property rights that shape firm incentives to innovate combinations, and highlight the importance of drug ownership in determining the amount, direction, and timing of combination innovation.

2.3.2 Who Funds Combination Innovation? Single versus Combination Trials

Our first fact explores incentives to innovate single versus combination therapies, motivated by a prediction of the stylized model: all else equal (e.g., regimen quality, price), market expansion externalities imply that a firm will be less likely than the social planner to trial a particular combination compared to a single-agent therapy. The data support this prediction:

Fact 1. A firm's research portfolio consists of a smaller share of combination trials than that of a social planner or its proxy (social-welfare minded public researchers).

To provide evidence for this fact, we calculate how many cancer combination and single-drug

different firms. Knowledge spillovers are pervasive but challenging to measure in the pharmaceutical industry, so we limit our focus to pecuniary externalities.

Figure 2.1: Clinical Trials by Funding Type

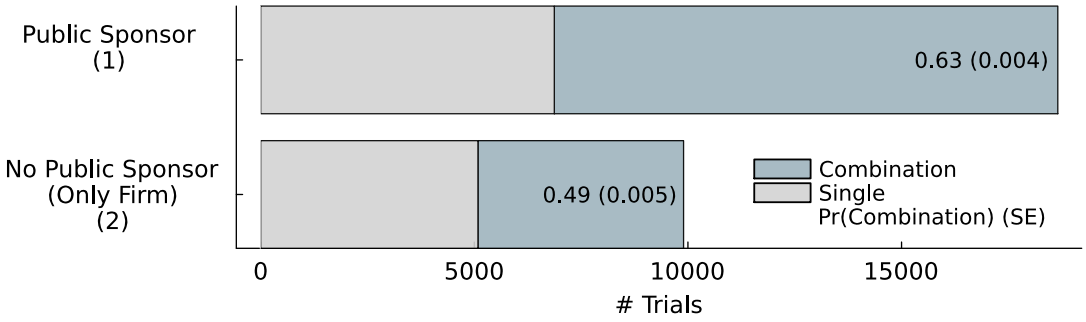

Notes: Figure shows the number of trials run by different sponsors broken into combination (blue) and single drug trials (gray) and the fraction of that sponsor's trials that are for combinations (black number). The "Public Sponsor" bar includes trials with at least one public sponsor. The "No Public Sponsor (Only Firm)" bar includes trials that are solely sponsored by firms. Robust standard errors are in parentheses on each bar, which come from a trial-level regression of an indicator for being a combination trial on sponsor type. The corresponding regression table is shown in Appendix B.2, where we also show robustness with respect to the time period of trials included, additional controls (cancer, trial submission year, and trial size), and estimating separately for collaborative trials.

trials are sponsored by publicly-funded innovators versus firms. In Figure 2.1, we see that trials with public funding are the most numerous and more likely to trial combination therapies (63% of trials are for combinations)[37] compared to trials with only firm sponsors (49% of trials are for combinations).[38] Standard errors of a trial-level regression of an indicator of being a combination trial on indicators for different funding types are also displayed, indicating statistically significant differences between the combination trialing behaviors of different funding types.[39]

This fact demonstrates that firms are biased away from clinical trials for combination therapies relative to publicly-funded researchers. Firms may be dissuaded from trialing a regimen if many of the resulting profits accrue to other firms that own drugs in the same regimen, unlike publicly-funded researchers with other (potentially aggregate welfare-minded) objectives. However, this fact is also consistent with the idea that firms are "crowded out" of funding clinical trials for combination therapies by public institutions, so this fact does not provide direct evidence of *inefficient* underinvestment. We disentangle these effects in the remaining facts.

[37] We also note that public institutions run many single drug clinical trials, often for second uses or indications that were not tested by the drug's original manufacturer. We decompose these types of trials in Appendix B.1.

[38] Collaborative trials between firms and public sponsors are included in the "Public Sponsor" bar (22% of total combination trials), while collaborations between 2 or more firms (and no public sponsors) are included in the "No Public Sponsor" bar (5% of total combination trials).

[39] We show robustness of this fact and the remaining facts in this section to focusing on clinical trials run after 2007 in Appendix B.2. We also show robustness to controlling for other trial characteristics, such as cancer type and trial (submission) year.

2.3.3 What is the Direction of Combination Innovation? Trials and Ownership

Our next fact characterizes the direction of combination innovation by calculating the probability of running a particular combination clinical trial as a function of drug ownership. In the stylized model, all else equal, a firm earns the highest expected profits from combination regimens that consist primarily of its own drugs, compared to drugs owned by other firms or generic drugs. Compared to a firm, the social planner is more likely to trial generic drugs, since it considers the positive externality on consumers. The data support this prediction:

Fact 2. Firms are more likely to trial combinations of their own branded drugs than combinations with other firms' branded drugs or generic drugs. Compared to firms, publicly-funded researchers are more likely to trial combinations with generic drugs.

To demonstrate this empirically, we construct a panel of trialing decisions for two-drug combinations and calculate the probability of trialing a combination as a function of drug ownership. We focus on two-drug combinations for ease of exposition, which include approximately 50% of the combination trials in our data.[40] Let $\mathcal{D}_t$ denote the set of all drugs discovered by the beginning of year t, and let $\mathcal{L}_{2t} \subset \mathcal{D}_t^2$ denote the set of all two-drug clinical trials $r = (d_1, d_2)$ that have yet to be run (but can be run) by the beginning of year t.[41] As discussed earlier, an unusual and attractive feature of this setting is that we are able to observe the risk set of potential innovations; in other words, we can observe combinations that have not been trialed since we can enumerate the set of potential combinations for any given set of drugs. We classify two-drug combinations in $\mathcal{L}_{2t}$ into groups depending on whether the components are branded or generic and owned by the same or different firms.

We then measure how clinical trial decisions deviate from the benchmark in which clinical trials are chosen uniformly from all possible clinical trials $\mathcal{L}_{2t}$. Let L_{it} denote the total number of trials run by innovator of type i in year t, where i is either private (firms) or public. If these trials were chosen uniformly at random, the probability that any given trial r is chosen (with replacement) would be

$$\beta^{\text{unif}}(L_{it}) = 1 - \left(1 - \frac{1}{|\mathcal{L}_{2t}|}\right)^{L_{it}}.$$

We then estimate the following equation for private innovators i, pooling across two-drug

[40] An extension of this fact to 3 or more drugs amplifies the general patterns of this exercise, and is discussed in Appendix B.2.

[41] For computational tractability, we restrict to drugs that are trialed in at least one Phase 2 trial (or later phases). This drops drugs that are investigational and fail Phase 1 trials.

combination regimens r and years t:

$$\frac{\text{Trial Now}_{rit}}{\beta^{\text{unif}}(L_{it})} = \gamma_1 (\text{2-Brand Same})_{rit} + \gamma_2 (\text{2-Brand Different})_{rit} + \gamma_3 (\text{Has Generic})_{rit} + \varepsilon_{rit}, \tag{2.3}$$

where Trial Now_{rit} is an indicator that regimen r is trialed by innovator type i in year t. Each coefficient γ_k measures how many times more likely an innovator of type i is to run a given trial of type k in a year relative to uniform selection from the set of available combinations. We estimate a similar equation for public innovators, where the types of combinations included are 2-Brand, Brand and Generic, and 2-Generic.

Figure 2.2 plots the estimated coefficients, where firms are shown on the left and public innovators the right. Firms have a significantly higher relative probability of trialing combinations consisting of two branded drugs owned by the same firm (approximately 11 times more likely than uniformly at random) compared to combinations that contain two branded drugs owned by different firms (relative probability of .46) or combinations with a generic drug (relative probability of 1.03). Compared to firms, public innovators are more likely to trial combinations with a branded drug and a generic drug (relative probability of 1.35) or two generic drugs (relative probability of 5.05).

While this fact is consistent with firms directing combination innovation based on financial incentives towards high-profit own-drug combinations and away from combinations with market expansion externalities on other firms, it could also stem from a number of confounding factors. First, firms might be more familiar with their own drugs and therefore more likely to trial them in combinations. Second, running trials with a firms' own drugs might be significantly less expensive. Third, drugs produced by the same firm might be more likely to be complementary because of specialization in particular research areas.[42] Finally, public innovation in combinations with drugs owned by different firms or generic drugs could cause crowd-out. We present evidence against each of these alternative hypotheses.

To address the potential for intra-firm familiarity, we refine the types of combinations considered in Equation (2.3) to exploit variation in ownership over time. Focusing on combinations that consist of a branded drug and a generic drug, we can decompose these combinations into two groups: those where the two drugs were both initially owned (i.e., patented) by the same firm or not. Greater familiarity with own drugs suggests that firm innovation decisions would be

[42]The other direction is also possible: as many combinations are often effective through reducing drug resistance, if a firm is more likely to produce multiple drugs with similar mechanisms of action, then its drugs may be less likely to form effective combinations.

Figure 2.2: Relative Probability of Trialing 2 Drugs Together by Drug Ownership Status

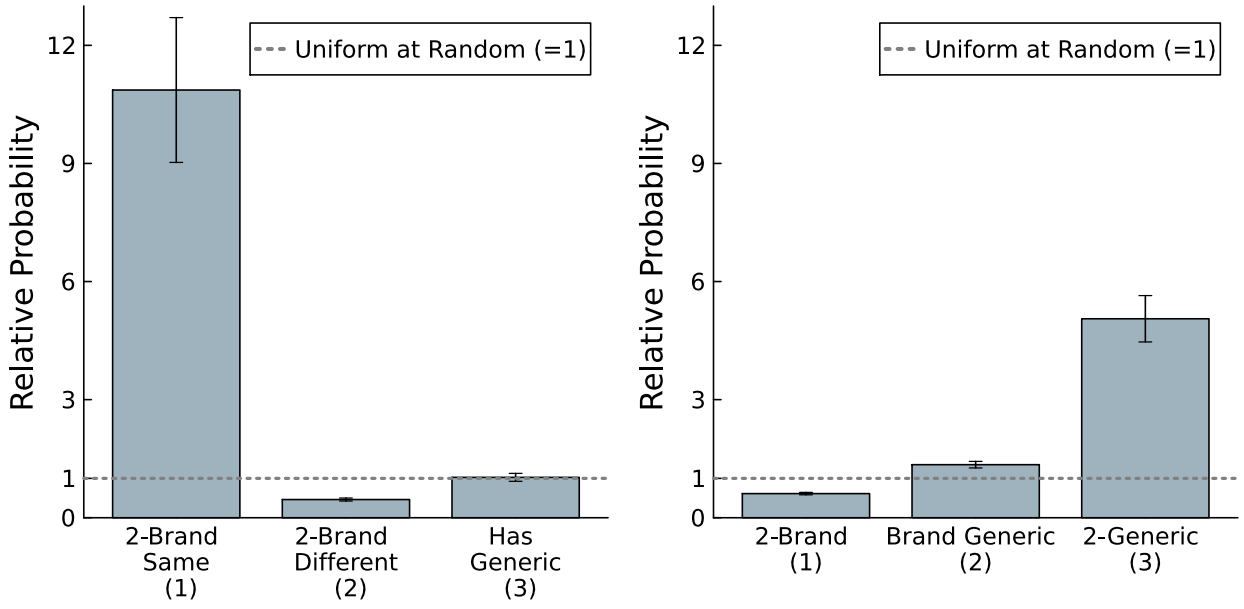

(a) Firm

(b) Public

Notes: Figure shows estimated γ_k coefficients of Equation (2.3). Private innovators (firms) are on the left panel, and public innovators on the right. The dotted line at 1 displays what trialing behavior would look like if innovators selected combinations uniformly at random from available combinations. 95% confidence intervals for the regression coefficients calculated from robust standard errors are displayed on each bar. Regression tables and additional robustness checks are given in Appendix B.2.

significantly biased towards combinations where they initially had patents for both drugs versus only one. Estimating an expanded version of Equation (2.3) largely rejects this hypothesis. Figure 2.3(a) plots the key coefficients. Firms trial combinations of a branded drug and a generic drug where both drugs were initially owned by the same firm 1.76 (.56) times more than uniformly at random (left bar), compared to 1.01 (.05) times more than uniformly at random (right bar) when patented by different firms. This difference is small relative to the bias towards combinations of two branded drugs by the same firm (11 times more than uniformly at random).[43]

Reports of the different contributors to trialing costs suggest that the cost of the clinical procedure, which in the case of combination trials includes the cost of the drugs and any costs

[43]This difference in relative probabilities between can also be explained by the small but positive profits original patent holders may still earn from their drugs that have experienced generic entry. Generic entry erodes monopoly profits over a number of years as patients substitute away from the branded drug, leaving higher financial incentives combinations with their own generic drugs for a few years.

associated with administering the drugs, makes up at most 20% of total costs on average.[44] This fraction is likely too small to rationalize the large difference in observed trialing decisions.

To address the potential for intra-firm complementarity of drugs, we provide two additional pieces of evidence. First, we can conduct a similar test to the intra-familiarity test but with public innovation decisions, where we decompose the 2-Generic category into two groups: combinations where both generic drugs were initially patented by the same firm or not. Estimating an expanded version of Equation (2.3) reveals that publicly-funded innovators trial these two types of generic drug combinations with similar relative probabilities, as shown in Figure 2.3(b). They trial combinations with two generic drugs originally patented by the same firm 4.66 (1.34) times more than uniformly at random (left bar) and two generic drugs originally patented by different firms 5.07 (.31) times more than uniformly at random (right bar). Thus, publicly-funded researchers trial combinations consisting of two generic drugs with similar probabilities, regardless of original ownership.

The second piece of evidence against intra-firm complementarity of drugs uses data on laboratory measures of two-drug combination efficacy. In 2017, the National Cancer Institute (NCI) released the NCI ALMANAC ("A Large Matrix of Anti-Neoplastic Agent Combinations"): a database of laboratory tests of all two-drug combinations derived from a set of approximately 100 FDA approved cancer drugs (Holbeck et al., 2017). The laboratory tests measure tumor growth rates on 60 different tumor cell lines for various dosages, resulting in approximately 3 million tests. We merge drug ownership information with this database and show that combinations consisting of drugs owned by the same firm or different firms result in similar tumor growth rates, displayed in Figure 2.4. While the mapping between laboratory results and clinical trials is not immediate, this fact suggests ex-ante measures of complementarity are not significantly different for combinations consisting of two drugs owned by the same firm or not.[45]

Finally, crowd-out by publicly-funded researchers is a less compelling explanation for low firm activity in 2-Brand Different trials, because public institutions also run few trials with two branded drugs (relative probability of .6) relative to other two-drug combination therapies. Public researchers, whose objective may be closer to aggregate welfare than firms, are more likely than firms to run trials consisting of generic drugs (relative probability of 5.05).

[44] We present a detailed breakdown of these trialing costs in Appendix B.7.

[45] Robustness checks of this result are described in detail in Appendix B.2, which test additional measures of complementarity and find similar results.

Figure 2.3: Relative Probability of Trialing 2 Drugs Together – Alternative Explanations

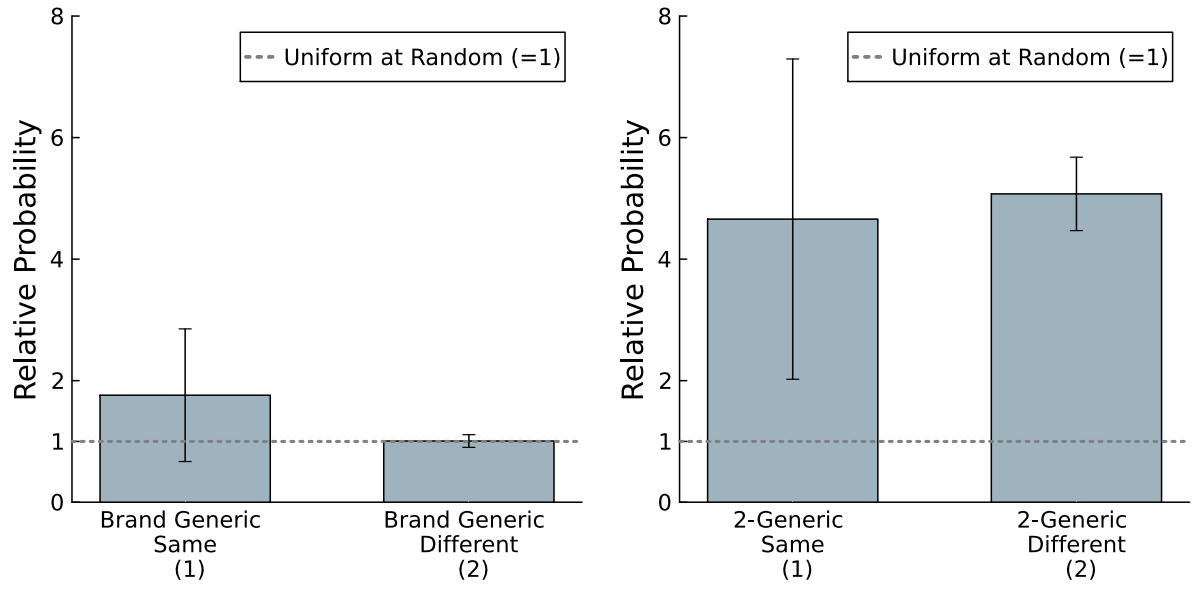

(a) Firm (b) Public

Notes: Figure shows estimated γ_k coefficients of Equation (2.3), expanded to include (and only showing coefficients for) the additional categories of Brand + Generic Same and Brand + Generic Different for private innovators and 2-Generic Same and 2-Generic Different for public innovators. Private innovators (firms) are on the left panel, and public innovators on the right. The dotted line at 1 displays what trialing behavior would look like if innovators selected combinations uniformly at random from available combinations. 95% confidence intervals for the regression coefficients calculated from robust standard errors are displayed on each bar. Regression table given in Appendix B.2.

2.3.4 When Does Combination Innovation Occur? Trials and Generic Entry

Our final fact characterizes how the regulatory environment affects the timing of combination innovation. In the stylized model, generic entry into a drug reduces the original owner's incentive to trial it in combinations. In contrast, the decline in price due to generic entry may incentivize other firms and publicly-funded innovators to trial it in new combinations.[46] The data support this prediction:

Fact 3. A cancer drug is involved in more combination therapy clinical trials after (and in anticipation of) generic entry. The original firm runs fewer trials, but both other firms and

[46]This holds formally when some drugs d and d' become "more complementary" in demand after the introduction of a new regimen r^+:

$$\frac{\partial \Delta\pi_d(p)}{\partial p_{d'}} < 0 \iff \frac{\partial \tilde{s}_{r^+}(p)}{\partial p_{d'}} + \sum_{r\in\mathcal{R}:d\in r} \frac{\partial \tilde{s}_r(p)}{\partial p_{d'}} < \sum_{r\in\mathcal{R}:d\in r} \frac{\partial s_r(p)}{\partial p_{d'}}.$$

Figure 2.4: NCI ALAMANC Distribution of Tumor Growth Rates by 2-Drug Combination

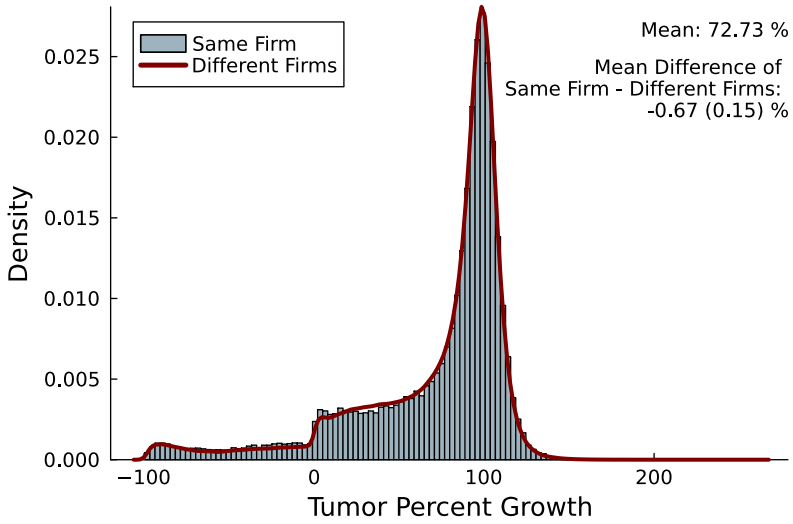

Notes: Figure shows the distribution of tumor growth rates for combinations in the NCI ALMANAC, separately for combinations consisting of drugs owned by the firm or not. Each distribution is normalized to show a density. The mean growth rate is 72.73%, and the coefficient (SE) of a combination-level regression of the tumor growth rate on an indicator for the combination having two drugs owned by the same with drug and cancer fixed effects is -.67% (.15%). Regression tables and additional results are included in Appendix B.2.

public researchers run more.

To provide evidence for this fact, we construct a panel of drugs that experience generic entry before 2023. For each drug d and year t, we calculate the total number of combination trials involving that drug that start in that year, and how many of those combination trials are run by (either as lead sponsor or as a collaborator) the firm that originally patented the drug, publicly-funded innovators, or other firms. These groups are not mutually exclusive since trials can potentially be run with collaborators. We then estimate

$$Y_{dt} = \beta \mathbf{1}(\text{Generic Indicator}_{dt}) + \delta_t + \delta_d + \varepsilon_{dt}, \tag{2.4}$$

where Y_{dt} is one of the counts of combination clinical trials listed above, $\text{Generic Indicator}_{dt}$ is an indicator of whether drug d is within five years of experiencing generic entry at time t, δ_t is a year fixed effect to control for different rates of trialing over time, and δ_d is a drug d fixed effect to control for unobservable drug characteristics that affect the rate of trialing. We define the generic indicator starting 5 years before generic entry as combination trials may start in anticipation of generic entry, with results from the trial coming as generic entry occurs.

Table 2.3 summarizes the results. A drug is used in more combination trials on average in the five years leading up to generic entry (i.e., in anticipation) and after generic entry, compared to

Table 2.3: Combination Trials and Generic Entry Regression

	Original Firm	Public	Other Firms	Total Trials
	(1)	(2)	(3)	(4)
Generic Indicator	-0.136	0.307	0.610	0.530
	(0.020)	(0.181)	(0.112)	(0.227)
Year Fixed Effects	Yes	Yes	Yes	Yes
Drug Fixed Effects	Yes	Yes	Yes	Yes
N	3,881	3,881	3,881	3,881
Mean Pre	0.229	3.894	1.462	4.815

Notes: Table shows estimates of β in Equation (2.4). Each column shows a different count of combination trials involving a drug: those run by the original firm that patented the drug (Column 1), publicly-funded researchers (Column 2), other firms (Column 3) and total trials (Column 4). Robust standard errors are in parentheses. Additional robustness checks are included in Appendix B.2.

the period more than five years before generic entry, an 11.0% increase. The original owner runs fewer (-59.4% decrease) and publicly-funded innovators and other firms run more (7.9% and 41.7% increase, respectively). The increase in total trialing is driven by an increase in the probability of a newly generic drug being trialed with an on-patent drug, as shown in Table 2.4.

This fact clarifies that the ownership status of a given drug is critical not only for the original owner's clinical trial decisions, but also for other firms' and public innovators' decisions. Public crowd-out does not provide an explanation for this fact since overall more private trials are run with a drug after generic entry. The timing of combination innovation is shaped by the existing intellectual property protections, and inefficient investment by firms might result in delayed arrival of life-saving combination therapies to market.

To summarize, these facts demonstrate that firms are generally biased away from combination therapy trials, and in particular those trials that involve other firms' branded drugs. This is consistent with the stylized model, which together with these facts suggests inefficient under-investment by firms in clinical trials for combination therapies because of market expansion externalities on other firms and positive spillovers to patients. We next quantify these externalities using an empirical model of demand and pricing of cancer drugs.

Table 2.4: Combination Trials and Generic Entry Regression

	With All Generic Total	With at Least One Brand Total
	(1)	(2)
Generic Indicator	0.028	0.502
	(0.038)	(0.208)
Year Fixed Effects	Yes	Yes
Drug Fixed Effects	Yes	Yes
N	3,881	3,881
Mean Pre	0.490	4.325

Notes: Table shows estimates of β in Equation (2.4) where Y_{dt} is either total combination trials with all generic drugs (i.e., all *other* drugs in the combination are generic) (Column 1) or total combination trials with at least one branded drug (Column 2). Robust standard errors are in parentheses.

2.4 Static Model of Drug Consumer Surplus and Profits

In this section, we develop and estimate an empirical model of cancer drug demand and price setting, which we use to recover patient substitution patterns, consumer surplus, and firm profits. We estimate the model using microdata on patient drug usage and prices, and in Section 2.5 we use the estimated model to quantify the externalities associated with observed introductions of combination therapies.

2.4.1 Cancer Drug Demand

The main role of the demand model is to measure patient substitution patterns and willingness to pay for different regimens, which are important inputs into measuring the externalities from combination innovation. The model has two key features. First, it allows for complementarities between drugs by modeling demand over bundles of drugs (similar to Gentzkow, 2007), which we refer to as *regimens*, using insights from the medical literature (Chu and DeVita, 2019) to define these regimens. Allowing for complementarities across drugs is crucial for estimating the value of trialing combinations. Second, it allows for heterogeneity in substitution patterns by patient demographics, including type of insurer, which is important for predicting substitution patterns and price effects of combination innovation.

At each time t, each cancer patient j makes a joint decision with her doctor over what drug regimen r, either single or combination, to take from the set of available regimens $\mathcal{R}_t$.[47] The

[47]By modeling this choice as a joint decision between patients, doctors, and in some cases, insurers, when interpreting the model's implied welfare, we make the assumption that these other agents are trying to maximize

patient j has cancer $c \in \mathcal{C}$ and will consider taking only regimens recommended for that cancer, which is subset a $\mathcal{R}_{ct} \subseteq \mathcal{R}_t$ of available regimens.[48] The patient j also has an insurance type $\iota \in \mathcal{I}$, where the set of insurance types is defined as

$$\mathcal{I} = \{\text{Medicare, Medicare + Medicaid, Private}\}.$$

Regimen r characteristics at time t for insurer ι are given by $(p_{rt\iota}, x_{rt}, \xi_{r\iota})$, where $p_{rt\iota}$ is the regimen price, x_{rt} is a vector of observable characteristics, and $\xi_{r\iota}$ is unobserved (to the econometrician) regimen quality.[49] Given that our demand system is over regimens, similar to Gentzkow (2007), this regimen quality term captures complementarity of the regimen components. Patient j's characteristics are given by $(z_{jt}, \nu_{jt}, \varepsilon_{jt})$, where z_{jt} is a vector of patient demographics, ν_{jt} is a patient unobservable, and ε_{jt} is a vector of patient-regimen preference shocks, where each element is distributed Type I extreme value. We suppress conditioning on cancer type c in what follows. We take t to be a month.

The utility u_{jrt} of patient j, who has insurer ι, from taking regimen $r \in \mathcal{R}_t$ at time t is given by the sum of a regimen-insurer mean utility term $\delta_{rt\iota}$, a patient heterogeneity term μ_{jrt}, and a patient-regimen preference shock ε_{jrt},

$$u_{jrt} = \delta_{rt\iota} + \mu_{jrt} + \varepsilon_{jrt}.$$

The regimen-insurer mean utility $\delta_{rt\iota}$ is given by

$$\delta_{rt\iota} = \alpha_\iota p_{rt\iota} + \xi_{r\iota} + \xi_{y(t)\iota} + \Delta\xi_{rt\iota},$$

where $\xi_{y(t)\iota}$ is a year fixed effect that captures changes in the quality of the outside option and $\Delta\xi_{rt\iota}$ is a structural error that captures unobserved demand shocks.[50] The patient heterogeneity component of utility μ_{jrt} is given by

$$\mu_{jrt} = \theta^z_{\iota 1} a_{jt} p_{rt\iota} + \theta^z_{\iota 2} a_{jt} \mathbf{1}_{r \text{ biologic}} + \theta^z_{\iota 3} a_{jt} \mathbf{1}_{r \text{ combo}} + \theta^z_{\iota 4} \nu_{jt} p_{rt\iota},$$

patient utility. In reality, the utility function is some combination of patient, doctor, and insurer utility. This assumption is common in work estimating demand for pharmaceuticals (e.g., Dubois and Lasio, 2018).

[48] In cases where a patient has more than one type of cancer, she will appear in each cancer's demand system. Generally, a patient just has one type of cancer in any given year.

[49] We allow regimen quality to differ across insurers as variation in cancer incidence and patient demographics among their respective populations may influence the effectiveness of specific regimens.

[50] Important examples of demand shocks we do not explicitly model include drug shortages (negative shocks), which are becoming increasingly common for certain injectable drugs (Yurukoglu et al., 2017), and marketing of cancer drugs (a positive shock) to physicians (Carey et al., 2024).

where age a_{jt} is an element of z_{jt}, $\mathbf{1}_{r \text{ combo}}$ is an indicator of whether regimen r is a combination regimen (rather than single-drug), $\mathbf{1}_{r \text{ biologic}}$ is an indicator of whether r contains at least one biologic drug,[51] and ν_{jt} is a random coefficient on price. Throughout, prices are normalized to be in thousands of dollars, and age is divided by 100.

We include a random coefficient ν_{jt} on price in the patient heterogeneity term to capture unobserved supplemental insurance, heterogeneity in insurance schedules, and other heterogeneity in regimen prices. We specify ν_{jt} to be a Bernoulli random variable, which takes the value of 1 with probability $\psi^{\nu}_{t\iota}$.[52] This probability is set to the average fraction of patients with low (instead of high) cost-sharing. For Medicare and dual-enrolled Medicare patients, $\psi^{\nu}_{t\iota}$ is the probability of having supplemental insurance, derived from annual aggregate data from the Medicare Current Beneficiary Survey (MCBS). For privately insured patients, $\psi^{\nu}_{t\iota}$ is the average fraction of patients who have reached their out-of-pocket maximum.[53]

The coefficient α_ι and patient heterogeneity coefficients $\theta^z_{\iota 1}$ and $\theta^z_{\iota 2}$ allow price sensitivity to vary by insurance type.[54] This is motivated by the different out-of-pocket costs associated with different types of insurance and potential heterogeneity in price sensitivity that occurs through the insurer (via e.g., prior authorization) or doctor. Though some Medicare patients with supplemental insurance and privately insured patients may have low out of pocket costs for much of their drug usage, price sensitivity may also occur through doctors or insurers. For example, many chemotherapy and biologic drugs require prior authorization. Nevertheless, there is evidence that many patients face high out-of-pocket costs that may affect treatment decisions.[55]

The patient heterogeneity term allows patient substitution patterns to vary by patient charac-

[51] We do not include the individual terms (i.e., not interacted) $\mathbf{1}_{r \text{ combo}}$ and $\mathbf{1}_{r \text{ biologic}}$ as they are absorbed in the regimen fixed effect. Similarly, the mean effect of price is already included in $\delta_{rt\iota}$. Finally, we omit the mean effect of age.

[52] This specification is similar to "discrete-type" random coefficients models (e.g., Berry and Jia, 2010).

[53] In 2022, approximately 89% of Traditional Medicare beneficiaries had some form of supplemental coverage (Ochieng et al., 2024). In the Marketscan data, approximately 80% of claims for cancer drugs list the patient as already having hit her out-of-pocket maximum at the time of the claim.

[54] Note that the price sensitivity coefficient α_ι can represent $\tilde{\alpha}_\iota \times \zeta_\iota$, where ζ_ι is the fraction of drug costs (i.e., from co-insurance, co-payments, or deductibles) that the patient pays.

[55] Narang and Nicholas (2017) documents significant out-of-pocket costs for Medicare patients with cancer for a range of types of supplemental insurance using data from the Health and Retirement Study: "$2116 among those insured by Medicaid, $2367 among those insured by the Veterans Health Administration, $5976 among those insured by a Medicare health maintenance organization, $5492 among those with employer-sponsored insurance, $5670 among those with Medigap insurance coverage, and $8115 among those insured by traditional fee-for-service Medicare but without supplemental insurance coverage." Furthermore, these OOP costs were on average 23.7% of household income. More generally, the "financial toxicity" of cancer is a growing concern (Smith et al., 2022). The American Cancer Society Cancer Action Network (2024) reports 47% of patients surveyed in their 2024 Survivor Views research panel have had medical debt related to their cancer, and 25% of patients delayed or skipped care to avoid further debt.

teristics and regimen characteristics, reflecting clear patterns in the data. For example, in the Medicare data, older patients tend to receive less expensive drugs, are less likely to receive regimens containing biologic drugs, and have a lower likelihood of taking combination therapies, even when controlling for cancer type. We estimate the parameters in the patient heterogeneity term by matching micro moments (discussed in detail in Section 2.4.4).

We define the outside option to be taking a drug regimen that does not appear in treatment guidelines. We discuss the construction of the outside option in Section 2.4.3. With this outside option, we take the decision to treat a patient with drugs rather than other means, such as surgery or radiation, as exogenous.[56] We normalize the utility of the outside option as $u_{i0t} = \varepsilon_{i0t}$.

The share of patients taking regimen r at time t conditional on insurance type ι is

$$s_{rt\iota}(p_{rt\iota}) = \int_{\nu}\int_{z} \frac{\exp\left(\delta_{rt\iota} + \mu_{jrt}\right)}{1 + \sum_{r' \in \mathcal{R}_t}\left(\exp\left(\delta_{r't\iota} + \mu_{jr't}\right)\right)} dz_{jt}\, d\nu_{jt}.$$

Let $M_{ct\iota}$ be the market size cancer type c patients at time t with insurance type ι and let M_{ct} be the market size of cancer type c at time t, summing over insurance segments. We describe the calculation of market size, regimen market shares, and regimen prices in detail in Section 2.4.3.

2.4.2 Cancer Drug Price Setting

The price setting model serves two key purposes. First, we use the model to estimate marginal costs of drug production, which are important inputs into calculating firm profits and externalities from innovation. Second, we estimate price setting conduct that enables predicting counterfactual drug prices after combination innovation.

We model drug price setting as simultaneous bilateral Nash bargaining between manufacturers and a single insurer in each year.[57] This single insurer serves as a stand-in for the private insurance market. A bargaining model captures the role of the insurer as an intermediary in determining drug prices for its relatively price inelastic beneficiaries. We assume gross drug prices p_t are bargained each year and that this gross price applies to both publicly and

[56] A larger fraction of patients are taking cancer drugs over time (instead of exclusively other treatment modalities like radiation or surgery), so our analysis likely underestimates this market expansion effect of new drug innovations. Appendix B.1 presents these trends. Accounting for an increasing fraction of patients taking drugs in the analysis would require us to model innovations in other treatment modalities, which is beyond the scope of this paper.

[57] We abstract from dynamic pricing concerns, including those created by Medicare Part B's lagged-price reimbursement contracts (Acquatella et al., 2023).

privately insured patients. Indeed, since 2005, maximum reimbursement rates for drugs covered by Medicare Part B are set to be the Average Sales Price (ASP) of the drug plus 6%, and thus the government does not directly negotiate with manufacturers over drug prices for its beneficiaries.

Prices p_t are the gross drug prices for privately insured patients, and individual patients may pay different amounts to receive the drug because of different insurance benefits. The difference between the gross price and the patient price contributes to insurer costs. Some of these costs are paid by the government (e.g., the cost after coinsurance for Medicare patients) and some are paid by private insurers (e.g., those for privately insured non-Medicare beneficiaries).

The surplus of the private insurer, denoted by ι_{private}, in the bargaining problem is a weighted sum of consumer surplus for its beneficiaries and drug costs paid by the insurer (following, e.g., Capps et al., 2003, Gowrisankaran et al., 2015, Dafny et al., 2023):[58]

$$\mathrm{V}_{t\iota_{\text{private}}}(\mathcal{R}_t, p_t) = \rho \sum_{c \in \mathcal{C}} \mathrm{CS}_{t\iota_{\text{private}}}(\mathcal{R}_t, p_t, c) - \sum_{c \in \mathcal{C}} \mathrm{TC}_{t\iota_{\text{private}}}(\mathcal{R}_t, p_t, c),$$

where ρ is the relative weight the insurer places on consumers relative to its own costs. The idea behind this objective function is that the insurer sets premiums to extract the expected difference between surplus for its consumers net of insurer costs in a first-stage problem that we do not model.

The weight ρ, similar to the welfare weight in Gowrisankaran et al. (2015), may be different than 1 for a number of reasons. For example, if consumers underestimate their drug costs (Abaluck and Gruber, 2011) or overestimate their surplus from taking drugs (e.g., overestimating the probability of requiring chemotherapy), then ρ may be greater than 1. We microfound these reasons in Appendix B.4.

Consumer surplus for the private insurer's patients with cancer c given regimens $\mathcal{R}_{ct}$ and prices p_t at time t is given by

$$\mathrm{CS}_{t\iota_{\text{private}}}(\mathcal{R}_t, p_t, c) = \\ -M_{ct\iota_{\text{private}}} \int_{\nu} \int_{z} \frac{1}{\alpha_{\iota_{\text{private}}} + \theta^{z}_{\iota_{\text{private}}1} a_{jt} + \theta^{z}_{\iota_{\text{private}}4} \nu_{jt}} \log\left(1 + \sum_{r \in \mathcal{R}_{ct}} \exp\left(\delta_{rt\iota_{\text{private}}} + \mu_{jrt}\right)\right) dz_j \, d\nu_j.$$

[58] We assume the consumer surplus of privately insured patients is the sum of two components. First, the consumer surplus of patients under the age of 65, which comes from Marketscan. Second, the consumer surplus of patients enrolled in Medicare Advantage. We do not have data on the claims of these patients, so we assume the demand estimates for this group are the same as Traditional Medicare beneficiaries. Our Marketscan data end in 2013, so we assume the privately insured demand primitives remain constant after this date.

Total insurer costs for the private insurer's patients with cancer c are

$$\text{TC}_{t\iota_{\text{private}}}(\mathcal{R}_t, p_t, c) = M_{ct\iota_{\text{private}}} \sum_{r \in \mathcal{R}_{ct}} (1 - \zeta_{\iota_{\text{private}}}) p_{rt} s_{rt\iota_{\text{private}}}(p_t),$$

where $\zeta_{\iota_{\text{private}}}$ is the expected patient payment fraction (i.e., combining any co-payments, co-insurance, or deductibles) from taking a regimen.

For each firm $f \in \mathcal{F}$, let $f_t \subset \mathcal{D}$ denote the set of drugs owned by the firm at time t. Firm f's profits at time t are

$$\pi_{ft}(p_t) = \sum_{\iota} \sum_{c \in \mathcal{C}} M_{ct\iota} \sum_{d \in f_t} \sum_{r \in \mathcal{R}_{ct}} (p_{dt} - mc_{dt}) s_{rt\iota}(p_t) \mathbb{1}[d \in r],$$

where we assume that firms have constant marginal costs of production mc_{dt} for each of their drugs. This profit function sums across profits from each insurance segment ι, which includes both privately and publicly insured patients. We assume that a firm only "owns" a drug when it is currently patented, and we do not include generic price setting in the model. Regimens including generic drugs are still in the demand system, but the price of a generic drug is fixed at the average price we observe for that drug in the data each year.[59]

The drugs in our sample can be classified into two types: small molecules and biologic drugs. Small molecules are manufactured from chemicals and generally considered to have very low marginal costs of production. Biologic drugs are generally manufactured from living cells, and the costs associated with production can be substantial. We will assume that marginal costs are zero for all small molecule drugs, but we allow for and estimate positive marginal costs for biologic drugs.[60,61]

The prices p_{ft} of drugs produced by firm f at time t satisfy

$$\max_{p_{ft}} \quad \pi_{ft}(p_t)^{\gamma} \left[\text{V}_{t\iota_{\text{private}}}(\mathcal{R}_t, p_t) - \text{V}_{t\iota_{\text{private}}}\left(\mathcal{R}_t \setminus \mathcal{R}_{ft}, p_t\right) \right]^{1-\gamma}.$$

The parameter γ gives the bargaining weight of a firm relative to the insurer. The insurer's outside option is the value when drugs from a given manufacturer are not available, holding all remaining drug prices fixed but allowing demand to adjust. The manufacturer's value is simply

[59] While we do not model generic pricing, we allow generic drug prices to adjust exogenously as the number of generic competitors changes.

[60] There is some evidence (Hill et al., 2016) that even relatively expensive small molecule drugs like small molecule inhibitors have relatively low marginal costs of production.

[61] We do not consider advertising (such as detailing) to be a marginal cost. Both types of drugs may accrue significant costs from detailing, though whether these costs are marginal is unclear.

total profits, and the manufacturer's outside option is zero.[62]

2.4.3 Market Shares and Prices

We measure cancer regimen usage and prices by combining data from Medicare (for publicly insured patients) and Marketscan (for privately insured patients), as described in Section 2.2.2. We use data from both publicly and privately insured patients for three reasons. First, in order to estimate the value of combination innovation, we must estimate a demand system for cancer drugs that is representative of the target market, which we define as the US population. The Medicare data contain relatively few individuals under the age of 65, and may not capture drug usage patterns for cancers that skew younger or have different recommended treatments for younger patients. Second, our model of cancer drug price setting must match important institutional features. Prices for drugs in Medicare Part B are set based on average national sales prices of the drugs, including sales to privately insured patients. Third, a concern with estimating price elasticities of demand from the Medicare data is the presence of unobserved supplemental insurance (Medigap). We allow for unobservables to affect price sensitivity, but use the privately insured patients to benchmark the price elasticities we estimate in the Medicare data.

To identify patients with a particular cancer in a year in the Medicare data, we combine claims from the inpatient, outpatient, and carrier (professional provider service) fee-for-service files. For the Marketscan data, we combine claims from the outpatient, inpatient services, and inpatient admissions files. We classify a patient has having a particular cancer in a given year if she has at least one claim (line) with a primary diagnosis code of that cancer, where we make a crosswalk between diagnosis codes (ICD9 and ICD10) and types of cancers in Chu and DeVita (2019). The vast majority of cancer drugs are delivered in a clinical outpatient setting and thus covered by Part B for Medicare beneficiaries and medical benefits for privately insured patients (there is a small subset of regimens that include prescription drugs that would be included in Part D plans for Medicare beneficiaries or pharmacy benefits for privately insured patients).[63]

Calculating regimen usage is complicated by the fact that combination regimens are typically

[62] Our specification of the bargaining model assumes that all the profits associated with the drug go to the manufacturer, rather than dividing profits between manufacturers and providers that deliver the drugs (e.g., hospitals). Hospitals likely charge positive markups on drugs (Robinson et al., 2021, Robinson et al., 2024), and we abstract from the split of profits between hospitals and manufacturers.

[63] For Medicare beneficiaries, enrollment in Part D is voluntary and thus we do not observe usage of prescription drugs for patients who elect to not enroll in Part D. Furthermore, Part D was created in 2006, so we do not see usage for these drugs before 2006. The Marketscan data include prescription drug benefits in all years of our sample.

not packaged together (and therefore do not have their own billing codes, and instead must be identified via the usage of component drugs in microdata) and drugs in the regimen are taken over the course of many days or weeks. To assign patients to regimens, we use the following algorithm. For each year and cancer, we subset to patients who we identify to have that cancer in that year. We further subset to patients taking at least one of the drugs in a recommended regimen for that cancer in that year. The cardinality of this set of patients (scaled by the corresponding sampling weights) is the size of the market. We pick a number of days N, where $N = 30$ is our baseline assumption. For each patient and each day she takes a cancer drug, we make a list of all drugs she takes N days before or N days after the current day. We assign the patient to the largest (in terms of number of component drugs) regimen she takes during that window, if any.[64] For days the patient does not take drugs but are within N days before or after a day she does and is assigned to a particular regimen, we "fill in" these days with the closest (in terms of days) regimen.[65]

We calculate each regimen's market share at a monthly level, where the numerator is the number of patient-days taking a given regimen in a month, and the denominator is the number of patient-days taking any drug for that cancer in a month. A regimen's total price (summing patient and insurer payments) is the total spending on component drugs during the rolling window, averaged over patient-days assigned to that regimen. This regimen assignment procedure performs well, and patients take relatively few "extra" drugs during the rolling window (where $N = 30$ is our baseline assumption), less than .3 drugs on average, not included in their regimen (but included in other recommended regimens for that cancer).[66] The average share of the inside option (taking one of the recommended regimens) is 0.77 for the Medicare dataset and 0.76 for the Marketscan dataset.

There is both observed and unobserved regimen price dispersion across patients. In the Medicare and Marketscan datasets, there is observed priced dispersion across patients because of variation in quantities of component drugs taken by patients and provider markups.[67] For privately insured

[64] For example, suppose the set of recommended regimens for a cancer is given by $\{\{A,B,C\},\{A,D\},\{A,C\}\}$. Suppose the patient takes drugs A,B,C within N days of a date τ. The algorithm would assign the patient to regimen $\{A,B,C\}$ (not regimen $\{A,C\}$) on date τ. If the patient takes drugs A,B,C,D within N days of date τ the algorithm would still assign the patient to regimen $\{A,B,C\}$ at date τ, denoting drug D as an "extra drug."

[65] More precisely, suppose a patient takes drugs for regimen r and is thus assigned to regimen r on dates she takes the component drugs. Suppose these dates are January 20th and 25th. The algorithm would "fill in" dates and assign the patient to take regimen r for N days before January 20th, between January 20th and 25th, and for N days after January 25th. If a patient is assigned to more than one regimen on different dates, the algorithm fills in the remaining dates using the closest (in terms of days) regimen.

[66] See Appendix B.1 for additional details.

[67] One important source of unmodeled price dispersion is the 340B drug pricing program (see Levengood et al., 2024 for a review of its effects on drug access and providers), which provides certain healthcare organizations (those serving many uninsured or low-income patients) with discounts on outpatient drugs.

patients, there is observed price dispersion in the patient's financial responsibility (the sum of deductibles, co-pays, and co-insurance). For Medicare beneficiaries, there is additional price dispersion because of Medicaid dual enrollment and private supplemental insurance (Medigap) plans, the latter of which we do not observe. In our analysis, we abstract away from price dispersion due to differences in quantities of component drugs and provider markups by considering average prices across all patients, and we account for both observed and unobserved price incidence on patients in our demand model.

We obtain patient characteristics (e.g., age) from the enrollment files in each dataset. We use aggregate data from the Medicare Current Beneficiary Survey to calculate rates of supplemental insurance over time.

We define the market size of a particular insurance segment, cancer type, time to be the total number of patients in the US taking drugs for that cancer at time t. We scale the number of patients observed in the Medicare data to match the total population of traditional Medicare and Medicare Advantage patients, and we scale the number of patients observed in the Marketscan data to match the remainder of the US population.

Table 2.5 presents summary statistics of drug usage, prices, and patients for the Medicare and Marketscan data. We include cancers with at least one thousand Traditional Medicare patients in our sample taking drugs for that cancer per year on average, leaving 19 cancers. The average total market size (summing across all cancers) is 823,839 Medicare patients and 549,694 Marketscan patients, with an average 43,360 Medicare patients per cancer and 28,931 Marketscan patients per cancer. Additional summary statistics about patient regimen usage are presented in Appendix B.1.

2.4.4 Estimation and Identification

Demand: We estimate the demand model using a generalized method of moments (GMM) estimator, combining aggregate sample and micro moments, following the method of Berry et al. (1995) and best practices in Conlon and Gortmaker (2020, 2023). We use a nested fixed point algorithm to optimize over nonlinear parameters governing patient heterogeneity, and we concentrate out the linear parameters. We estimate the demand model separately for each insurance type. We approximate the distribution of patient demographics z_i by creating age bins by insurance type.[68]

[68]We do not impose that the regimen fixed effects are the same across insurance types. It is possible that the quality of the regimen for a patient depends on patient demographics, so allowing the regimen fixed effect to vary across insurance type captures this heterogeneity. In the dynamic model in Section 2.6, we will make a simplifying assumption that regimen quality is drawn from a distribution that represents average quality over

Table 2.5: Cancer Drug Usage, Price, and Patient Summary Statistics

Cancer	Patients			Shares		Regimens		Price
	Total	Drug Fraction	Market Size	Inside	Combo	Total	Combo	$
Medicare								
Prostate	202,962	0.2	293,025	0.97	0.03	26	11	942
Colorectal	80,160	0.17	76,634	0.83	0.64	35	27	7,127
Head and Neck	73,280	0.05	21,252	0.59	0.19	20	12	950
Bladder	52,901	0.09	25,986	0.44	0.3	20	12	1,268
Breast	48,581	0.14	41,656	0.94	0.14	63	42	2,997
Non-Hodgkin Lymphoma	47,019	0.33	94,912	0.61	0.09	33	18	6,842
Non-Small Cell Lung	32,537	0.33	67,184	0.82	0.43	47	28	3,561
Malignant Mesothelioma	23,244	0.2	26,302	0.39	0.12	12	7	2,440
Chronic Lymphocytic Leukemia	20,030	0.13	15,360	0.97	0.19	22	10	6,179
Hepatocellular	18,235	0.24	24,370	0.61	0.17	10	3	2,968
Endometrial	17,240	0.11	9,399	0.74	0.45	15	10	720
Multiple Myeloma	16,358	0.2	15,049	0.96	0.24	27	19	9,814
Biliary Tract	16,032	0.27	26,286	0.42	0.06	8	5	1,787
Brain	14,521	0.16	12,391	0.39	0.03	9	4	3,710
Hodgkin Lymphoma	14,192	0.22	17,554	0.52	0.04	14	6	7,891
Ovarian Epithelial	13,741	0.32	26,336	0.89	0.47	24	11	2,456
Pancreatic	11,854	0.25	14,817	0.93	0.24	12	9	1,658
Gastric	7,981	0.19	9,054	0.67	0.46	18	14	806
Esophageal	4,005	0.29	6,272	0.62	0.25	9	7	349
Mean (patient weighted)	37,625	0.19	43,360	0.77	0.21	27	16	3,001
Marketscan								
Prostate	55,842	0.08	42,481	0.99	0.04	25	12	1,437
Colorectal	31,978	0.25	69,599	0.87	0.63	32	26	13,416
Head and Neck	31,160	0.06	21,292	0.6	0.23	18	12	2,751
Bladder	15,282	0.07	9,617	0.49	0.4	15	12	5,125
Breast	48,604	0.19	83,918	0.98	0.16	57	38	4,902
Non-Hodgkin Lymphoma	32,120	0.29	81,790	0.45	0.13	28	17	7,922
Non-Small Cell Lung	13,092	0.4	41,438	0.87	0.51	33	21	6,195
Malignant Mesothelioma	11,948	0.23	24,252	0.35	0.13	10	7	5,360
Chronic Lymphocytic Leukemia	5,909	0.15	8,202	0.97	0.31	19	8	3,299
Hepatocellular	9,182	0.31	25,460	0.68	0.07	8	3	6,683
Endometrial	11,095	0.1	9,778	0.82	0.34	15	10	2,077
Multiple Myeloma	5,714	0.18	9,215	0.85	0.29	21	15	7,141
Biliary Tract	8,478	0.38	27,286	0.65	0.11	9	6	6,096
Brain	12,542	0.16	16,895	0.73	0.07	10	5	7,265
Hodgkin Lymphoma	15,790	0.2	28,606	0.24	0.12	10	4	4,867
Ovarian Epithelial	12,336	0.23	25,267	0.91	0.48	22	11	5,242
Pancreatic	4,242	0.31	11,683	0.95	0.29	12	9	5,984
Gastric	3,502	0.29	8,665	0.74	0.52	18	16	4,886
Esophageal	1,568	0.37	4,250	0.64	0.41	10	9	3,251
Mean (patient weighted)	17,389	0.19	28,931	0.76	0.23	26	17	5,411

Notes: Table shows summary statistics for Medicare patients (top panel) and Marketscan patients (bottom panel) by cancer for the 19 cancers included in our sample. The table is sorted in descending order of the median (over years) number of Medicare patients we observe with that cancer. Patients Total is the median (over years) number of patients we observe in each dataset that have a diagnosis code for a particular type of cancer in a year. The Medicare data is a 20% sample of beneficiaries, and we observe claims for traditional Medicare beneficiaries. The number of individuals included in the Marketscan dataset grows considerably during the sample period, as discussed in the main text. Drug Fraction is the median (over years) fraction of those patients that receive at least one cancer drug in that year. Market Size is the median (over years) market size we consider for each cancer, defined to be the number of patients taking drugs for that cancer. We calculate this measure by subsetting to patients who take drugs for that cancer and apply sampling weights. Shares Inside is the median (over months and years) sum of market shares of recommended regimens in Chu and DeVita (2019), where the market share for a regimen is calculated for each month as the number of patient-days (within a particular month) assigned to that regimen divided by the total number of patient-days taking drugs (for a particular cancer-month). Combo is the median (over months and years) sum of market shares of recommended combination regimens in Chu and DeVita (2019). Regimens Total is the total number of recommended regimens we observe taken for that cancer, and Combo is the number of combination recommended regimens we observe taken for that cancer. Price is the median (over regimens, time windows, and years) total price, in dollars, of taking a regimen for that cancer for the N ($= 30$) day rolling window. The last row in each panel gives a patient weighted mean (except for total patients) of each column.

We use two types of instruments for regimen prices. The first is a baseline price-weighted average of the number of manufacturers producing drugs in the regimen interacted with indicator variables for the baseline revenue share bin of the regimen. More precisely, this instrument b_{rt}^{D1} for regimen r and time t is given by

$$b_{rt}^{D1} = \frac{\sum_{d \in r} p_{d1} n_{dy(t)}}{\sum_{d \in r} p_{d1}} \times h_1(r),$$

where p_{d1} is the (median) price of drug d in the first period it appears in our data (either the first period of our data if the drug was already being marketed, or the price the first year it is marketed), $n_{dy(t)}$ is the number of manufacturers that produce drug d in year $y(t)$, and $h_1(r)$ an indicator for the baseline bin of the revenue share of the regimen (calculated across all cancers), where we use three bins in our baseline specification. Intuitively, an increase in the value of the instrument suggests greater generic entry at t, which we expect to reduce the price p_{rt} of the regimen r. We interact with baseline regimen share bin to allow the effect of generic entry to vary flexibly based on size (determined as baseline revenue share) of the regimen.

The second instrument is a baseline price-weighted average of the time relative to patent expiry (measured in years) of drugs in the regimen interacted with indicator variables for the baseline revenue share bin of the regimen. Time before patent expiry is coded as negative values. More precisely, this instrument b_{rt}^{D2} for regimen r and time t is given by

$$b_{rt}^{D2} = \frac{\sum_{d \in r} p_{d1} g_{dt}}{\sum_{d \in r} p_{d1}} \times h_1(r),$$

where g_{dt} is the time since generic entry time of drug d at time t.[69] The idea for this instrument is similar to the first.

To target the interactions between observable patient characteristics and regimen characteristics, we match the covariance between (i) regimen prices and patient age (ii) regimen combination status and patient age (iii) regimen biologic status and patient age. Micro moment targets and patterns are discussed in Appendix B.3.

Price Setting: Given estimates of patient demand for cancer treatment regimens, we then estimate the price setting model. We specify the marginal costs of biologic drug d at time t

different insurance types. We find that the fixed effects are highly correlated, making this assumption reasonable given the patterns in the data (see Appendix B.3). An alternative assumption would be to estimate the demand model jointly for different insurance types and impose common regimen fixed effects.

[69] We allow g_{dt} to be negative before generic entry rather than bottom coding it at 0 for all values of t before generic entry as this captures a trend in the price (often increasing) before generic entry. For example, suppose drug d had generic entry in 2005. Then g_{dt} is -1 in 2004, 0 in 2005, and 1 in 2006.

as $mc_{dt} = \overline{mc}_d + \eta_{dt}$, where $\overline{mc}_d$ is the baseline cost of drug d and η_{dt} is a structural error unobserved by the econometrician that represents changes in production costs over time (e.g., changes in input prices, shocks to manufacturing). We set $mc_{dt} = 0$ and $\eta_{dt} = 0$ for all non-biologic drugs.[70] We also assume that there is one bargaining weight γ for all drugs. Therefore, the parameters to be estimated are γ, ρ, and one baseline marginal cost per biologic drug.

This setting has the standard identification problem of bargaining models with unobserved costs. For example, observing high prices could be either because costs are high and the manufacturer bargaining weight is low, or because costs are low and the bargaining weight is high. To separately identify the bargaining weight, costs, and consumer surplus weight, we follow the outline of the identification argument in Lee et al. (2021). We construct instruments that shift demand (and are uncorrelated with demand except through the effects on prices) and use them as surplus shifters. Movement in prices in response to these shifters will be informative about the bargaining weight and costs. For example, suppose there is a positive shock to demand for a particular drug. If the insurer has all the bargaining power, the negotiated drug price will equal marginal cost, and there will not be a price change in response to the positive demand shock. If instead the manufacturer has all the bargaining power, then we expect to observe increases in price.

We construct drug demand shifts for a focal drug d by calculating how many manufacturers are marketing *other* drugs used in regimens containing the focal drug. More precisely, we construct this surplus shifter b_{dt}^{S1} for drug d at time t as

$$b_{dt}^{S1} = \sum_{r \in \mathcal{R}_{ct} : d \in r} \frac{\sum_{d' \in r \setminus \{d\}} p_{d'1} n_{d'y(t)}}{\sum_{d' \in r \setminus \{d\}} p_{d'1}},$$

using the same data on baseline prices and number of generic entrants as in the demand instrument construction. Intuitively, an increase in the instrument suggests generic entry in drugs that are complementary to d, likely lowering their prices $p_{d't}$ and hence increasing demand and raising the surplus to be split between the manufacturer of d and the insurer.

We simplify estimation by optimizing only over the bargaining weight γ and consumer surplus weight ρ. For biologics, at each possible value of γ and ρ, the vector of biologic costs (mc_{dt}) can be solved in closed form via the bargaining first-order conditions. Assuming the error terms (η_{dt}) are mean independent of the baseline costs allows us to calculate the baseline costs $\overline{mc}_d$

[70] An alternative assumption we could make is to assume mc_{dt} is the lowest price we see in the data once the drug has generic entry. For small molecules that never have generic entry during our sample period, we could calculate the average markup for small molecules that have generic entry and set costs to match this average markup.

for each drug d in closed form. Then, we can calculate moments of the form $\mathbb{E}\left[\eta_{dt} B^S_{dt}\right] = 0$ and estimate via GMM, while only optimizing over γ and ρ. Here, B^S_{dt} is a vector of instruments we use for drug d, where $B^S_{dt} = b^{S1}_{dt} \times \mathbf{1}_{d'=d}$ (i.e., interacted with an indicator for each drug).

2.4.5 Estimation Results

Demand: Table 2.6 shows estimated parameters and summary statistics of own-price elasticities of demand. Each column presents estimates for a different insurance segment: Medicare (not dual-enrolled), Medicare dual-enrolled in Medicaid, and privately insured (Marketscan). With regimen price instruments, we estimate a median own price elasticity of demand of -1.16 for Medicare, -0.17 for Medicare dual-enrolled, and -2.49 for Marketscan. Previous papers estimating Medicare patient demand for cancer drugs report own-price elasticities of demand ranging between $-.7$ and -2.7 (Jung et al., 2017, Song et al., 2017),[71] and our median elasticity for Medicare patients falls within this range.

The ranking of these elasticities across insurance types is what we would expect based on the price incidence of drugs on patients and the price sensitivity of insurers. Medicare dual-enrolled patients face little to no cost-sharing, and we estimate that they are the most price inelastic group. Medicare patients likely face the highest cost-sharing depending on supplemental insurance but do not face prior authorization for drug usage. And while privately insured patients in our sample typically hit the out-of-pocket maximum relatively quickly after diagnosis, prior authorization is often required, so that the relative elastic demand curve reflects the price sensitivity of the insurer.

Similar to Gentzkow (2007), regimen fixed effects are informative about the complementarity of component drugs. We find that combination regimens often have higher fixed effects than single-agent regimens of the components. These patterns are detailed in Appendix B.3.

The model's implied micro moments are shown in Rows 15-18. The model fit is reasonable, with the model implied micro moments having a similar magnitude to the targets. Appendix B.3 presents additional estimation results.

Price Setting: We estimate a bargaining weight of 0.69 and a consumer surplus weight of 8.04. These parameter estimates are similar to those estimated in other health contexts. Dafny et al. (2023) estimates a bargaining weight of .69 in a similar bargaining problem over the prices of drugs for multiple sclerosis. Gowrisankaran et al. (2015) estimates a weight on consumer surplus of $2.79 - 6.69$, depending on the specification, in a similar bargaining problem over

[71] Jung et al. (2017) estimate price elasticities for cancer drugs covered under Medicare Part D. Song et al. (2017) estimate price elasticities for colorectal combination therapies.

hospital prices.[72] For these parameters, the mean markup over marginal cost for biologic drugs is .25, with a standard deviation (across drugs and years) of .44. While there is limited information about the accounting manufacturing costs of biologic drugs, estimates of the price declines after biosimilar entry range from 1-25%, suggesting our estimated markups fall in a reasonable range.[73] We present robustness checks of the estimated bargaining weight and implied markups to different values of ρ in Appendix B.4.

2.5 Externalities from Combination Innovation

This section estimates externalities from successful drug regimen innovation, comparing the distinctive externalities that arise from combination innovation to externalities present in standard innovation settings. To calculate the externalities for regimens introduced in the data, we augment the stylized model in Section 2.3.1 to incorporate additional details of our setting and parameterize it using the demand and price setting models estimated in Section 2.4.

2.5.1 Measuring Externalities after Regimen Introduction

We start by outlining how we measure innovation externalities after regimen introduction, extending the stylized model in Section 2.3.1 to incorporate fiscal externalities on the insurer. Suppose firm f introduces new regimen r^+ at time t. Since consumers do not pay the full price of treatment, we must add the change in total costs of the insurer (ΔTC) to the changes in consumer surplus and total profits to recover the change in welfare at time t:[74]

$$\Delta W = \underbrace{\text{CS}\left(\mathcal{R} \cup r^+, \tilde{p}\right) - \text{CS}\left(\mathcal{R}, p\right)}_{\Delta\text{CS}} - \underbrace{\left[\text{TC}\left(\mathcal{R} \cup r^+, \tilde{p}\right) - \text{TC}\left(\mathcal{R}, p\right)\right]}_{\Delta\text{TC}} + \sum_{d \in \mathcal{D}} \underbrace{\pi_d\left(\mathcal{R} \cup r^+, \tilde{p}\right) - \pi_d\left(\mathcal{R}, p\right)}_{\Delta\pi_d}. \tag{2.5}$$

As in the stylized model, $\tilde{p}$ is the observed price vector after introduction of the new regimen r^+, while p is the counterfactual price vector without r^+. Firm f internalizes only the change

[72] Our estimate falls within the confidence interval of their estimate of 6.69 with standard error 5.53.

[73] Price et al. (2015) reports European manufacturers expending between $100 million to $250 million to reverse engineer biologic drugs, with prices dropping by about 25% after biosimilar entry. In the US, estimated price declines range from 5-9% percentage points (Frank et al., 2022, Stern et al., 2021). Compared to the US, competition resulting from biosimilar entry in Europe is much stronger. For comparison, after generic competitors enter, small-molecules may see a decline in price of over 95% (Conrad and Lutter, 2019).

[74] Throughout this section, we suppress the dependence of all values on t.

Table 2.6: Demand Model Estimates

	Medicare	**Medicare Dual**	**Marketscan**
	(1)	(2)	(3)
α_ι	-0.489	-0.311	-0.628
	(0.012)	(0.006)	(0.047)
θ_1^z (age × price)	0.0867	0.2437	-0.0405
	(0.02389)	(0.01525)	(6.0e-5)
θ_2^z (age × biologic)	0.0039	0.0078	0.0101
	(0.00083)	(0.00059)	(0.00091)
θ_3^z (age × combo)	0.0014	0.0	-0.0182
	(0.00072)	(0.00028)	(0.0002)
θ_4^z (r.c. × price)	0.0122	0.0497	0.0028
	(0.00013)	(0.0002)	(0.00088)
Median Own Price Elasticity	-1.156	-0.175	-2.492
Median Own Price Elasticity (no price ins)	0.068	-0.076	0.114
Median Own Price Elasticity Logit	-1.1	-0.454	-5.402
Age Price Covariance	-0.0361	-0.0375	0.0006
Age Biologic Covariance	-0.0019	-0.0022	0.0008
Age Combo Covariance	0.0011	-0.0008	-0.0008
Regimen FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: Table shows parameters and summary statistics of the demand system, estimated separately for each insurance type in each column. Column (1) is traditional Medicare beneficiaries, Column (2) is traditional Medicare patients who are dual-enrolled in Medicaid, and Column (3) is privately insured patients in Marketscan. The α row shows the price sensitivity coefficient. Rows θ_1^z through θ_4^z are the patient heterogeneity coefficients. Rows Median Own Price Elasticity through Median Own Price Elasticity Logit are own price elasticities of demand for the baseline, no price instruments, and Logit specifications (using the same price instruments, but without demographic heterogeneity), respectively. Rows Age Price Covariance through Age Combo Covariance show the estimated micro moments. The final two rows indicate the inclusion of regimen and year fixed effects. Standard errors are in parentheses.

Table 2.7: Bargaining Model Estimates

Variable	Mean	Standard Error (Deviation)
Bargaining Weight γ	0.69	0.29
Consumer Surplus Weight ρ	8.04	4.41
Markup over Marginal Cost (Biologics)	0.25	0.44

Notes: Table shows parameter estimates and summary statistics of the bargaining model. The first row is the estimated bargaining weight. The second row is the estimated weight on consumer surplus in the insurer's objective function. Standard errors for each of these parameters are given in the second column. The final row computes the mean and standard deviation markup over marginal cost for biologic drugs.

in its own profits after regimen introduction, so the net externality is now

$$\text{Net Externality}_{\mathcal{R}^+} = \Delta\text{CS} + \Delta\text{TC} + \sum_{d \notin f} \Delta\pi_d. \tag{2.6}$$

This definition differs from (2.2) only because of the new insurer cost term.

The expressions above describe the externalities at the time of introduction t, but profit externalities will persist throughout the exclusivity periods of the affected drugs, while the consumer surplus and insurer cost externalities will persist until demand for the new regimens falls to zero.[75] Extending the calculation of these externalities requires assumptions on the future path of innovation, which we address with the dynamic model in Section 2.6. Here we focus instead on calculating each externality at the time of introduction, recognizing that this likely understates the magnitude of the full dynamic externality because of gradual adoption dynamics.

We observe 131 combination therapy introductions and 92 single-agent therapy introductions between 1999 and 2019.[76] Table 2.8 summarizes these events. Of the combination therapy introductions, 66 contain at least two branded drugs owned by different firms (at the time of introduction). The remaining 65 either consist of at least two branded drugs owned by the same firm or at most one branded drug (e.g., a branded drug in combination with a generic drug).

A significant fraction—85%—of trials for the combinations that are introduced are sponsored by public innovators. These publicly-funded trials are likely to generate positive profit externalities on firms (as we demonstrate below), but may overstate these externalities compared to a

[75]There will also be effects on generic drugs, but significantly smaller given that prices for generic drugs are substantially lower than branded drugs.

[76]There are more entry events that occur according to guidelines such as Chu and DeVita (2019), but not all regimens included the treatment guidelines are taken in the Medicare or Marketscan data.

counterfactual scenario where a firm conducts the trial itself (since the firm would own at least one drug in the regimen). To separately identify public spillovers and likely externalities arising from private innovation of combinations, we compute the effects of new combination introductions using two methods. For method 1, we calculate the externalities from new combination introductions from the perspective of each firm that would earn positive profits from the trial—that is, firms owning at least one branded drug in the combination. This approach allows us to leverage the publicly funded trials we observe to estimate externalities over a broad set of combinations. For method 2, we compare these results with estimates of externalities based on the actual innovator who conducted each trial in our data, analyzing publicly and privately funded trials separately. The former set is informative about public spillovers, while the latter set show realized externalities for the selected set of combinations that firms trial.

For each approach, we compute each term in the net externality definition (2.6) by introduction event. As in the stylized model, we decompose the net profit externality component into business stealing, market expansion, and price adjustments. We use the estimated demand and bargaining models to simulate these changes the year after entry of a particular regimen. We hold all other primitives constant.

As discussed in Section 2.3.1, several factors prevent us from calculating the regimen introduction externalities directly from the data. First, measuring the change in consumer surplus requires an estimated demand model. Measuring profit changes also requires marginal costs, which we estimate from the bargaining model. Second, there are often multiple regimen introductions in each year for a given cancer, which prevents us from observing the post-introduction prices that would prevail upon introduction of a single regimen. Additionally, idiosyncratic events including patent expiry or generic entry also affect drug prices at each time, again contaminating our observation of post-introduction prices. We use the bargaining model to compute post-introduction prices holding all other primitives constant.[77]

Finally, it is important to note that our estimates of the externalities are computed for the set of regimens that were ultimately introduced, which likely implies selection on being privately worth trialing for the innovator (whether public or private), among other unobserved factors.

Table 2.8: Average One-Year Externalities by Regimen Type

	Regimen Type		
	Single Drug	**≥ 2 Firms' Branded**	**Other Combinations**
	(1)	(2)	(3)
# Events	92	66	65
Fraction Firm Trials		0.15	0.15
Business Stealing + Market Expansion ($M)	-40.2	27.01	-16.77
Market Expansion ($M)	0.0	52.46	0.0
Business Stealing ($M)	-40.2	-12.57	-16.77
Price Adjustment ($M)	-1.9	-0.48	-1.12
Δ CS ($M)	92.77	31.02	35.73
Δ TC ($M)	16.22	23.7	19.79
Net Externality ($M)	6.04	34.65	3.13
Fraction Positive Net Externality	0.43	0.9	0.6
Fraction Positive BS + CS - TC	0.43	0.25	0.6
Firm Profit (Innovator) ($M)	58.18	32.58	26.83

Notes: Table shows summary statistics and estimates of externalities separately for single drugs (1), combinations with at least 2 firms' branded drugs at the time of introduction (2) and all other combination (3). Combinations of this third type include: combinations with at least two branded drugs owned by the same firm, or combinations that consist of at most one branded drug. Externalities and own profit effects are computed separately for each event and from the perspective of each firm that would have earned positive profits from the new regimen (i.e., firms with at least one branded drug in the regimen). All effects are calculated one year after the introduction of the new treatment. # Events is the number of introduction events. Fraction Firm Trials is the fraction of trials for that regimen type that were run by firms (rather than public innovators). Business Stealing + Market Expansion is the mean sum (in millions of dollars) of the business stealing and market expansion terms in the drug profit change decomposition from (2.1), summed over all drugs not owned by the innovating firm. Market Expansion and Business Stealing show these two terms separately. Price Adjustment is the mean change in drug profits due to the price adjustment term from (2.1), summed over all drugs not owned by the innovating firm. Δ CS is the mean change in consumer surplus. Δ TC is the mean change in insurer costs. Net Externality is the mean net externality (summing consumer surplus, insurer costs, and net profit externality). Fraction Positive Net Externality is the fraction of events with positive net externalities. Fraction Positive BS + CS - TC is the fraction of events with the sum of the business stealing externality, consumer surplus spillover, and insurer cost spillover being positive. Firm Profit (Innovator) is the mean profit change of the drugs owned by the innovating firm.

2.5.2 Estimates of Externalities after Regimen Introduction

Table 2.8 summarizes the externalities after single-drug and combination introduction, computed from the perspective of each firm that would have earned positive profits from that new regimen (firms with at least one branded drug in the combination), i.e., using method 1. We show effects

[77] We compare the model predicted market shares and realized market shares in the data after innovation events in Appendix B.5. These shares are highly correlated, but we use the model estimates as our baseline specification given the concerns discussed in this paragraph.

Figure 2.5: Business Stealing + Market Expansion Externalities

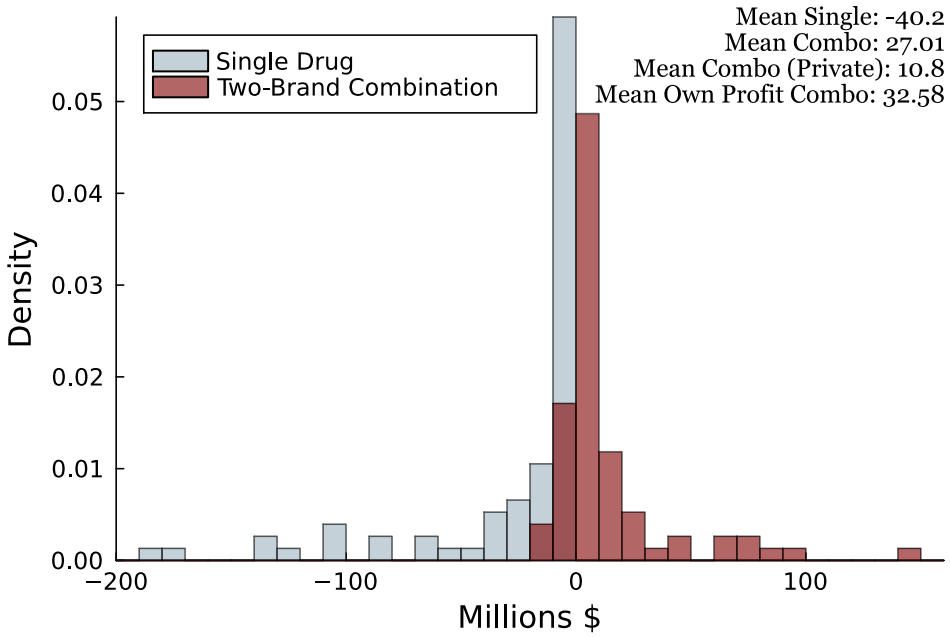

Notes: Figure shows the distribution of the sum of business stealing and market expansion externalities over introduction events. Combination entry events with at least two firms' branded drugs are in red and single-agent entry events are in gray.

separately for three types of regimens: single drugs, combinations with at least two different firms' branded drugs, and all other combinations.

The introduction of single drugs, for which there is no market expansion effect, results in large business stealing effects on other drugs (not owned by the trial sponsor) on average: negative $40 million per year per new drug. For combinations with at least two firms' branded drugs, market expansion dominates business stealing on average, with the sum of these effects averaging $27 million per year per combination. These externalities are significant compared to the own-profit effect of introducing the combination, which averages $33 million per year. Extrapolated over the average patent length of affected drugs implies positive externalities upwards of $200 million over the life-cycle of each new combination therapy. Figure 2.5 shows the distribution of the sum of business stealing and market expansion by regimen entry event, where combinations with at least two firms' branded drugs are shown in red and single-drugs are shown in gray. This sum is positive for approximately 80% of combination entry events, indicating that the market expansion effect dominates the business stealing effect.

All types of new regimens also have positive spillovers on consumers and negative spillovers on insurers, and the price adjustment terms in the profit change decomposition are relatively small in magnitude. In sum, the net externality is positive on average for each type of new regimen,

Figure 2.6: Public Innovation Profit Spillovers

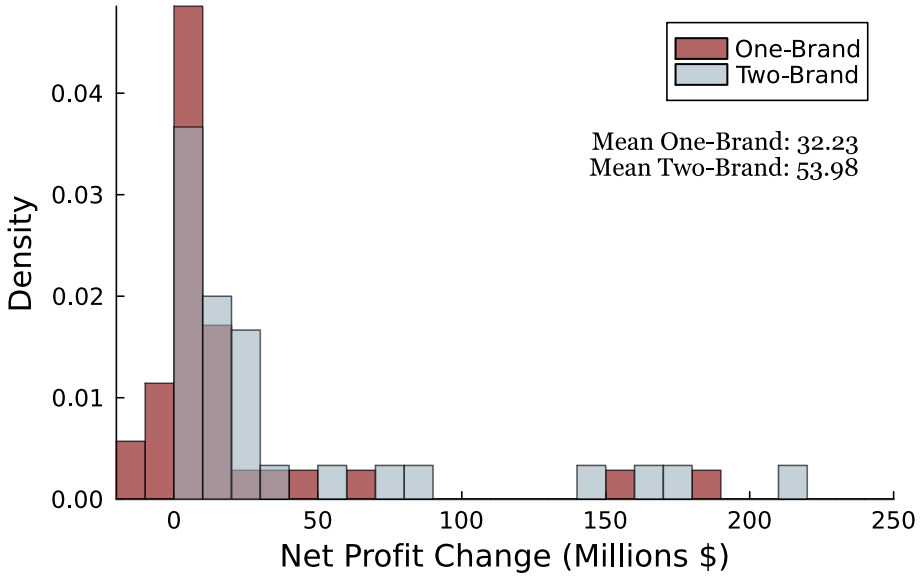

Notes: Figure shows the distribution of the net profit externality (i.e., spillovers on firms) for combinations trialed by public innovators. Combinations with one-branded drug are in red and combinations with two-branded drugs are in gray.

but largest for combinations with at least two firms' branded drugs, at $35 million per year per new combination.

We can also compute the externalities from the perspective of the innovator who ran the trial in the data (i.e., using method 2). Publicly-funded combination trials have large positive profit spillovers on firms on average, shown in Figure 2.6. Privately funded trials for combinations with at least two branded drugs owned by different firms have a mean sum of business stealing and market expansion externalities of $11 million per year per combination, as annotated on Figure 2.5.

Together, these results suggest that firms are often under-incentivized to conduct trials for combination therapies because of large positive externalities on other firms and patients. Having quantified the externalities associated with combination innovation, we now derive the implications of these externalities for innovation decisions by building a dynamic model of combination innovation decisions. We use this model to explore potential policies to support combination innovation.

2.6 Dynamic Model of Combination Innovation

The results in Section 2.5 indicate that significant externalities may arise from combination innovation. Analyzing policies to correct these externalities requires imposing additional structure on innovation decisions. In this section, we develop and estimate a dynamic model of combination innovation. We use this model to explore how externalities influence the path of innovation: market expansion externalities create incentives for firms to *free ride* off others' combination innovations, and innovation by public innovators may similarly *crowd out* private combination innovation. We first use the model to recover primitives of the innovation process, including the incentives for free-riding and public crowd-out. Then, in Section 2.7, we use the model to explore potential policies to support combination innovation.

2.6.1 Setup

We model the innovation decision for each combination regimen r as a dynamic discrete-choice game. This game involves a set of innovators I: firms with at least one patented drug in the regimen, and a public innovator.[78] Innovators can choose to run a clinical trial for regimen r to learn its quality. We refer to this regimen r as the *focal regimen* of its game. Time is discrete (yearly) with a finite horizon, which allows us to capture the fundamental nonstationarity of the setting arising from individual drug introduction and patent expiry.

At year t, the state variables relevant for the decision to trial focal regimen r are summarized by the vector $s_{rt} \in \mathcal{S}_r$. This state includes information about whether other potential regimens have been trialed and, if so, their revealed qualities. This state space suffers from the curse of dimensionality, and we discuss simplifications to facilitate estimation below.

In each year t until focal regimen r has been trialed, each innovator $i \in I$ takes action a_{rit}: trial the focal regimen, $a_{rit} = 1$, or not, $a_{rit} = 0$. Each action has some i.i.d. (across firms and time) private cost shock $\varepsilon_{rit,a}$, drawn from a Type 1 extreme value distribution, scaled by parameter θ^{ε}. Trialing focal regimen r by innovator i at time t has cost κ_{rit}. We assume this decision represents the cumulative efforts to bring the combination to market, and we abstract away from any choice of effort (e.g., trial size) conditional on trialing.[79]

A trial is successful with probability χ, where success means that the regimen is of sufficient quality to be taken by patients (i.e., appears in treatment guidelines such as Chu and DeVita,

[78] Therefore, I is the union of the set of firms $\mathcal{F}$ that have at least one patented drug in the regimen and a public innovator.

[79] We do not model different phases of trials. Combinations are involved in 1.3 different trials (phases) on average. We take the trialing decision for a particular regimen to be the first trial for that regimen in the data.

2019). If trialed successfully, regimen r has quality $\xi_r \sim G_c(\cdot)$, where c is the cancer the regimen is being considered for. The distribution G_c is estimated from our demand model in Section 2.4.

Each innovator maximizes the discounted sum of flow profits net of trial costs. For a firm, flow profits in state s_{rt} are defined to be its profits from regimens that have been successfully trialed. We model the flow profits of the public innovator as the sum of consumer surplus and weighted firm profits:

$$\pi_{rit}(s_{rt}) = CS(s_{rt}) + \lambda \sum_{i' \neq \text{ public}} \pi_{ri't}(s_{rt}),$$

where the weight on firm profits $\lambda \in [0, 1]$ will be estimated.

Equilibrium

For each focal regimen game, we solve for a type-symmetric pure strategy Markov Perfect Equilibrium (MPE), as is common in the literature following Maskin and Tirole (1988) and Ericson and Pakes (1995). Each innovator i's Markov strategy σ_i at time t is a mapping from the current state vector $s_{rt} \in \mathcal{S}_r$ and vector of private shocks ε into a trialing decision: $\sigma_i : \mathcal{S}_r \times \mathbb{R}^2 \to a_i$. The profile of Markov strategies is given by the vector $\sigma = (\sigma_i)_{i \in I}$. As we discretize the state space (discussed below), an equilibrium of the game exists (Doraszelski and Satterthwaite, 2010). However, we cannot guarantee uniqueness since multiple innovators could trial regimens in a particular period.

Note that equilibrium strategies are defined for the focal regimen of each game, and we take the trialing decisions of all other regimens as exogenous. This assumption imposes separability across the different focal regimen games: the state of the focal regimen evolves as a function of equilibrium strategies, while the states of other regimens evolve exogenously. Separability is critical to making the problem computationally tractable, yet preserves the key economic forces our dynamic combination innovation model seeks to capture: externalities on drug owners and patients, the incentive to free ride off others' combination innovation, and public innovation crowd-out. We discuss this assumption and its implications for estimation in further detail in Section 2.6.2.

Timing

For focal regimen r, the game starts when each component drug of the regimen has been trialed in at least one other cancer clinical trial, and ends 5 years past the last patent expiry event.[80] We assume individual drug arrival and patent expiry is exogenous and deterministic. Let t_{r1} denote the first year of the game for focal regimen r and denote the last year t_{rT}. The timing of each year t of the game for focal regimen r is as follows:

1. The state $s_{rt} \in \mathcal{S}_r$ is observed.
2. Equilibrium drug prices are determined via the static price setting game (as estimated in Section 2.4), and each innovator i realizes variable surplus $\pi_{rit}(s_{rt})$.
3. Each innovator i observes private cost shocks $\varepsilon_{rit,a}$ and innovators simultaneously decide whether to trial the regimen or not following strategy profile σ. If an innovator i trials the focal regimen, she pays cost κ_{rit}.
4. The state evolves to $s_{rt+1} \in \mathcal{S}_r$ on the basis of the outcome of the focal regimen trial (if applicable) and the exogenous trialing decisions and outcomes of other regimens.

Dynamic Game

The ex-ante choice-specific value function of not trialing focal regimen r by innovator i at time $t < t_{rT}$ at state s_{rt} is given by the following Bellman equation:

$$v_{rit}(s_{rt}, 0) = \pi_{rit}(s_{rt}) + \beta \mathbb{E}_{s_{rt+1}|\sigma}[V_{rit+1}(s_{rt+1})|s_{rt}, 0]. \tag{2.7}$$

The first term is the flow surplus for innovator i. When i is a firm, it only considers its own profits, ignoring the potential externalities on firms (business stealing and market expansion) and patients. The second term is the discounted expected future value V_{rit+1}. This expectation is calculated over potential states in the next period, capturing two key elements: (i) whether another innovator i' trials the focal regimen (following strategy $\sigma_{i'}$) and that trial's outcome (if applicable) and (ii) the exogenous trialing decisions and outcomes of other regimens. The possibility that another innovator may trial the focal regimen creates the incentive to free ride off that innovator's trial. Similarly, trials run by the public innovator may crowd out trials run by a firm.

[80] The regimens trialed in the data are typically trialed within this time horizon. As our demand data end in 2019, we extrapolate the profit functions we estimate in 2019 forward to all remaining years of the game. Alternative ways of extrapolating could include estimating a constant profit decay or growth factor in years before 2019 and applying in future years.

The ex-ante choice-specific value function of trialing focal regimen r by innovator i at time $t < t_{rT}$ at state s_{rt} is

$$v_{rit}(s_{rt}, 1) = \pi_{rit}(s_{rt}) - \kappa_{rit} \times \mathbf{1}_{r \text{ not trialed by } t} + \beta \mathbb{E}_{s_{rt+1}|\sigma}[V_{rit+1}(s_{rt+1})|s_{rt}, 1]. \tag{2.8}$$

The first term is flow surplus minus trial costs, where the costs are paid only if the focal regimen has not been trialed. We model trialing as an absorbing state so that, after trialing, innovators remain in a state in which the focal regimen has been trialed. The second term is the discounted expected future value, conditional on the focal regimen having been trialed. The expectation is again calculated over potential states in the next period, capturing: (i) the outcome of focal regimen trial (if applicable) and (ii) the exogenous trialing decisions and outcomes of other regimens.

At the final year of the game $t = t_{rT}$, innovators do not trial and simply receive flow surplus from the current state.

Conditional Choice Probabilities

Let $Pr^{\theta}(a_{rit}|s_{rit})$ be the conditional choice probability of innovator i taking action $a_{rit} \in \{0, 1\}$ in the focal regimen r game at time t, conditional on parameters θ and state s_{rt}. The conditional choice probability of trialing is then

$$Pr^{\theta}(1|s_{rt}) = \frac{\exp\left(\frac{v_{rit}(s_{rt},1)}{\theta^{\varepsilon}}\right)}{\exp\left(\frac{v_{rit}(s_{rt},0)}{\theta^{\varepsilon}}\right) + \exp\left(\frac{v_{rit}(s_{rt},1)}{\theta^{\varepsilon}}\right)}. \tag{2.9}$$

State Space

The state space summarizes variables relevant to the profitability of trialing the focal regimen. Without further restrictions, it is infinite dimensional: predicting profits requires knowing the states of all other regimens (i.e., indicators of whether regimens have been trialed and resulting qualities). This state space suffers from the curse of dimensionality even if we discretize the quality distribution, so we make an additional simplification to reduce the size of the state space.

We assume each innovator tracks the trialing status (or outcomes) of a subset of potential combination regimens and a "fringe" regimen. The fringe regimen represents trialing outcomes of all other regimens and has a relatively low probability of success. This assumption is similar to those made in partially oblivious equilibrium (Weintraub et al., 2008, Benkard et al., 2015)

or moment-based Markov equilibrium (Ifrach and Weintraub, 2016). We denote by $\mathcal{R}_r$ the set of regimens included in the state for focal regimen r, which we refer to as *tracked* regimens. For focal regimen r being trialed for cancer c, we let the set of tracked regimens include the focal regimen r, the regimens that were ultimately successfully trialed for that cancer c, and a fringe regimen with a relatively low probability that summarizes all remaining regimens. The expectations in (2.7) and (2.8) integrate over the exogenous evolution of these tracked regimens.

We discretize the regimen quality distribution G_c to have N_ξ points. The state is summarized by the quality levels of the tracked regimens, where the quality level is either not trialed, trialed and failed, or trialed and successful with some quality level. The state space $\mathcal{S}_r$ for the focal regimen r game is therefore given by

$$\mathcal{S}_r = \Big\{\text{not trialed},\ \text{trialed and failed},\ \text{trialed and quality } \xi_1, \ldots,\ \text{trialed and quality } \xi_{N_\xi}\Big\}^{|\mathcal{R}_r|}. \tag{2.10}$$

2.6.2 Estimation and Identification

The parameters in the model are the innovation fixed cost κ_{rit}, scale parameter θ^ε, profit weight λ, discount rate β, and success rate χ. We estimate a common fixed cost κ, scale parameter θ^ε, and profit weight λ via a full-solution method using maximum likelihood. We set a yearly discount rate of .9, and we set the success rate to be the rate observed in the data for each cancer.

The likelihood of a candidate parameter vector $\theta = (\kappa, \theta^\epsilon, \lambda)$ is computed as follows. For each focal regimen $r \in \mathcal{R}$, we approximate the solution to the game at parameters θ (details of this approximation procedure are given in Section 2.6.3), obtaining ex-ante choice specific value functions (2.7) and (2.8). These functions are used to compute the conditional choice probabilities $Pr^\theta(a_{rit}|s_{rit})$ as defined in (2.9). As is common in the literature on dynamic games, we fix competitor conditional choice probabilities at first-stage estimated values (via multinomial logit) in estimation.[81] This assumption is made to avoid issues with multiple equilibria when iterating over conditional choice probability profiles (Aguirregabiria and Mira, 2010; Sweeting, 2013; Bodéré, 2023). Let $\{(\hat{a}_{rit}, \hat{s}_{rit})\}_{r\in\mathcal{R}, i\in\mathcal{I}, t\in\mathcal{T}}$ denote the set of observations (over regimens, innovators, and time) of action-state tuples in data. Define the log-likelihood

[81] The specification of first-stage CCPs and results are shown in Appendix B.6.

function as

$$\sum_{r\in\mathcal{R},i\in\mathcal{I},t\in\mathcal{T}} \frac{1}{|\mathcal{R}|\,|\mathcal{I}|\,|\mathcal{T}|} \log\left(Pr^{\theta}(\hat{a}_{rit}|\hat{s}_{rit})\right). \tag{2.11}$$

We find the θ that maximizes (2.11).

The estimates of profits and consumer surplus from the introduction of a new regimen (from Sections 2.4 and 2.5), combined with observed trialing decisions, are key inputs into identifying the dynamic parameters. For example, conditional on expected profits, a high fixed cost of innovation will lower the CCP of trialing. Similarly, a high weight on aggregate profits in the public innovator's objective will increase the public innovator's CCP of trialing certain types of regimens (e.g., those including patented drugs) relative to others (e.g., those with all generic drugs). Finally, the scale parameter is identified since revenues are taken as given when estimating the dynamic model.

We make additional assumptions to ease the computational burden of estimation. First, we estimate the model for focal regimens for colorectal cancer. Colorectal cancer is not unique compared to the other cancers in our sample, but focusing a single cancer alleviates significant computational challenges.[82] Second, we focus on regimens that were trialed in the data rather than considering the full set of potential combinations given a set of available drugs. By focusing on regimens that were ultimately trialed in the data, the model is most useful for understanding the timing of trialing and which innovator runs the trial rather than the extensive margin of which regimens are trialed. Extending the model to consider the extensive margin requires estimating the (unobserved) quality distribution of this larger risk set of potential regimens. Such an extension is possible with more data about the potential quality of regimens as a function of characteristics such as mechanism of action, known interactions, etc., but is beyond the scope of our analysis.[83]

Within each focal regimen, taking the trialing decisions of other regimens as exogenous is critical to the computational tractability of the game. Given the separability across focal regimens, we can parallelize across focal regimens when approximating the game solution. This assumption is also motivated by institutional details: decisions in large pharmaceutical companies are likely decentralized, and a model with strategic interactions within each focal regimen, taking the evolution of other regimens as given, approximates this structure. This is especially true of public innovation decisions, which are made across many hundreds of different publicly-funded

[82] An extension to all cancers in our data will be included in future drafts.

[83] In future work, we can make progress on this extension by focusing on two-drug combinations and using the results of trials in the NCI ALMANAC to create a risk set of potential combinations based on having promising results in the laboratory tests.

research organizations.

We compare alternative ways of estimating the model in Appendix B.6. The model has the finite dependence property (Arcidiacono and Miller, 2019), and can be estimated through a two-step estimation approach (similar to Scott, 2014). The first step requires non-parametric estimates of the conditional choice probabilities, while the second estimates dynamic parameters using Euler perturbations. While this approach is computationally less demanding than the full-solution method, it is very demanding of data. We explore it in Appendix B.6, but prefer our current approach given the noise with which the first-stage conditional choice probabilities are estimated in the two-step approach.

2.6.3 Solution Procedure

For each candidate parameter vector in estimation (and counterfactuals), we apply sieve value function approximation (Arcidiacono et al., 2013), extended to a game rather than single-agent problem, to solve for approximate value and policy functions. This method approximates the integrated (i.e., expected) value function with a non-parametric sieve function of state variables. Applying approximation methods is required to make the problem computationally tractable: despite the assumptions made regarding the number of regimens to track in the state space and the discretized quality distribution, the dimension of the state space is prohibitively large to solve the model exactly via backwards induction.[84]

For notational convenience, define the flow surplus net of trial costs as

$$\Pi_{rit}(s_{rt}, a) = \pi_{rit}(s_{rt}) - a \times \kappa_{rit} \times \mathbf{1}_{r \text{ not trialed by } t}.$$

Let n be the dimension of the sieve. The approximation $\hat{v}_{rit,n}$ to the ex-ante value function for the focal regimen r game for innovator i at time t satisfies

$$\begin{aligned}\hat{v}_{rit,n}(s_{rt}) &= \theta^{\varepsilon} \ln\left(\sum_{a_{rit}\in\{0,1\}} \exp\left(\frac{1}{\theta^{\varepsilon}}\left(\Pi_{rit}(s_{rt}, a_{rit}) + \beta \mathbb{E}\left\{\hat{v}_{rit+1,n}(s_{rt+1}) | s_{rt}, a_{rit}\right\}\right)\right)\right) + \tilde{\gamma}, \\ &\approx \mathbb{E}\left\{\max_{a_{rit}\in\{0,1\}} \{v_{rit}(s_{rt}, a_{rit}) + \theta^{\varepsilon}\varepsilon_{rit}(a_{rit})\}\right\},\end{aligned}$$

[84] For example, suppose there are 8 points in the quality distribution in addition to the states of not trialed and trialed and failed. With 10 tracked regimens (including) the focal regimen, there are 10 billion states.

where the sieve $\hat{v}_{rit,n}(s_{rt})$ is

$$\hat{v}_{rit,n}(s_{rt}) = \omega_{rit,1} w_1(s_{rt}) + \ldots + \omega_{rit,n} w_n(s_{rt}) = \omega_{rit} \cdot W_n(s_{rt}),$$

and $\tilde{\gamma}$ is the Euler-Mascheroni Constant.

The vector of coefficients ω_{rit} characterizes the sieve, along with the set of functions W. These functions are all linear terms of the state space, all quadratic interactions, all cubic interactions, etc., until the dimension of the sieve n is reached.[85] We can recover the coefficients of the sieve function via backwards recursion at each t by solving the following program:

$$\hat{\omega}_{rit} = \underset{\omega_{rit}}{\operatorname{argmin}} \sum_{s_{rt} \in \hat{\mathcal{S}}_r} [\omega_{rit} \cdot W_n(s_{rt}) - \theta^{\varepsilon} \ln\left(\sum_{a_{rit} \in \{0,1\}} \exp\left(\frac{1}{\theta^{\varepsilon}} \left(\Pi_{rit}(s_{rt}, a_{rit}) + \beta \hat{\omega}_{rit+1} \mathbb{E}\left\{ W_n(s_{rt+1}) | s_{rt}, a_{rit} \right\} \right) \right) \right) - \tilde{\gamma}]^2.$$

The sieve function can be used to derive the ex-ante choice value functions (2.7) and (2.8), which in turn give the CCPs in (2.9). When computing counterfactuals, we iterate over CCP profiles until they are consistent.

The key assumption for computational tractability is that we use a subset of states $\hat{\mathcal{S}}_r \subseteq \mathcal{S}_r$ for this approximation. This choice of approximation states $\hat{\mathcal{S}}_r$ is crucial to the accuracy of the approximation. In order to increase accuracy of the approximation around the states in the data and states likely reached in counterfactual simulations, we generate the approximation states as follows, similar to Sweeting (2013). For each state observed in the data, we choose a random subset of tracked regimens (including the focal regimen) and change those regimens' quality levels to random values (using weights based on the likelihood of those quality levels being realized). We repeat until the desired number of approximation states is reached. In our baseline specification, we set $\left|\hat{\mathcal{S}}_r\right| = 15,000$ and $n = 1,000$.

We present results about approximation error from Monte Carlo simulations of the model with a small number of tracked regimens in the state space and quality levels in Appendix B.6. We can solve this model exactly via backwards induction and compute approximation error for various choices of the number of approximation states and sieve specification. The approximation error is generally small, even for sets of approximation states with relatively small cardinality compared to the dimension of the full state space.

[85] Additional details about the exact construction of the sieve function are given in Appendix B.6.

Table 2.9: Dynamic Model Parameter Estimates

Parameter	Method	Estimate
Cost κ (million \$)	Estimated – Maximum Likelihood	28.0
		(13.19)
Scale θ^{ε}	Estimated – Maximum Likelihood	10.91
		(0.28)
Profit Weight λ	Estimated – Maximum Likelihood	0.25
		(0.23)
Discount Rate β	External Estimate	0.9
Success Rate χ	Estimated – Mean Success	0.05

Notes: Table presents estimates of the dynamic parameters. Each row is a parameter. Column 2 shows the method used to estimate the parameter. Column 3 contains parameter estimates. Standard errors are in parentheses, calculated from bootstrap samples of focal regimens. Note that these standard errors currently only capture variation in estimates of the dynamic game and not the demand or bargaining models.

2.6.4 Parameter Estimates

Table 2.9 presents estimates of the dynamic parameters. We assume a yearly discount rate of .9, and we observe a 5% success rate for colorectal combination therapies in the data. We estimate trial costs of \$28 million, a scale parameter θ^{ε} of 10.91, and a weight on aggregate profits in the public innovator's objective of .25.

Our cost estimates are similar to expected accounting costs of oncology trials. Sertkaya et al. (2016) estimates that oncology trial costs average \$4.5 million in Phase I, \$11.2 million in Phase II, and \$22.1 million in Phase III, totaling \$37.8 million.[86] While similar, our estimate of \$28 million may differ from the sum of these phases for two key reasons. First, as discussed in Section 2.2.1, successful combinations often do not go through all three phases, often combining Phase I/II or Phase II/III if the component drugs have already been shown to be safe. Our model does not distinguish between stages and instead assumes the innovation costs are the total cost of bringing combinations to market across all phases. Many of the combinations in our data fail after one or two phases, reducing the expected trialing costs. Second, the costs recovered by the model are economic costs rather than accounting costs, so additionally capture the possible opportunity costs associated with not trialing other regimens.

Model Fit: We compare model predicted trial times to the data in Appendix B.6. The model fit is reasonable with high correlation between model predicted and actual trial times, though our

[86] Additional external estimates of the accounting costs of a single stage of an oncology trial are approximately \$100,000 per patient (Sertkaya et al., 2014). In our data, combination trials enroll an average of 116 patients (across all phases), implying an average per-phase cost of \$12 million. This estimate is similar to the Phase 2 trial costs cited above.

model predicts trials occurring approximately 1.5 years earlier than the data on average.

2.7 Designing Policy to Support Combination Innovation

In this section, we use the dynamic model to design cost-effective policies to support combination innovation. The policies we consider include research subsidies, varying the amount of public innovation, and varying the direction of public innovation.

2.7.1 Counterfactual Policies

The large set of externalities that arise from combination innovation suggests that private innovation of new combinations is unlikely to yield a socially efficient outcome, perhaps justifying policy intervention. In principle, a social planner could ensure efficient innovation decisions by giving each firm f a subsidy equal to the net externality after it trials a new regimen r^+. But directly implementing this subsidy is impractical, because it requires correct forecasts of the changes in demand, prices, and profits after the introduction of *each* possible regimen by *each* firm. Moreover, such subsidies would need to account for the dynamic effects on patients and firms, which also depend on a correct forecast of the future path of innovation. Given these difficulties, we study a more limited collection of policies that can improve welfare and are simple to implement: a regimen- and firm-independent research subsidy for combination therapy trials and changes in the amount and direction of public innovation.

We note that crafting these policies to balance public and private innovation is particularly important given the large role of public innovation in pharmaceutical markets. Our analysis suggests an important role for public researchers to study combination therapies neglected by firms, particularly those involving only generic drugs. But policies that expand the role of public innovation can also crowd out private innovation, which may be inefficient to the extent that public innovation is more costly (i.e., through the marginal cost of public funds).

We compute counterfactuals by solving the model at the estimated parameters and counterfactual policies and forward simulating the model to predict trial times (and which innovator trials the regimen), described in more detail in Appendix B.6. To implement a research subsidy we reduce the trial cost κ by some fixed amount. To vary the amount of public innovation, we set the CCPs of public innovation at either lower or higher (exogenous) levels than we estimate. To vary the direction of public innovation, we set the CCPs of public innovation of certain regimens at higher or lower levels than we estimate.

2.7.2 Results

Table 2.10 summarizes the effects of counterfactual innovation policies for funding colorectal cancer combination trials.[87] For each counterfactual policy, we compute the total cost to the government of implementing the policy, the change in consumer surplus, the change in firm profits, and the net change in welfare. Figure 2.7 compares the policy cost to the government with the gains in consumer surplus and firm profits for a subset of the policies we consider.

Research Subsidy: The first set of policies we consider is a constant research subsidy to firms for running the trial at a percentage of the estimated trial cost. The government cost of implementing this subsidy consists of two components. First, offering a subsidy may increase the probability of private trialing and result in fewer publicly funded trials relative to baseline. Second, the government must pay a research subsidy on all privately run trials irrespective of success. A 20% subsidy increases total welfare by approximately $442 million ($642 per patient-year), with larger increases in firm profits compared to consumer surplus because of the large transfer that occurs through the subsidy.

Amount of Public Innovation: The next policy we consider varies the probability of public trialing by setting the public innovator's CCPs as exogenous and at higher levels than the estimated levels. For example, the NIH would reduce the threshold for approving grants by some fixed amount. Increasing the probability of the public innovator has two key effects: First, the public innovator is more likely to trial, and therefore more likely to trial earlier. Second, other innovators may be less likely to trial because of the incentive to free-ride and public innovation crowd-out. Increasing public CCPs by 4 percentage points (uniformly across all states) increases total welfare by approximately $424 million ($616 per patient-year). We use 4 percentage points as it gives a similar policy implementation cost to the research subsidy, and we show results for other policies in Table 2.10. On average, a firm's CCP of trialing decreases by approximately 18% relative to baseline, implying positive but limited public crowd-out.

Direction of Public Innovation: The final set of policies we consider varies the direction of public innovation. As public innovation may crowd out private innovation, it may be possible to redirect public innovation towards combinations that firms are particularly underincentivzed to trial and increase trialing speed while keeping the public budget fixed. We implement this type of directional policy similar to the previous counterfactual, where we set the public innovator's CCPs as exogenous and at either or lower levels than what we estimate, depending on characteristics of the combination.

First, we construct a policy that redirects public innovation towards regimens with at least one

[87] In future drafts, we will include extensions to all cancers in our sample.

generic drug. This type of regimen may feature higher consumer surplus spillovers because of the lower drug prices, and firms may be particularly underincentivzed to trial a regimen when a sufficiently high fraction of component drugs are generic. We reduce public CCPs by a fixed amount for all regimens that do not involve any generic drugs, while we increase public CCPs for regimens that do. It is important to note that in this counterfactual we do not compute the optimal policy for redirecting public innovation based on this regimen characteristic, but rather choose the magnitude of the change in CCPs to deliver similar welfare gains as the previous policies. To obtain a budget neutral policy, we reduce public CCPs by approximately 1 percentage point for regimens that do not involve generic drugs and increase by approximately 5 percentage points for those that do. Despite being budget neutral, this policy increases total welfare by $367 million ($533 per patient-year).

The second policy we consider similarly redirects public innovation towards regimens with at least two on-patent drugs owned by different firms. These regimens are likely to have positive market expansion externalities, resulting in reduced firm incentives to trial them relative to what is socially optimal. We decrease public CCPs by approximately 4 percentage points for regimens that have at most one firm's patented drug and increase by approximately 3 percentage points for regimens with at least two firms' patented drugs. This policy is again approximately budget neutral yet increases total welfare by $337 million ($490 per patient-year).

These policies demonstrate that it is possible to obtain similar welfare gains to untargeted subsidies or increases in public innovation at much lower costs by taking advantage the features of the setting that reduce firm incentives to trial certain regimens: missing property rights and market expansion externalities.

2.8 Conclusion

Market expansion externalities, missing property rights for combinations, and the incentive to free ride off others' combination innovation are forces that tend to reduce the private value of combination innovation below its social value.

This paper presents descriptive evidence consistent with underinvestment in combinations because of these forces, develops an empirical framework to quantify combination innovation externalities, and evaluates alternative innovation funding policies in the context of cancer drug combination therapies. Market expansion externalities often dominate business stealing, showing that there is often underinvestment in combination therapies. Redirecting public innovation towards combinations that firms are particularly underincentized to trial—those with the potential for large consumer surplus spillovers, such as combinations of generic drugs,

Figure 2.7: Policy Cost and Welfare Gains

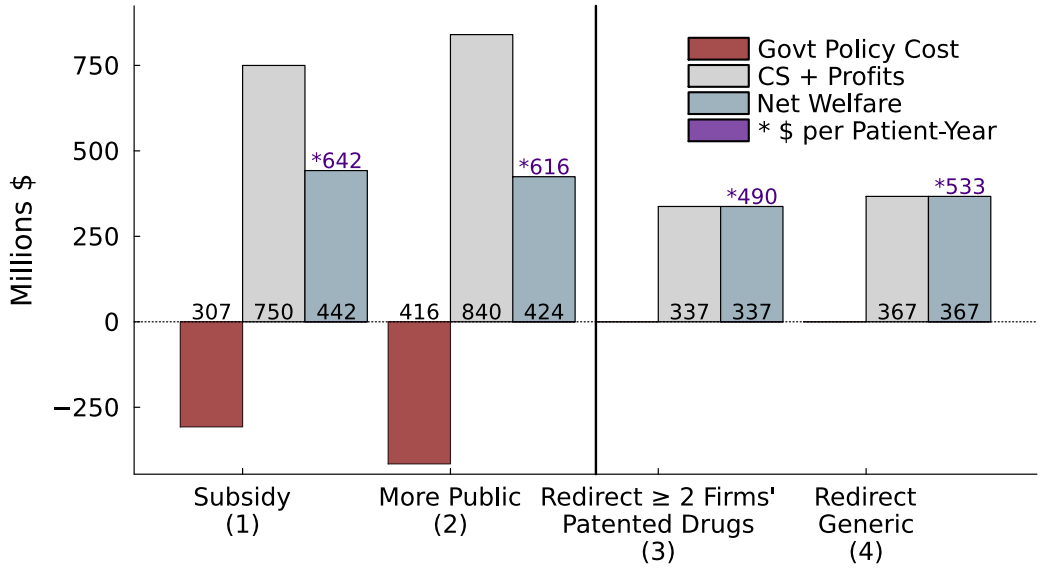

Notes: Figures shows estimates of the government policy cost and change in consumer surplus and profits for counterfactual innovation policies.

and market expansion externalities—minimizes free-riding and public crowd-out and provides a set of relatively simple innovation funding policies that increase total welfare yet are budget neutral.

This analysis focuses on particular aspects of combination innovation, leaving other important features for future research. Mergers and acquisitions of drugs in the early stages of development may play a key role in the shaping the risk set of potential combinations and the amount and direction of combination innovation. And underinvestment in combination therapies suggests potential inefficiencies in the amount, and direction, of innovation in new drugs themselves. Technology (e.g., artificial intelligence) that helps screen potential combinations for predicted success probabilities and externalities could be an important tool in correcting the amount and direction of innovation.

Moving beyond pharmaceuticals, the efficiency of combination innovation will depend critically on the nature of product market competition, the possibility of joint ventures, and the particular property rights institutions. The empirical framework in this paper can be extended to match alternative institutions and provide evidence of how these features shape the efficiency of combination innovation.

Table 2.10: Policy Counterfactual Summary

Counterfactual	**Policy Cost (Govt)** Millions $	**Δ CS** Millions $	**Δ Profits** Millions $	**Δ Welfare** Millions $
Research Subsidy 20%	307	255	495	443
More Public + 1 pp	116	5	120	9
More Public + 2 pp	155	124	292	261
More Public + 3 pp	337	151	507	321
More Public + 4 pp	416	200	640	424
More Public + 5 pp	571	241	837	507
Redirect Generic	0	160	207	367
Redirect $\geq$ 2 Firms' Patented Drugs	0	120	218	338

Notes: Table shows estimates of the government policy cost and change in consumer surplus and profits for counterfactual innovation policies. Column 1 is the cost to the government of implementing the policy. Column 2 is the change in consumer surplus. Column 3 is the change in profits. Column 4 is the change in total welfare (summing the government policy cost and change in consumer surplus and profits). Values are in millions of dollars, rounded to the nearest million.

Chapter 3

Regulating Transformative Technologies

This chapter was co-authored with Daron Acemoglu.

3.1 Introduction

Recent breakneck advances in (generative) artificial intelligence have simultaneously raised hopes of productivity gains in many sectors and fears that this technology will be used for nefarious purposes, even posing an existential risk comparable to nuclear war.[1] Some experts have called to slow down or pause the development and adoption of AI technologies,[2] partly because a slower rollout might provide time to identify danger areas and craft appropriate regulations. However, there is little economic analysis of these issues, and it is unclear whether slowing the development and adoption of a promising, transformative technology ever makes sense.

In this paper, we develop a framework to provide a first set of insights on these questions. We consider a multi-sector economy that initially uses an old technology but can switch to a new, transformative technology. This technology is *transformative* both because it enables a higher growth rate of output, and because it is general-purpose and can be adopted across all sectors of the economy. It also poses new risks. We model these by assuming that there is a positive probability of a *disaster*, meaning that the technology will turn out to have many harmful uses. If a disaster is realized, some of the sectors that had started using the new technology may not be able to switch away from it, despite the social damages. Whether there will be a disaster is initially unknown, and society can learn about it over time. Critically, we assume that the

[1] https://www.nytimes.com/2023/05/30/technology/ai-threat-warning.html
[2] https://futureoflife.org/open-letter/pause-giant-ai-experiments/

greater are the new technology's capabilities, the more damaging it will be when used for harmful purposes.[3]

In this environment, we study (socially) optimal and equilibrium adoption decisions. We first show that it is optimal to adopt the new technology gradually, because this enables greater learning. If all sectors immediately adopted and the disaster transpired, many of them would not be able to switch back and avoid the social damages. Gradual adoption allows society to gain from the new technology while updating its beliefs about whether it will have socially damaging uses. As more time passes without disaster, the belief that there will be a disaster declines ("no news is good news"). As society becomes more optimistic, it is optimal to adopt the new technology across a larger number of sectors. Under weak conditions this adoption path is slow and convex, accelerating only after society is fairly certain that a disaster will not occur. A simple quantitative example indicates that, for reasonable parameters for the new technology's growth advantage and disaster risk, optimal adoption can be very slow.

Perhaps surprisingly, we demonstrate that adoption should be slower when the new technology has a higher growth rate and damages from a disaster are large. This is for two reasons. First, since damages after a potential disaster increase with the new technology's capabilities, a higher growth rate means that damages also grow more quickly. Second, with a higher growth rate the effective discount rate for future output declines, so that short delays in adoption are not very consequential for discounted utility.

Compared to optimal adoption, equilibrium adoption is inefficiently fast if private firms internalize only part of the social damages from a disaster. Even the order in which sectors adopt the new technology can differ between the equilibrium and the optimum—sectors that have high social damages are not necessarily those that have high *private* damages for adopters.[4] Finally, we discuss how regulatory schemes can help to close the gap between optimal and equilibrium adoption. Pigovian taxes, use taxes, or adoption taxes that are sector-specific can fully implement optimal adoption. When sector-specific policies are not feasible, it is generally not possible to implement optimal technology choices, but regulation can still increase welfare by prohibiting use of the new technology in the sectors with the largest potential for harm until the risk of a disaster is sufficiently low.

[3]These assumptions can be motivated with generative AI applications. For irreversibility, once large language models like ChatGPT are deployed in secondary education, it may be impossible to roll back their use, even after it becomes clear that they harm student learning. For the damages rising with productivity, many experts fear that these technologies either pose existential risks or will be misused, both of which would be more damaging when they have greater capabilities (e.g., Shevlane, et al., 2023).

[4]For example, if AI is used to create pervasive disinformation on social media, this may be disastrous for democracy but profitable for social media platforms.

This paper is a first attempt to study the consequences and regulation of transformative technologies that can be used for good or bad. Our conclusions naturally depend on our modeling assumptions and should be interpreted with caution.

There are three literatures on which we build. The first is a growing literature on economic disasters (e.g., Rietz, 1988; Barro, 2006, 2009; Weitzman, 2009, 2011; Martin and Pindyck, 2015, 2021), which explores how the risk of rare economic disasters affects asset prices and cost-benefit analysis, but does not focus on questions of technology adoption.

The second is a literature on technology adoption (e.g., Katz and Shapiro, 1986; Parente and Prescott, 1994; Foster and Rosenzweig, 1995, 2010; Acemoglu et al., 2006; Acemoglu et al., 2007; Comin and Mestieri, 2014). Early work touching on AI includes Galasso and Luo (2018) and Agrawal et al. (2019), but these papers do not focus on issues of learning about social damages from new technologies.

Third, there is a nascent literature focusing on damages from certain technologies (e.g., Bovenberg and Smulders, 1995; Acemoglu et al., 2012). Most closely related to our paper are a few works that discuss the dilemma between growth and existential risk from new technologies, including AI. C. I. Jones (2023) develops a one-sector growth model in which AI can be used to raise the aggregate growth rate, but with small probability causes human extinction. Whether it is optimal to use AI depends crucially on the coefficient of relative risk aversion and whether consumption utility is bounded. Aschenbrenner (2020) incorporates existential risk into C. I. Jones's (2016) model of growth and mortality, and argues that existential risk rises with consumption unless new mitigation technologies are developed. His model thus exhibits an "existential risk Kuznets curve" in which existential risk optimally increases until sufficient R&D resources are shifted toward mitigation. These two papers share our focus on the costs and benefits of transformative technologies, but they do not address the speed of adoption across sectors and do not feature learning about risks over time.

The rest of the paper is organized as follows. Section 3.2 presents our benchmark model. Sections 3.3 and 3.4 characterize optimal and equilibrium technology choices. Section 3.5 discusses the conditions under which optimal technology choices can be restored through regulatory taxes, and Section 3.6 concludes. Omitted proofs and extensions are in the online Appendix.

3.2 Setup

We consider a continuous-time economy that linearly produces a final good from a continuum of sectors $i \in [0,1]$:

$$Y = \int_0^1 Y_i di.$$

A representative household has risk-neutral preferences defined over this final good and discounts the future at rate $\rho > 0$.

Each sector can use an old technology O or a new, transformative technology N. We write $Q_j(t) > 0$ for the quality of technology $j \in \{O, N\}$ at time t, $x_i(t) = 1$ if sector i switches its production process to technology N, and $x_i(t) = 0$ otherwise. Sectoral output is

$$Y_i = (1 - x_i) Q_O + x_i \alpha_i Q_N,$$

where α_i designates the comparative advantage of the new technology, which may vary if the new technology is better-suited for some sectors than others. Given technology choices $x = (x_i)_{i \in [0,1]}$ and qualities $Q = (Q_O, Q_N)$, final output is

$$Y(x, Q) = \int_0^1 (1 - x_i) Q_O + x_i \alpha_i Q_N di.$$

The new technology is *transformative*, both because it is general-purpose and can be applied across all sectors, and because it enables not just the production of more output, but a higher growth rate:

$$g_N > g_O \geq 0.$$

As a result of its restructuring impact on the economy, it also poses new risks. We model these by assuming that there may be a *disaster* whereby the new technology generates negative effects. If a disaster happens, then there will be *damages* of $\delta_i Q_N > 0$ (in units of the final good) in the sectors that are using the technology. We assume that use of the new technology may be irreversible, so that with probability $\eta_i \in (0,1)$ sector i cannot switch to technology O if it is using technology N when the disaster strikes. The realization of this reversibility event is independent across sectors. We assume that damages are proportional to Q_N because the negative effects correspond to misusing the better capabilities of the new technology.

In what follows, we reorder sectors so that δ_i is increasing and assume that i denotes the quantiles of the δ distribution, so that we can take this distribution to be uniform over some

interval $[\underline{\delta}, \overline{\delta}]$. Overall damages then become

$$D(x,Q) = \left(\int_0^1 \delta_i x_i di \right) Q_N.$$

The economy will experience a disaster with probability $\bar{\mu} \in (0,1)$, and if there is a disaster, its arrival time T is distributed exponentially with rate λ. We let $\mu(t)$ denote the (the planner's or society's) posterior belief at t that there will be a disaster, assuming one has not yet arrived. We impose rational expectations, so that $\mu(0) = \bar{\mu}$ and the posterior belief evolves according to Bayes's rule:

$$\dot{\mu}(t) = -\lambda \mu(t)(1-\mu(t)). \tag{3.1}$$

A few comments are in order. First, we model damages in each sector i by the reduced-form function $\delta_i Q_N$ to capture a broad range of potential harms. In the context of AI, these include the spread of disinformation that harms democracy; mass unemployment; and the disruption of production in many sectors from AI-aided cyberattacks.[5] Second, as suggested above, the assumption that damages are proportional to Q_N is related to the transformative nature of this new technology. For example, damages from disinformation from AI will be higher when it can generate better language. Third, we assume that the arrival rate of the disaster—and hence learning about the negative effects of the new technology—is independent of how many sectors switch to the new technology. This is for simplicity, but is not unreasonable since many of the potential misuses of a new technology can be gradually recognized without widespread adoption.[6] Fourth, it can be verified that our results remain identical if, instead of a single economy-wide disaster, there are sector-specific disasters and beliefs about each sector's disaster follow (3.1).

3.3 Socially Optimal Technology Choice

In this section, we set up, solve, and provide comparative statics for the (social) planner's problem.

[5] Our functional form assumptions also impose that the rate of substitution between gross consumption and damages in utility is constant and equal to one. C. I. Jones (2023) points out that this may not hold in the case of existential risk and explores the implications for optimal use of a life-threatening new technology.

[6] Alternative assumptions are discussed in Section 3.6.

3.3.1 Social Planner's Problem

Given risk neutrality, the planner's objective is

$$V(0) = \mathbb{E}_{\mu(0)}\left[\int_0^\infty \exp(-\rho t)[Y(t) - D(t)]\,dt\right], \tag{3.2}$$

where $Y(t)$ and $D(t)$ denote output and damages at time t and the expectation $\mathbb{E}_{\mu(0)}$ is with respect to the prior belief $\mu(0)$ over the disaster's arrival time T. To ensure that the objective is well-defined, we assume

$$\rho > g_N, \tag{3.3}$$

which rules out the case in which the new technology grows so quickly that discounted utility becomes infinite.

It is more convenient to work with the recursive formulation of (3.2), which has the following state variables: the posterior belief of disaster, μ; the time-varying qualities of the old and new technologies, Q; and, after the disaster, the set of sectors that were already using the new technology and for which this use is irreversible. We track these sectors using the vector $\bar{x} = (\bar{x}_i)_{i\in[0,1]}$, where $\bar{x}_i = 1$ if sector i uses technology N irreversibly and $\bar{x}_i = 0$ otherwise. Let $V(\mu, Q)$ denote pre-disaster social welfare, and let $W(\bar{x}, Q)$ denote post-disaster welfare. Then the Hamilton-Jacobi-Bellman (HJB) equations for the planner are

$$\rho V(\mu, Q) = \max_{x_i\in\{0,1\}} \{Y(x, Q) + \mu\lambda(\mathbb{E}[W(\bar{x}, Q)|\,x] - V(\mu, Q))\} + \dot{V}(\mu, Q), \tag{3.4}$$

$$\rho W(\bar{x}, Q) = \max_{x_i\in\{\bar{x}_i,1\}} \{Y(x, Q) - D(x, Q)\} + \dot{W}(\bar{x}, Q). \tag{3.5}$$

Equation (3.5) imposes that x_i cannot be less than $\bar{x}_i$, because $\bar{x}_i = 1$ implies that sector i's use of the new technology is irreversible. V then depends on the conditional expectation of welfare after a disaster given the current technology choices x, denoted by $\mathbb{E}[W(\bar{x}, Q)|\,x]$.[7] In (3.4) we also use the fact that the arrival rate of the disaster, given the posterior μ, is $\mu\lambda$.

To characterize the planner's technology choices, suppose first that the disaster has occurred. The planner's problem in (3.5) is linear, so the solution is

$$x_i = \begin{cases} 1 & \text{if} \quad \bar{x}_i = 1 \quad \text{or} \quad (\alpha_i - \delta_i)Q_N > Q_O, \\ 0 & \text{else.} \end{cases}$$

This expression assumes, without loss of generality, that the planner sticks with the old tech-

[7] To determine this conditional expectation, we use $\mathbb{P}(\bar{x}_i = 1|\,x_i = 1) = \eta_i$ and $\mathbb{P}(\bar{x}_i = 1|\,x_i = 0) = 0$.

nology if indifferent. It also imposes the constraint that $x_i = 1$ when $\bar{x}_i = 1$. Even when unconstrained, it may be optimal to set $x_i = 1$ if the output produced by technology N exceeds its damages plus the output that can be produced by technology O. We first assume that damages are sufficiently large that, whenever possible, the planner chooses technology O after a disaster:

$$\alpha_i \leq \delta_i. \tag{3.6}$$

This enables us to focus on the most interesting case where damages exceed the benefits of the new technology. We return to the general case in Section 3.3.3.

Integrating the HJB equation (3.5) and taking expectations with respect to $\bar{x}$, we have

$$\mathbb{E}[\,W(\bar{x},Q)|\,x\,] = \int_0^1 \left[(1-x_i\eta_i)\frac{1}{\rho-g_O}Q_O + x_i\eta_i\frac{\alpha_i-\delta_i}{\rho-g_N}Q_N\right]di.$$

Before the disaster, it is optimal from (3.4) to use technology N in sector i iff

$$\alpha_i Q_N - Q_O > \mu\lambda\eta_i\left[\frac{1}{\rho-g_O}Q_O - \frac{\alpha_i-\delta_i}{\rho-g_N}Q_N\right]. \tag{3.7}$$

Intuitively, the left-hand side is the flow gain from using technology N in sector i, while the right-hand side is the expected loss due to the disaster, including both the discounted value of lost output and the irreversible damages. These losses are multiplied by the posterior arrival rate of the disaster $\mu\lambda$ and the probability of irreversibility η_i. Since μ is decreasing and Q_N/Q_O is increasing, for any initial state $(\mu(0),Q(0))$ there exists a time $t_i < \infty$ such that technology O is used in sector before t_i and technology N is used thereafter.

3.3.2 Socially Optimal Technology Adoption

To determine how (socially) optimal use of technology N changes over time, denote the fraction of sectors that use technology N, or total *adoption*, by

$$X(\mu,q) = \int_0^1 x_i(\mu,q)\,di.$$

Here $q = \log(Q_N/Q_O)$ is the *quality gap* between the technologies, and $x_i(\mu,q) = 1$ iff it is optimal to use technology N in sector i in state (μ,q). For simplicity, we assume that α_i and η_i are constant across sectors and equal to α and η (the general case is studied in Appendix C.2). This implies that there exists a *damage threshold* $L(\mu,q)$ such that it is optimal to adopt the new technology in sector i iff $\delta_i < L(\mu,q)$. Letting F denote the cumulative distribution function of

the uniform distribution over $[\underline{\delta}, \bar{\delta}]$, total adoption is then just the fraction of sectors below the damage threshold:

$$X(\mu, q) = F(L(\mu, q)).$$

The following proposition is immediate from (3.7), and we omit its proof:

Proposition 7. *Suppose (3.6) holds and α_i and η_i are constant across sectors. It is socially optimal to use technology N in sector i iff $\delta_i < L(\mu, q)$, where*

$$\frac{L(\mu, q) - \alpha}{\rho - g_N} = \frac{\alpha - \exp(-q)}{\mu\lambda\eta} - \frac{\exp(-q)}{\rho - g_O}. \tag{3.8}$$

$L(\mu, q)$ (and thus $X(\mu, q)$) is increasing in α and q; decreasing in g_O, λ, and μ; and decreasing in g_N, provided that $L(\mu, q) > \alpha$.

Given (3.6), the condition $L(\mu, q) > \alpha$ is satisfied as soon as there is any adoption. Proposition 7 then implies that when the new technology enables *faster growth*, its adoption should be *slower*. This is because of a *precautionary motive*—even though the planner is risk-neutral, she would like to avoid irreversible damages from the new technology. The faster the new technology grows, the greater are the potential net output losses, strengthening this precautionary motive.

The comparative statics in Proposition 7 are partial because they hold the state (μ, q) fixed. Full comparative statics must account for how parameter changes affect the evolution of the state $(\mu(t), q(t))$. The belief $\mu(t)$ does not depend on the growth rates g_O and g_N, but the quality gap $q(t) = q(0) + (g_N - g_O)t$ does. The damage threshold $L(\mu, q)$ is increasing in the quality gap, so any change in growth rates affects adoption at each $t > 0$ through both the direct effects described in Proposition 7 and the indirect effects through changes in the quality gap $q(t)$. The next proposition characterizes these total effects.

Proposition 8. *Suppose (3.6) holds and α_i and η_i are constant across sectors.*

1. *$X(\mu(t), q(t))$ is decreasing in g_O.*
2. *There exists an earliest time $\bar{t} < \infty$ such that $X(\mu(t), q(t))$ is decreasing in g_N if $t > \bar{t}$. The time $\bar{t}$ is decreasing in g_N.*
3. *Adoption falls to zero as g_N approaches ρ, i.e., $\lim_{g_N \uparrow \rho} X(\mu(t), q(t)) = 0$.*

The first part of Proposition 8 establishes that the comparative static for g_O from Proposition 7 generalizes in the presence of the indirect effects through $q(t)$—the quality gap $q(t)$ is declining in g_O, reinforcing the direct effect and decreasing adoption. The second part shows that the new technology's growth rate has more nuanced implications: Adoption is not always decreasing in g_N, but it is after some critical time $\bar{t}$, and this time itself is a decreasing function of g_N. This

holds because the precautionary motive highlighted above must compete with the fact that the quality gap $q(t)$ is increasing in g_N, but this indirect effect can dominate only at short time horizons.

The third part of proposition establishes that as g_N increases towards the discount rate, adoption almost stops. This might appear paradoxical initially, but is also intuitive. When g_N is approximately equal to ρ, the benefits from the new technology are very high, leading to nearly infinite discounted utility provided no disaster arrives. Delay in adoption thus has little effect on these benefits. However, a disaster will have huge negative consequences, and avoiding it now takes precedence.

The next proposition further characterizes the shape of the adoption curve. Since F is uniform, $\dot{X}(\mu,q) = f\dot{L}(\mu,q)$, where f is the constant density of F. Hence, the *curvature* of technology adoption is

$$\frac{\ddot{X}(\mu,q)}{\dot{X}(\mu,q)} = \frac{\ddot{L}(\mu,q)}{\dot{L}(\mu,q)}.$$

We therefore have:

Proposition 9. *Suppose (3.6) holds.*

1. *$\dot{L}(\mu,q) > 0$ is decreasing in g_O, and it is decreasing in g_N iff the quality gap is sufficiently large, i.e.,*
$$\alpha\exp(q) - 1 > \frac{(\rho - g_N) - (g_N - g_O)}{1-\mu}\left(\frac{1}{\lambda} + \frac{\mu\eta}{\rho - g_O}\right).$$

2. *There exists a positive constant $G(\mu,q)$ such that if $\alpha\exp(q) > 1$, $\ddot{L}(\mu,q) > 0$ is positive iff $g_N - g_O > G(\mu,q)$. $G(\mu,q)$ is independent of g_N and increases to infinity over time.*

The intuition for the first part is the same as for Proposition 8: The damage threshold increases as the posterior belief μ falls and the quality gap q grows. Faster growth for technology O slows the rate of increase of the quality gap and raises the opportunity cost of using technology N after the disaster. Consequently, the damage threshold grows less quickly in each state. Faster growth for technology N raises both the rate of increase in the quality gap and the net output losses from technology N after the disaster. The latter effect dominates when the quality gap is sufficiently large because additional improvements in technology N relative to O have only a negligible impact on the planner's technology choice.[8]

The second part of the proposition proves that when the new technology's growth advantage is sufficiently large, its adoption will have a convex segment where adoption accelerates. This

[8]The latter effect also dominates regardless of the quality gap whenever $L(\mu,q) > 0$ and $g_N - g_O \geq \rho - g_N$.

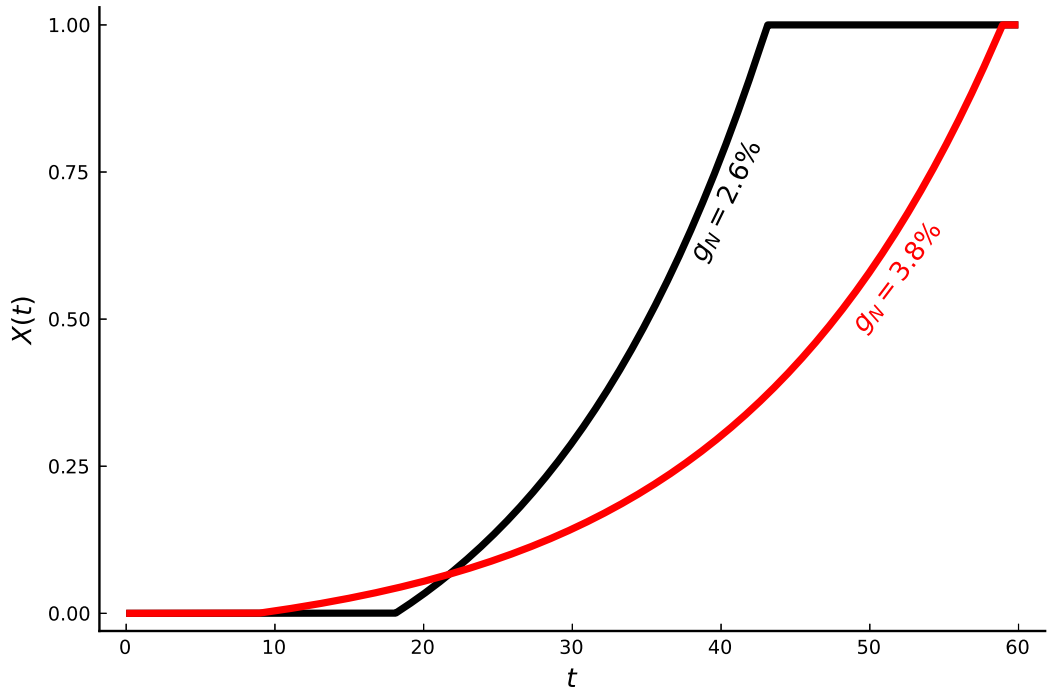

Figure 3.1: Adoption curves $X(t) \equiv X(\mu(t), q(t))$ for different values of g_N. The remaining parameter values are $\rho = 0.04$, $\lambda = 0.05$, $\eta = 0.5$, $\alpha = 1$, $g_O = 0.02$, $\underline{\delta} = 1$, and $\bar{\delta} = 5$. The initial state is $\mu(0) = 0.2$ and $q(0) = 0$.

result holds even though the learning rate $|\dot{\mu}|$ falls at a greater than exponential rate when $\mu < \frac{1}{2}$ (in particular, $\frac{d}{dt}|\dot{\mu}| = -\lambda|\dot{\mu}|(1-2\mu)$). This is because expected damages from technology N in sector i are proportional to the posterior μ, and as μ declines, larger increases in the damage threshold $L(\mu, q)$ are needed to balance the expected damages and benefits in the "marginal" sector.[9]

To illustrate these results, we depict the time path of adoption in a couple of parameterized cases in Figure 3.1. We set $g_O = 2\%$ in line with trend GDP growth in developed economies and $\rho = 0.04$ to produce a risk-free interest rate of 4%. We choose two values for g_N based on Chui, Roberts, Yee, Hazan, et al. (2023), who forecast an increase in the growth rate of 0.6-3.6% in the United States between 2023 and 2040 from AI and other automation technologies. We take the lower end of this range, $g_N - g_O = 0.6\%$, and a higher but still conservative estimate from the middle of the range: $g_N - g_O = 1.8\%$ (while still satisfying (3.3)). We take the two technologies to have the same quality in year $t = 0$, thus $q(0) = 0$. We suppose that damages range from one to five times gross sectoral output ($\underline{\delta} = 1$, $\bar{\delta} = 5$), and we set $\eta = 0.5$ so that half of all sectors using the new technology cannot switch back after a disaster. We set the expected

[9] In Appendix C.2, we verify this intuition by showing that learning dynamics favor *concave* adoption when sectors are heterogeneous according to α_i instead of δ_i.

arrival time of a disaster (if one exists) to be 20 years, which gives $\lambda = 0.05$. Finally, a recent survey of AI experts reports a median estimate of existential risk of about 10%,[10] and since we are interested in non-existential misuses of AI as well, we choose the initial disaster probability to be twice as large, $\mu(0) = 20\%$. Figure 3.1 shows that optimal adoption is slow, taking about 40 years until full adoption when $g_N = 2.6\%$ and almost 60 years when $g_N = 3.8\%$.

3.3.3 Optimal Adoption with Small Damages

We have so far imposed (3.6), ensuring that the post-disaster damages from the new technology are *large* and exceed its gross output within each sector. This is a natural benchmark, since our analysis is motivated by significant potential harms from AI. We now relax this assumption and allow a sector's damages to be *small* relative to its output under the new technology ($\delta_i < \alpha$).

In Appendix C.3, we show that socially optimal adoption is again characterized by a damage threshold $L(\mu, q)$, and we prove the following analogue to Proposition 8 for small damages.

Proposition 10. *Suppose α_i and η_i are constant across sectors. For all t with $L(\mu(t), q(t)) < \alpha$:*

1. *$X(\mu(t), q(t))$ is decreasing in g_O.*
2. *$X(\mu(t), q(t))$ is increasing in g_N.*
3. *If $q(0)$ is sufficiently low and $X(\mu(t), q(t)) < F(\alpha)$, adoption is bounded below $F(\alpha)$ as g_N approaches ρ, i.e., $\lim_{g_N \uparrow \rho} X(\mu(t), q(t)) < F(\alpha)$.*

Adoption among sectors with small damages is still decreasing in g_O, but in contrast to the case with large damages, it is increasing in g_N. Gradual adoption remains optimal even when g_N increases toward the discount rate ρ. With small damages, using technology N is always optimal in the long run. Nevertheless, gradual adoption is optimal to learn about the probability of a disaster (before one occurs) and to delay the adoption of technology N in case of a disaster until the quality gap becomes sufficiently large. This strategy thus avoids *temporary* costs of irreversibility. Further analysis of this case is presented in Appendix C.3.

Finally, we note that if damages are uncertain, any chance of large damages leads to longer optimal delay, even if expected damages are small, in order to avoid the possibility that damages turn out to be large *and* adoption is irreversible.

In summary, the optimal adoption of a new, transformative technology should be gradual,

[10] https://aiimpacts.org/2022-expert-survey-on-progress-in-ai

particularly when its superior capabilities also make its potential damages greater and there is learning about the likelihood of misuse (a "disaster").

3.4 Equilibrium Technology Choice

We now characterize equilibrium technology adoption when private firms do not fully internalize social damages.

3.4.1 The Firm's Problem

Suppose now that in each sector, the choice of technology is made by a private (representative) firm that seeks to maximize expected discounted profits. To simplify, we assume that the firm in sector i appropriates all output of its intermediate as profits, but only internalizes *private damages* $\gamma_i \leq \delta_i$. This textbook externality leads to excessively fast adoption of the new technology before the disaster, and our main results below describe how the equilibrium and socially optimal adoption curves differ.

Firm i's profit maximization problem can be formulated recursively in the same way as the planner's problem in the previous section. The state variables before the disaster are again μ and Q, and after the disaster they are $\bar{x}_i$ and Q. Let $\Pi_i(\mu, Q)$ denote the firm's pre-disaster value, $\Phi_i(\bar{x}_i, Q)$ its post-disaster value, and $Y_i(x_i, Q)$ its (gross) output. The HJB equations for the firm are

$$\rho \Pi_i(\mu, Q) = \max_{x_i \in \{0,1\}} \{Y_i(x_i, Q) + \mu\lambda(\mathbb{E}[\Phi_i(\bar{x}_i, Q) | x_i] - \Pi_i(\mu, Q))\} + \dot{\Pi}_i(\mu, Q), \tag{3.9}$$

$$\rho \Phi_i(\bar{x}_i, Q) = \max_{x_i \in \{\bar{x}_i, 1\}} \{Y_i(x_i, Q) - x_i \gamma_i Q_N\} + \dot{\Phi}_i(\bar{x}_i, Q). \tag{3.10}$$

These value functions differ from the planner's (3.4) and (3.5) because the firm internalizes only a fraction γ_i/δ_i of the flow damages from technology N.

We now impose a stronger version of (3.6): private damages are sufficiently large that firm i will always choose technology O after the disaster if possible:[11]

$$\alpha_i \leq \gamma_i. \tag{3.11}$$

[11] Without this assumption, an additional inefficiency would arise in equilibrium as firms would use the new technology in some (reversible) sectors even after a disaster.

Similar to the planner's solution, it is privately optimal for firm i to use technology N iff

$$\alpha_i Q_N - Q_O > \mu\lambda\eta_i \left[\frac{1}{\rho - g_O} Q_O - \frac{\alpha_i - \gamma_i}{\rho - g_N} Q_N \right].$$

The only difference between this condition and the planner's optimality condition (3.7) is that private damages γ_i appear instead of social damages δ_i on the right-hand side. Firm i internalizes fewer damages from technology N and thus begins using it earlier.

3.4.2 Equilibrium Technology Adoption

We denote total equilibrium adoption by

$$\tilde{X}(\mu, q) = \int_0^1 \tilde{x}_i(\mu, q)\, di,$$

where $\tilde{x}_i(\mu, q) = 1$ iff firm i uses technology N in state (μ, q). Again assuming that α_i and η_i are constant across sectors, it is immediate that firm i will adopt the new technology iff *private* damages are lower than the damage threshold, $\gamma_i < L(\mu, q)$. Equilibrium adoption is then

$$\tilde{X}(\mu, q) = F_\gamma(L(\mu, q)),$$

where F_γ is the cumulative density function of γ_i.

This characterization implies that all comparative statics results from Section 3.3.2 apply to equilibrium adoption. The results in Propositions 7 and 9 concern only the damage threshold $L(\mu, q)$ and hold exactly as stated, while Proposition 8 applies after replacing $X(\mu, q)$ with $\tilde{X}(\mu, q)$:

Proposition 11. *Suppose (3.11) holds and α_i and η_i are constant across sectors.*

1. *$\tilde{X}(\mu(t), q(t))$ is decreasing in g_O.*
2. *There exists an earliest time $\tilde{t} < \infty$ such that $\tilde{X}(\mu(t), q(t))$ is decreasing in g_N if $t > \tilde{t}$. The time $\tilde{t}$ is decreasing in g_N.*
3. *Adoption falls to zero as g_N increases to ρ: $\lim_{g_N \uparrow \rho} \tilde{X}(\mu(t), q(t)) = 0$.*

In the remainder of this section, we characterize how the optimal and equilibrium adoption curves differ. We first observe that similar adoption curves do not imply that the equilibrium is optimal, because the order in which sectors adopt the new technology matters. For example, private and social damages may be *negatively affiliated*, meaning that high social damage sectors

have low private damages. In this case, the order in which the new technology spreads in equilibrium is exactly the opposite of the optimal order.

Even when the equilibrium and optimal orders of adoption coincide, the equilibrium can be inefficient. To see this, suppose that social and private damages are *positively affiliated*, so that there exists a non-negative and (strictly) increasing function κ with $\gamma_i = \kappa(\delta_i) \leq \delta_i$. We can then write equilibrium adoption as

$$\tilde{X}(\mu, q) = F\left(\kappa^{-1}(L(\mu, q))\right).$$

This equation implies that the equilibrium adoption curve $\tilde{X}(\mu(t), q(t))$ is a distorted version of the optimal adoption curve, with an *equilibrium damage threshold* $\tilde{L}(\mu, q) = \kappa^{-1}(L(\mu, q))$. In this case, knowing how the equilibrium and social damage thresholds differ is sufficient to fully characterize equilibrium inefficiencies. The next proposition determines how the level, rate of change, and curvature of the equilibrium damage threshold $\tilde{L}(\mu, q)$ differ from its social counterpart $L(\mu, q)$.

Proposition 12. *Suppose (3.11) holds and α_i and η_i are constant across sectors.*

1. *The equilibrium damage threshold is always greater than the social damage threshold: $\tilde{L}(\mu, q) \geq L(\mu, q)$.*

2. *The equilibrium damage threshold increases more quickly than the social damage threshold when $\kappa'\left(\tilde{L}(\mu, q)\right) < 1$:*
$$\dot{\tilde{L}}(\mu, q) = \frac{\dot{L}(\mu, q)}{\kappa'\left(\tilde{L}(\mu, q)\right)}.$$

3. *The equilibrium damage threshold is more convex than the social damage threshold when κ is locally concave:*
$$\frac{\ddot{\tilde{L}}(\mu, q)}{\dot{\tilde{L}}(\mu, q)} = \frac{\ddot{L}(\mu, q)}{\dot{L}(\mu, q)} - \frac{\kappa''\left(\tilde{L}(\mu, q)\right)}{\kappa'\left(\tilde{L}(\mu, q)\right)}\dot{L}(\mu, q).$$

These results follow from the definition of the equilibrium damage threshold $\tilde{L}(\mu, q)$. We illustrate them in Figure 3.2 by depicting socially optimal and equilibrium adoption curves for the benchmark parameterizations in Figure 3.1 and a concave affiliation function κ. The equilibrium damage threshold is always greater than its social counterpart and increases more quickly (for the marginal sectors where $\kappa'\left(\tilde{L}(\mu, q)\right) < 1$). Consequently, equilibrium adoption is inefficiently rapid and accelerates when there are high social damages.

In summary, equilibrium adoption of transformative technologies is determined by the same

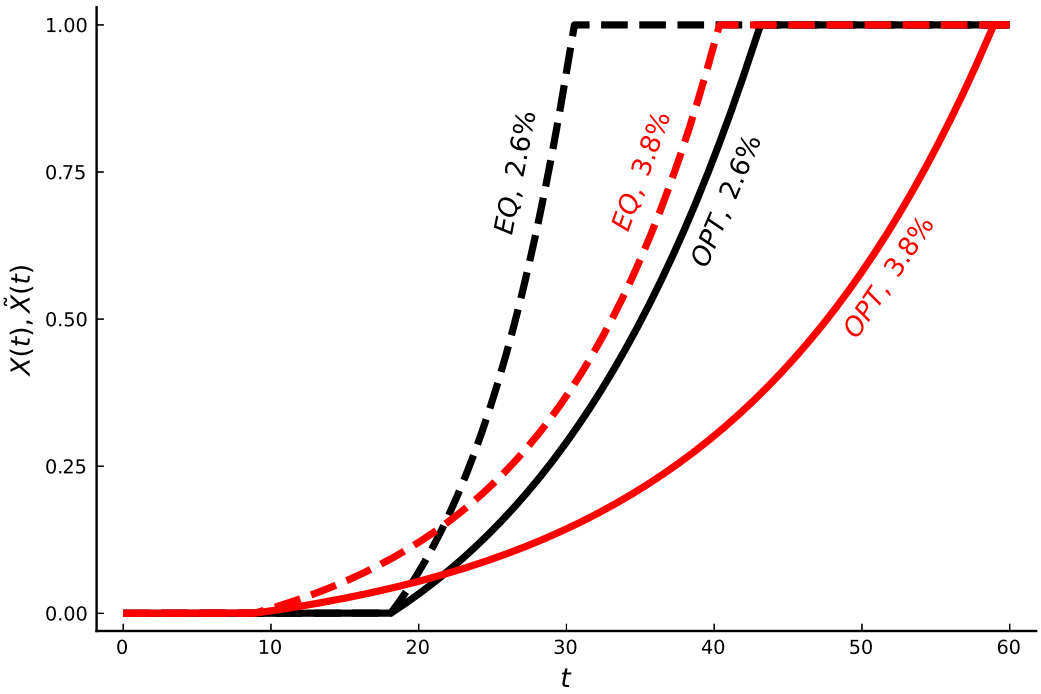

Figure 3.2: Socially optimal and equilibrium adoption curves, $X(t)$ and $\tilde{X}(t) \equiv \tilde{X}(\mu(t), q(t))$. The calibration is the same as in Figure 3.1. The affiliation function is $\kappa(\delta) = \delta^{1/2}$.

forces that shape optimal adoption. However, because firms are motivated by higher productivity and discouraged only by private damages, equilibrium adoption is generally suboptimal: Firms do not fully internalize social damages from potential disasters, so equilibrium adoption is typically too high and rises too quickly, and the order in which sectors adopt the new technology may differ from the optimal one.

3.5 Regulating Technology Choice

Since equilibrium adoption is potentially inefficient, a natural question is whether government regulation can close the gap between equilibrium and optimal adoption decisions. Throughout this section, we continue to assume that (3.11) holds, and we simplify the analysis by focusing on *ex ante* regulations.[12]

Socially optimal and equilibrium technology choices differ because the planner and private firms internalize different damages after the disaster and hence different *expected* damages before the disaster. A straightforward way to correct firms' incentives is through a *use tax* that

[12]We ignore *ex post* ("Pigovian") taxes both because their analysis is essentially identical to our characterization of use taxes, and also because they may not be credible as they do not affect technology choice after the disaster—the private sector already stops using the new technology whenever possible.

raises firms' costs of using the new technology *before the disaster*.[13] When sector-specific taxes are feasible, the tax that implements the optimal technology choice for sector i is equal to the difference between expected discounted social and private damages:

$$\tau_i(\mu, Q_N) = \mu\lambda\eta_i \frac{\delta_i - \gamma_i}{\rho - g_N} Q_N. \tag{3.12}$$

The next proposition notes several properties of these optimal taxes.

Proposition 13. *The optimal use tax $\tau_i(\mu, Q_N)$ is larger in sectors with a larger probability of irreversibility η_i and a larger difference between social and private damages $\delta_i - \gamma_i$. It is log-concave in time and limits to zero as $t \to \infty$ iff $\lambda > g_N$.*

The cross-sector comparative statics follow immediately from (3.12). Differentiating (3.12) with respect to time yields

$$\frac{\dot{\tau}_i(\mu, Q_N)}{\tau_i(\mu, Q_N)} = \frac{\dot{\mu}}{\mu} + \frac{\dot{Q}_N}{Q_N} = -\lambda(1-\mu) + g_N.$$

Since μ declines before the disaster, $\tau_i(\mu(t), Q_N(t))$ is log-concave. The difference between social and private damages from a disaster is increasing in Q_N, pushing taxes higher, while growing optimism about the absence of a disaster pushes taxes lower. The tax is eventually decreasing to zero iff learning about the disaster risk is sufficiently fast, $\lambda > g_N$.

Sector-specific taxes require detailed information about damages and may generally be difficult to implement. Even in the benchmark case in which α_i and η_i are constant across sectors, the next proposition shows that a sector-*independent* tax scheme cannot correct inefficient equilibrium adoption unless social and private damages are positively affiliated.

Proposition 14. *Suppose α_i and η_i are constant across sectors. Given any sector-independent use tax $\tau(\mu, Q)$, firm i begins using technology N earlier than firm j iff $\gamma_i \leq \gamma_j$. Socially optimal technology choices can be implemented for any initial state $(\mu(0), Q(0))$ iff social and private damages are positively affiliated. In this case, the following tax is optimal:*

$$\tau(\mu, Q) = \mu\lambda\eta \frac{L(\mu, q) - \kappa(L(\mu, q))}{\rho - g_N} Q_N. \tag{3.13}$$

This proposition clarifies that a sector-independent tax can differentially delay adoption for sectors with different private damages γ_i, but it cannot alter the order of adoption. When private and social damages are positively affiliated, the socially optimal and equilibrium orders of adoption coincide, so a sector-independent tax can fully correct equilibrium inefficiencies.

[13] Naturally, *adoption taxes* that are paid when new technologies are first introduced are equivalent.

When the optimal and equilibrium orders of adoption differ, a different policy that we refer to as a *regulatory sandbox* may be more effective. Under this policy, sectors with social damages below a threshold $\hat{\delta}$ ("inside the sandbox") can choose their technology freely, while sectors above the threshold are restricted from using the new technology until time $\hat{T}$. This policy allows the planner to ensure that sectors with high social damages adopt only after the new technology is established to be relatively safe. The next proposition demonstrates that the sandbox policy can improve upon the laissez-faire equilibrium.

Proposition 15. *Suppose α_i and η_i are constant and $\gamma_i < \delta_i$ across sectors. Then there exists a sandbox policy $\left(\hat{\delta}, \hat{T}\right)$ that strictly improves upon the laissez-faire equilibrium.*

In Appendix C.4, we provide additional details about optimal regulatory sandboxes and compare them to sector-independent taxes. In general, each of these policies can improve upon the laissez-faire equilibrium, and combining both is better: A sector-independent tax can differentially delay adoption for sectors with varying private damages γ_i, but it cannot alter the order of adoption. A regulatory sandbox can alter the order by delaying adoption for sectors with high social damages.

To implement welfare-improving use taxes or regulatory sandboxes, regulators must have some knowledge about the potential social damages from the new technology across different sectors. Although there is still substantial uncertainty about these damages from AI, assuming some knowledge is reasonable and consistent with current approaches to regulation. For example, the EU AI Act proposal outlines a framework in which an AI system is subject to different regulations depending on whether its intended use is considered "high risk" (e.g., the operation of critical infrastructure, employment and worker management processes, or law enforcement; see European Commission, 2021). The results in this section provide a foundation for this regulatory approach and also suggest that these policies should be updated as AI technologies become more capable and as society learns more about the risks.

3.6 Concluding Remarks

Advances in generative AI technologies, such as large language models, have intensified both hopes of more rapid economic growth and concerns about their potential negative consequences. Despite a robust public discussion on AI, there are currently no economic models of the regulation of transformative technologies. This paper has taken a first step in building such a model to provide novel insights for this debate.

We consider the adoption decision of a new, transformative technology that can increase productivity growth across all sectors of the economy but also raises risks of misuse, which

we model as the stochastic arrival of a "disaster". If a disaster occurs, some of the sectors using the new technology may be unable to switch back to the old, safe technology. Whether a disaster will occur is unknown, and society gradually learns about it over time. Consequently, adoption should be gradual and typically follows a convex path, initially growing slowly before accelerating later. Most surprisingly, a faster growth rate of the new technology should lead to slower adoption when potential damages are large: Although the planner is risk-neutral, she has a *precautionary motive* as irreversible damages imply that it is better to wait and learn about the likelihood of a disaster. These irreversible damages are greater when the new technology has a higher growth rate, strengthening the precautionary motive. Finally, if private firms internalize only part of the social damages from transformative technologies, equilibrium adoption is too fast and necessitates regulatory policies.

There are many interesting areas left for future work. First, in contrast to our baseline assumptions, early adoption may increase risks or may facilitate either general learning about potential misuses of the new technology or sector-specific learning about its "safe use". These considerations may motivate "experimentation" by adopting the technology in a few sectors or trying different uses in some areas, which is an important topic for future work.

Second, many of the misuses of new AI technologies depend on market structure and other aspects of regulation (e.g., concerning disinformation, discrimination, or privacy), and it would be interesting to explore how these affect optimal and equilibrium adoption.

Third, we simplified the analysis by assuming risk neutrality. C. I. Jones (2023) demonstrates that the extent of risk aversion and the precise form of damages have a first-order effect on the trade-off between higher growth and the likelihood of a disaster, and these can be incorporated in future analyses of learning about misuses of new transformative technologies.

Fourth, we abstracted from choices about how new technologies may be used. If regulations or other factors can prevent misuse of technology, then faster adoption can become optimal.

Finally, we showed that the optimal path of adoption depends on a few parameters, but there is currently a huge amount of uncertainty about their values. Careful empirical assessment of the costs and benefits of new, transformative technologies like generative AI is an obvious area for future research.

Appendix A

Appendix to Technology Choice, Spillovers, and the Concentration of R&D

A.1 Proofs

This appendix provides proofs for results in the main text.

A.1.1 Proofs for Section 1.3

Proof of Lemma 1: With linear, additively separable preferences over consumption sequences, the representative consumer's Euler equation and asset market clearing jointly imply $r(t) = \rho$. The intermediate price $p_\theta(q,t)$ must solve the maximization problem in the definition of flow profits (1.7), which implies a markup of $1/(1-\beta)$ over marginal cost: $p_\theta(q,t) = \gamma/(1-\beta)$. Given the final producer's intermediate demand curve, this implies the stated equations for intermediate quantities $x_\theta(q,t)$ and profits $\pi(q_\theta)$. Market-clearing for workers requires $L(t) = L$, and substituting the derived input quantities into the production function yields the state equation for output $Y(t)$. Market-clearing for goods then implies the corresponding equation for consumption $C(t)$. Finally, the wage for workers is recovered from the marginal product condition $w_L(t)L = \beta Y(t)$. ■

Proof of Proposition 1: The equations (1.17, 1.18) derived in the discussion before the Proposition characterize the growth rate g_θ^* and the relative quality distribution H_θ^*. I reproduce

these equations here:

$$g_\theta^* = \int (\lambda \sigma_E z_\theta + \sigma_I)\, \eta_\theta s dH_\theta^*(z_\theta), \tag{A.1}$$

$$0 = -\left[\left(\lambda - \frac{g_\theta^*}{\eta_\theta s}\right) z_\theta + \frac{\sigma_I}{\sigma_E}\right] \eta_\theta s h_\theta^*(z_\theta) + \delta N - H_\theta^*(z_\theta). \tag{A.2}$$

The evolution equation (1.15) for Q_θ additionally implies

$$g_\theta^* = \frac{K_\theta(t)}{Q_\theta(t)} \left(g_\theta^* + \delta N \sigma_E\right) - \delta,$$

which can be rearranged to give the expression for $Q_\theta(t)/K_\theta(t)$ stated in the Proposition. The inequality $Q_\theta(t)/K_\theta(t) < 1$ follows immediately from this expression and the assumed upper bound on entrant spillovers (1.19), recalling that market-clearing for scientists requires $S = Ns$.

To solve the differential equation (A.2) for H_θ^*, define the values

$$a \equiv \lambda \eta_\theta s - g_\theta^*,$$
$$b \equiv \frac{\sigma_I}{\sigma_E} \eta_\theta s.$$

Then (A.2) has the general solution

$$H_\theta^*(z_\theta) = \begin{cases} c\left(a z_\theta + b\right)^{-\frac{\delta}{a}} + N & a \neq 0, \\ c \exp\left(-\frac{\delta}{b} z_\theta\right) + N & a = 0. \end{cases}$$

The integration constant c is determined by the condition that H_θ^* equals zero at the lower bound of its support. To determine the support, consider the quality of an intermediate introduced at $t = 0$. In a BGP with $K_\theta(t) = K_\theta(0) \exp\left(g_\theta^* t\right)$, we can integrate the evolution equation (1.8) for q_θ to find

$$q_\theta(t) = \begin{cases} \left[-\frac{b}{a} + \left(1 + \frac{b}{a}\right) \exp(at)\right] \sigma_E K_\theta(t) & a \neq 0, \\ (1 + bt)\, \sigma_E K_\theta(t) & a = 0. \end{cases}$$

Hence

$$z_\theta(t) = \begin{cases} -\frac{b}{a} + \left(1 + \frac{b}{a}\right)\exp(at) & a \neq 0, \\ 1 + bt & a = 0. \end{cases}$$

With $a > 0$ or $b > 0$, z_θ is strictly increasing over time and limits to infinity. Since firms exit at a finite rate, the BGP distribution H^*_θ must then have support $[1, \infty)$. In fact, I show below that $a = 0$ implies $b > 0$, so that H^*_θ has support $[1, \infty)$ when $a \geq 0$. If instead $a < 0$, I conjecture and verify below that $b > -a$. The expression above then implies that z_θ is strictly increasing over time and limits to the finite value $-b/a > 1$, so the distribution H^*_θ has bounded support $[1, -b/a]$.

In all cases, the integration constant c is determined by the initial condition $H^*_\theta(1) = 0$. Solving this equation for c and substituting yields

$$H^*_\theta(z_\theta) = \begin{cases} N\left[1 - \left(1 + \frac{a}{\delta}\frac{z_\theta - 1}{\frac{a+b}{\delta}}\right)^{-\frac{\delta}{a}}\right] & a \neq 0, \\ N\left[1 - \exp\left(-\frac{\delta}{b}(z_\theta - 1)\right)\right] & a = 0. \end{cases}$$

Hence H^*_θ is a generalized Pareto distribution with location parameter 1, shape parameter $\varphi^*_\theta = \frac{a+b}{\delta}$, and tail parameter $\xi^*_\theta = \frac{a}{\delta}$. Provided that $\xi^*_\theta < 1$, the mean of this distribution (scaled by N) is

$$\int z_\theta dH^*_\theta(z_\theta) = N\frac{\delta + b}{\delta - a} = N\frac{\delta + \frac{\sigma_I}{\sigma_E}\eta_\theta s}{\delta - (\lambda\eta_\theta s - g^*_\theta)}$$

Substituting into (A.1), the growth rate g^*_θ must solve the fixed-point equation

$$g^*_\theta = \left[\lambda\frac{\sigma_E\delta + \sigma_I\eta_\theta s}{\delta - \left(\lambda\eta_\theta s - g^*_\theta\right)} + \sigma_I\right]\eta_\theta S.$$

Multiplying through by the denominator on the right side yields a quadratic equation in g^*_θ, and the solution is given by the expression for g^*_θ in the Proposition.[1] With this expression, we immediately observe that the assumption of positive spillovers from innovation $\lambda\sigma_E + \sigma_I > 0$ ensures that g^*_θ is positive, and in fact large enough to guarantee $\xi^*_\theta < 1$: This inequality holds if and only if $g^*_\theta > \lambda\eta_\theta s - \delta$. This is immediate if $\lambda\eta_\theta s < \delta$, and otherwise it follows from $\lambda\sigma_E + \sigma_I > 0$ after substituting the expression for g^*_θ from the Proposition.

[1] The lower root of the quadratic equation is negative and hence extraneous.

It remains to verify that $a = 0$ implies $b > 0$ and $b > -a$. For the former, note that $b \geq 0$, with equality if and only if $\sigma_I = 0$. If both $a = 0$ and $\sigma_I = 0$, the fixed-point equation for g_θ^* above simplifies to $s = \sigma_E S$, contradicting the assumed upper bound on entrant spillovers (1.19). To see that $b > -a$, note that this holds if and only if $\lambda + \frac{\sigma_I}{\sigma_E}\eta_\theta s > g_\theta^*$. After substituting the expression for g_θ^* from the Proposition, direct calculation shows that this inequality is also implied by the assumed upper bound on entrant spillovers (1.19). Note that this implies that the shape parameter φ_θ^* is positive.

To conclude, note that in all cases the assumed lower bound for ρ (1.20) ensures that the consumer's transversality condition (1.3) holds. Moreover, the expression for $q_\theta(t)$ derived avove implies an asymptotic growth rate of $\lambda\eta_\theta s$ for flow profits, so the assumed lower bound (1.14) on the exit rate δ ensures that the value function $V(q,t)$ remains finite. ■

Proof of Corollary 1: I first state the complete set of comparative statics for the tail parameter ξ_θ^*: The tail parameter ξ_θ^* is strictly decreasing in σ_E, σ_I, and S, and it is strictly increasing in λ and s. If σ_I is sufficiently close to zero, ξ_θ^* is strictly increasing in η_θ and strictly decreasing in δ, while the opposite holds for σ_I sufficiently large.

The comparative statics with respect to σ_E, σ_I, and S follow immediately by the formula for ξ_θ^* in Proposition 1. For the comparative static with respect to s, we can differentiate the equation for g_θ to find

$$\frac{\partial g_\theta^*}{\partial s} = \frac{\lambda\eta_\theta}{2}\left[1 - \frac{\frac{\delta - \lambda\eta_\theta s - \sigma_I\eta_\theta S}{2}}{\sqrt{\left(\frac{\delta - \lambda\eta_\theta s - \sigma_I\eta_\theta S}{2}\right)^2 + (\lambda\sigma_E + \sigma_I)\,\delta\eta_\theta S}}\right].$$

With $\lambda\sigma_E + \sigma_I > 0$ by assumption, the term in brackets is bounded strictly below 2. Thus $\partial g_\theta^*/\partial s < \lambda\eta_\theta$, and ξ_θ^* is strictly increasing in s. A similar calculation implies $\partial g_\theta^*/\partial\lambda < \eta_\theta s$, making use of the assumed upper bound to spillovers on entrants (1.19). This implies that ξ_θ^* is strictly increasing in λ. Finally, for η_θ,

$$\frac{\partial g_\theta^*}{\partial \eta_\theta} = \frac{\lambda s + \sigma_I S}{2} + \frac{-\frac{\lambda s + \sigma_I S}{2}\frac{\delta - \lambda\eta_\theta s - \sigma_I\eta_\theta S}{2} + (\lambda\sigma_E + \sigma_I)\,\delta S}{\sqrt{\left(\frac{\delta - \lambda\eta_\theta s - \sigma_I\eta_\theta S}{2}\right)^2 + (\lambda\sigma_E + \sigma_I)\,\delta\eta_\theta S}}.$$

If $\sigma_I = 0$, this value is below λs if and only if $s > \sigma_E S$, which holds by assumption (1.19). However, this value also limits to infinity as $\sigma_I \to \infty$. Thus ξ_θ^* is increasing in η_θ for σ_I small, while it is decreasing in η_θ for σ_I large. Since g_θ^* is linearly homogeneous in (δ, η_θ), the opposite comparative statics hold for δ: ξ_θ^* is decreasing in δ for σ_I small, while it is increasing

in δ for σ_I large.

Finally, note from the expression for the shape parameter φ_θ^* given in Proposition 6 that

$$\varphi_\theta^* = \xi_\theta^* + \frac{\sigma_I}{\sigma_E}\frac{\eta_\theta s}{\delta}.$$

The results above them imply that φ_θ^* is strictly increasing in s. ■

A.1.2 Proofs for Section 1.4

Proof of Proposition 2: We can directly integrate the quality evolution equation (1.8) to find that for $t' \geq t$,

$$q_\theta(t') = \exp\left(\lambda\eta_\theta \int_t^{t'} s_\theta(q(\tau),\tau)d\tau\right) q_\theta + \sigma_I \eta_\theta \int_t^{t'} s_\theta(q(t''),t'')K_\theta(t'')\exp\left(\lambda\eta_\theta \int_{t''}^{t'} s_\theta(q(\tau),\tau)d\tau\right) dt''.$$

Holding fixed the allocation of scientists $s(q(\tau),\tau)$ at $\tau \geq t$, we have

$$\frac{\partial q_\theta(t')}{\partial q_\theta} = \exp\left(\lambda\eta_\theta \int_t^{t'} s_\theta(q(\tau),\tau)d\tau\right).$$

Given an optimal allocation of scientists $s(q(\tau),\tau)$ at $\tau \geq t$ incumbent's value at t is

$$V(q,t) = \bar{\pi}\int_t^\infty \exp\big(-(\rho+\delta)(t'-t)\big)[q_A(t')+q_B(t')]dt'.$$

The Envelope Theorem then implies

$$\frac{\partial V(q,t)}{\partial q_\theta} = \bar{\pi}\int_t^\infty \exp\big(-(\rho+\delta)(t'-t)\big)\frac{\partial q_\theta(t')}{\partial q_\theta}dt' = \bar{\pi}\Psi_\theta(t).$$

The remaining results follow immediately. ■

Proof of Proposition 3: See Appendix A.2.2. ■

Proof of Proposition 4: See Appendix A.2.3. ■

Proof of Proposition 5: See the proof of Proposition A.3.4 in Appendix A.3. ■

Proof of Proposition 6: See the proof of Proposition A.3.5 in Appendix A.3. ■

A.2 Monotone Equilibria

This appendix characterizes and proves the existence of monotone equilibria.

A.2.1 Optimal Stopping: Proof of Lemma 2

To facilitate the analysis of monotone equilibria, I first prove several properties of solutions to the firm's problem (1.11) given an arbitrary but differentiable knowledge stock trajectory $[K(t)]_t$. Define the notation

$$\begin{aligned} k_\theta(t) &\equiv \lambda q_\theta(t) + \sigma_I K_\theta(t), \\ \Psi_\theta(T) &\equiv \int_T^\infty \exp\left(-\int_T^t [\rho + \delta - \lambda \eta_\theta s_\theta(q(\tau),\tau)]\, d\tau\right) dt, \\ \dot{k}_\theta(T-) &\equiv \lim_{t \uparrow T} \dot{k}(t). \end{aligned}$$

Note that $\Psi_\theta(T)$ satisfies the bounds

$$\Psi_\theta(T) \in \left[\frac{1}{\rho+\delta}, \frac{1}{\rho+\delta-\lambda\eta_\theta s}\right].$$

The following lemma provides a stopping time representation of solutions to the firm's problem (1.11), implying Lemma 2 in Section 1.4.1.

Lemma A.2.1. *Suppose $[K(t)]_t$ is differentiable, and suppose a piecewise-continuous solution $[s_\theta(q(t),t)]_t$ to (1.11) given initial qualities $q(t_0)$ at time $t_0 \geq 0$. Then without loss of generality $s_\theta(q(t),t) \in \{0,s\}$ for all $t \geq t_0$, and there exists a sequence of stopping times $t_0 < T_1 \leq T_2 \leq \ldots$ such that*

1. *if $s_\theta(q(t),t) = s$ for $t \in [T_m, T_{m+1})$, then $s_\theta(q(t),t) = 0$ for $t \in [T_{m+1}, T_{m+2})$;*
2. *every positive stopping time $T_m > 0$ satisfies the smooth-pasting condition*

$$k_B(T_m)\eta_B\Psi_B(T_m) = k_A(T_m)\eta_A\Psi_A(T_m);$$

3. *every positive stopping time $T_m > 0$ at which the firm transitions from A to B satisfies the necessary second-order condition*

$$\frac{\dot{k}_B(T_m-)}{k_B(T_m)} + \rho + \delta - \lambda\eta_B s - \frac{1}{\Psi_B(T_m)} \geq \frac{\dot{k}_A(T_m-)}{k_A(T_m)} + \rho + \delta - \frac{1}{\Psi_A(T_m)};\ \textit{and}$$

4. *every positive stopping time $T_m > 0$ at which the firm transitions from B to A satisfies the necessary second-order condition*

$$\frac{\dot{k}_B(T_m-)}{k_B(T_m)} + \rho + \delta - \frac{1}{\Psi_B(T_m)} \leq \frac{\dot{k}_A(T_m-)}{k_A(T_m)} + \rho + \delta - \lambda\eta_A s - \frac{1}{\Psi_A(T_m)}.$$

Proof of Lemma A.2.1: The optimality of corner allocations $s_\theta(q(t),t) \in \{0,s\}$ follows immediately from the linearity of the evolution equation (1.8) in $s_\theta(q,t)$. Hence any solution can be identified with a sequence of (potentially infinite) stopping times $t_0 < T_1 \leq T_2 \leq \ldots$ that prescribe when the firm should reverse its innovation direction.

Consider any positive stopping time $T_m > 0$, and suppose without loss of generality that the firm transitions from A to B at T_m. We can directly integrate the quality evolution equation (1.8) to find that for $t \geq T_m$,

$$\begin{aligned} q_\theta(t) = \exp\left(\lambda\eta_\theta \int_{T_m}^{t} s_\theta(q(\tau),\tau)d\tau\right) q_\theta(T_m) \\ + \sigma_I \eta_\theta \int_{T_m}^{t} s_\theta(q(t'),t')K_\theta(t')\exp\left(\lambda\eta_\theta \int_{t'}^{t} s_\theta(q(\tau),\tau)d\tau\right)dt'. \end{aligned}$$

The incumbent's value at t_0 is

$$V(q(t_0),t_0) = \bar{\pi}\int_{t_0}^{\infty} \exp\left(-(\rho+\delta)t\right)[q_A(t)+q_B(t)]\,dt.$$

The stopping time $T_m > 0$ must satisfy the interior first-order condition

$$0 = \frac{\partial}{\partial T_m}\frac{V(q(t_0),t_0)}{\bar{\pi}} = \int_{T_m}^{\infty} \exp\left(-(\rho+\delta)t\right)\left[\frac{\partial q_A(t)}{\partial T_m} + \frac{\partial q_B(t)}{\partial T_m}\right]dt.$$

For $t \geq T_m$, we can directly calculate

$$\frac{\partial q_\theta(t)}{\partial T_m} = \exp\left(\lambda\eta_\theta \int_{T_m}^{t} s_\theta(q(\tau),\tau)d\tau\right)[\dot{q}_\theta(T_m-) - \dot{q}_\theta(T_m+)].$$

Here $\dot{q}_\theta(T_m-)$ denotes the evolution of q_θ just before T_m, while $\dot{q}_\theta(T_m+)$ denotes the evolution just after T_m. The interior first-order condition for T_m simplifies to

$$0 = \sum_{\theta\in\{A,B\}} [\dot{q}_\theta(T_m-) - \dot{q}_\theta(T_m+)]\Psi_\theta(T_m).$$

Now given that the firm transitions from A to B at $T_m<$ we have

$$\dot{q}_A(T_m-) - \dot{q}_A(T_m+) = k_A(T_m)\eta_A s,$$
$$\dot{q}_\theta(T_m-) - \dot{q}_\theta(T_m+) = -k_B(T_m)\eta_B s.$$

Hence the interior first-order condition is exactly the smooth-pasting condition stated in the Lemma:

$$k_B(T_m)\eta_B\Psi_B(T_m) = k_A(T_m)\eta_A\Psi_A(T_m).$$

The second-order necessary condition for T_m requires that the left side of this equation be weakly increasing relative to the right side just before T_m. Log differentiating yields the necessary condition stated in the Lemma. The corresponding necessary condition for a transition from B to A is derived analogously. ■

A.2.2 Equilibrium Characterization: Proof of Proposition 3

In this section, I prove Proposition 3 and characterize all monotone equilibria converging to the BGP for technology B. I then state the analogous result for monotone equilibria converging to the BGP for A. The proof of this result is almost identical to that for Proposition 3, so I only sketch the differences.

Equilibrium Converging to B: Proof of Proposition 3

Step 0: T_E existence. As noted before the statement of Proposition 3, for any monotone equilibrium converging to B there exists a first time $T_E \geq 0$ at which entrants begin innovating immediately for technology B. The existence of this time is immediate: If the equilibrium converges to B, all entrants must eventually innovate for technology B exclusively. Hence there exists a smallest time $T_E < \infty$ after which this holds, and monotonicity implies that all entrants at $t \in [0, T_E)$ initially innovate for technology A. Any equilibrium with $T_E > 0$ is essentially an equilibrium with $T_E = 0$, but shifted forward in time. The distribution of relative qualities for A at T_E may stochastically dominate the initial distribution H_{A0} if the latter is not the BGP distribution H_A^*. But this simply implies incumbents are less willing to innovate for technology B at T_E, so whenever an equilibrium transitioning to B with $T_E > 0$ exists an equilibrium transitioning to B with $T_E = 0$ must also exist. This justifies my focus on the latter in the main text and throughout the proof below.

Step 1: $\chi(t)$ existence. To prove the existence of the cutoff $\chi(t)$, fix a time $t_0 \geq 0$ and a firm with qualities $q_\theta(t_0)$. Given trajectories for the knowledge stocks $[K(t)]_t$ and the firm's allocation of scientists $[s_\theta(q(t),t)]_{\theta,t}$, we can directly integrate the quality evolution equation (1.8) to find

$$q_\theta(t) = \exp\left(\lambda\eta_\theta \int_{t_0}^{t} s_\theta(q(\tau),\tau)d\tau\right) q_\theta(t_0) + \sigma_I \eta_\theta \int_{t_0}^{t} s_\theta(q(t'),t')K_\theta(t')\exp\left(\lambda\eta_\theta \int_{t'}^{t} s_\theta(q(\tau),\tau)d\tau\right) dt'.$$

This value is convex and supermodular in $[s_\theta(q(\tau),\tau)]_{\tau\in(t_0,t]}$, with increasing differences in $q_\theta(t_0)$ and $[s_\theta(q(\tau),\tau)]_{\tau\in(t_0,t]}$. Since flow profits are linear in qualities, the objective of the firm's problem (1.11) must also be supermodular in $[s_A(q(t),t)]_t$, with increasing differences in $q_A(t_0)$ and $[s_A(q(t),t)]_t$. Theorem 4 of Milgrom and Shannon (1994) then implies that the optimal value of $s_A(q(t),t)$ is non-decreasing in $q_A(t_0)$ at each $t \geq t_0$.

This observation immediately implies the existence of the cutoff $\chi(t)$: First setting $t_0 = 0$, we can observe that initial incumbents with higher relative qualities $z_A(0)$ transition later to technology B (if at all). The expression for $q_\theta(t)$ above implies that $q_A(t)$ is strictly increasing in $z_A(0)$ for these firms, so that the ratio

$$\frac{k_B(t)}{k_A(t)} = \frac{\lambda q_B(t) + \sigma_I K_B(t)}{\lambda q_A(t) + \sigma_I K_A(t)}$$

is always strictly decreasing in $z_A(0)$. We can simply define $\chi(t)$ as the least upper bound of this ratio across all initial incumbents still innovating for technology A.

Step 2: $\chi(t)$ characterization. The construction above implies that $\chi(0)$ must be weakly above $\frac{K_B(0)}{K_A(0)}$. To describe the evolution of $\chi(t)$, note that at $t \geq 0$ new entrants permanently innovate for technology B, while all remaining firms must choose whether to continue innovating for A or transition back to B. Lemma A.2.1 implies that any firm transitioning to B at T must satisfy the smooth-pasting condition

$$\frac{k_B(T)}{k_A(T)} = \frac{\eta_A}{\eta_B}\frac{\rho+\delta-\lambda\eta_B s}{\rho+\delta}, \tag{A.3}$$

with the left side increasing at T. I claim that these conditions are met at most once for each firm that initially innovates for technology A.

To prove this fact, it is helpful to first derive evolution equations for the knowledge stocks

$K(t)$. Let $T(z_A) \geq 0$ denote the transition time for a firm with initial relative quality $z_A(0)$. Let $\tilde{z}_A(t)$ denote the inverse of this function, taking value 1 outside its image. Given the initial distribution of relative qualities H_{A0}, the evolution equation (1.9) for K_A can be written

$$\dot{K}_A(t) = \exp(-\delta t) \int_{\tilde{z}_A(t)}^{\infty} [\lambda q_A(t; z_A) + \sigma_I K_A(t)] \eta_A s dH_{A0}(z_A),$$

where $q_A(t; z_A)$ denotes the quality for a firm with initial relative quality $z_A(0)$. Let $N_\theta(t)$ denote the total mass of firms innovating for technology θ at t, where

$$N_A(t) = \exp(-\delta t)[N - H_{A0}(\tilde{z}_A(t))] \quad \text{and} \quad N_B(t) = N - N_A(t). \tag{A.4}$$

Then we can differentiate again to derive the second-order differential equation

$$\begin{aligned} \ddot{K}_A(t) = &-[\delta - (\lambda + \sigma_I N_A(t))\eta_A s]\dot{K}_A(t) \\ &- \frac{\rho + \delta}{\rho + \delta - \lambda\eta_B s}[\lambda\sigma_E K_B(0) + \sigma_I K_B(t)]\eta_B s \exp(-\delta t) h_{A0}(\tilde{z}_A(t))\dot{\tilde{z}}_A(t). \end{aligned} \tag{A.5}$$

Turning to technology B, we can similarly write the evolution equation (1.9) as

$$\begin{aligned} \dot{K}_B(t) = \delta N &\int_{T_E}^{t} \exp(-\delta(t - t_0))[\lambda q_B(t; t_0) + \sigma_I K_B(t)]\eta_B s dt_0 \\ &+ \int_1^{\tilde{z}_A(t)} \exp(-\delta t)[\lambda q_B(t; z_A) + \sigma_I K_B(t)]\eta_B s dH_{A0}(z_A), \end{aligned}$$

where I abuse notation by letting $q_B(t; t_0)$ denote the quality for a firm that entered at time $t_0 \geq 0$ and $q_B(t; z_A)$ the quality of a firm with initial relative quality $z_A(0)$. Differentiating again yields

$$\begin{aligned} \ddot{K}_B(t) = &-[\delta - (\lambda + \sigma_I N_B(t))\eta_B s]\dot{K}_B(t) + \delta N(\lambda\sigma_E + \sigma_I)\eta_B s K_B(t) \\ &+ [\lambda\sigma_E K_B(0) + \sigma_I K_B(t)]\eta_B s \exp(-\delta t) h_{A0}(\tilde{z}_A(t))\dot{\tilde{z}}_A(t). \end{aligned} \tag{A.6}$$

The evolution equations (A.5, A.6) have two important implications for the remainder of the proof. First, (A.5) easily implies that $\ddot{K}_A(t) < 0$ because of the assumed lower bound (1.14) on δ. Second, (A.6) implies that the growth rate $g_B(t) \equiv \frac{\dot{K}_B(t)}{K_B(t)}$ is strictly increasing to its BGP value

g_B^*. This growth rate evolves according to

$$\begin{aligned}\dot{g}_B(t) &= \frac{\ddot{K}_B(t)}{K_B(t)} - g_B(t)^2 \\ &\geq \delta N(\lambda\sigma_E + \sigma_I)\eta_B s - [\delta - (\lambda + \sigma_I N_B(t))\eta_B s + g_B(t)]\, g_B(t).\end{aligned}$$

The right side of this inequality must be strictly positive for $t < \infty$: Define a new function $\check{g}(t)$ such that $\check{g}(0) = g_B(0)$, but with evolution equation

$$\dot{\check{g}}(t) = \delta N(\lambda\sigma_E + \sigma_I)\eta_B s - [\delta - (\lambda + \sigma_I N_B(0))\eta_B s + \check{g}_B(t)]\,\check{g}(t).$$

Direct calculation implies $\dot{\check{g}}(0) > 0$, making use of the assumed upper bound (1.19) on spillovers to entrants σ_E. With $\dot{N}_B(t) > 0$, we then have $\dot{\check{g}}(t) > 0$ and hence $\dot{g}_B(t) > 0$.

Returning to the smooth-pasting condition (A.3), consider any firm innovating for technology A at $t = 0$. If $\lambda\eta_A s > g_B^*$, it is easy to see that the firm never innovates for technology B: Lemma A.2.1 implies that the necessary second-order condition corresponding to the smooth-pasting condition (A.3) is

$$\frac{\sigma_I K_B(T)}{k_B(T)} g_B(T) \geq \lambda\eta_A s + \frac{\sigma_I K_A(T)}{k_A(T)} g_A(T). \tag{A.7}$$

But this inequality must be violated when $\lambda\eta_A s > g_B^*$, because the growth rate $g_B(T)$ is strictly lower than its BGP value g_B^* and $\sigma_I K_B(T) < k_B(T)$ by definition.[2]

Suppose instead $\lambda\eta_A s < g_B^*$ and $\sigma_I > 0$. I claim that any firm innovating for technology A at $t = 0$ must have a finite stopping time $T > 0$ after which it permanently innovates for technology B. The conditions $\lambda\eta_A s < g_B^*$ and $\sigma_I > 0$ easily imply that such a stopping time exists: With $g_B(t) \uparrow g_B^*$ and $g_A(t) \downarrow 0$, the technology B research productivity $k_B(t)$ eventually grows faster than the technology A research productivity $k_A(t)$. Moreover, it is straightforward to see that T is the unique solution to the smooth-pasting condition (A.3) that satisfies the corresponding second-order condition (A.7). While the firm innovates for technology A, the left side of this condition is strictly increasing because both $g_B(t)$ and $\sigma_I K_B(T)/k_B(T)$ are strictly increasing. The right side satisfies

$$\frac{\partial}{\partial T}\frac{\dot{K}_A(T)}{k_A(T)} = \frac{\dot{K}_A(T)}{k_A(T)}\left(\frac{\ddot{K}_A(T)}{\dot{K}_A(T)} - \frac{\lambda\eta_A s + \sigma_I \dot{K}_A(T)}{k_A(T)}\right)$$

[2]Here $g_B(t) < g_B^*$ follows by noting that at each time $t < \infty$, fewer than N firms innovate for B, and the distribution of relative qualities $H_B(z_B, t)$ is first-order stochastically dominated by the BGP distribution $H_B^*(z_B)$.

This value is negative because $\ddot{K}_A(T) < 0$. Thus for any firm that innovates for technology A at $t = 0$, the smooth-pasting and necesssary second-order conditions for a permanent transition to technology B are satisfied exactly once, and they provide a complete characterization of the transition time T.

Returning to the characterization of $\chi(t)$: The first initial incumbent to transition to technology B has initial knowledge ratio equal to $\chi(0)$, so we can define $\chi(t)$ to coincide with this firm's knowledge ratio $\frac{k_B(t)}{k_A(t)}$ until it reaches the right side of the smooth-pasting condition (A.7), after which χ remains constant. This ratio $\frac{k_B(t)}{k_A(t)}$ satisfies

$$\begin{aligned}\frac{d}{dt}\log\left(\frac{k_B(t)}{k_A(t)}\right) &= \frac{\sigma_I \dot{K}_B(t)}{k_B(t)} - \frac{\sigma_I \dot{K}_A(t)}{k_A(t)} - \lambda\eta_A s \\ &= \sigma_I \frac{\dot{K}_B(t) - \frac{k_B(t)}{k_A(t)}\dot{K}_A(t)}{\lambda K_B(0) + \sigma_I K_B(t)} - \lambda\eta_A s.\end{aligned}$$

Hence χ satisfies

$$\frac{\dot{\chi}(t)}{\chi(t)} = \left(\sigma_I \frac{\dot{K}_B(t) - \chi(t)\dot{K}_A(t)}{\lambda\sigma_E K_B(0) + \sigma_I K_B(t)} - \lambda\eta_A s\right) \mathbb{1}\left[\chi(t) \geq \frac{\eta_A}{\eta_B}\frac{\rho + \delta - \lambda\eta_B s}{\rho + \delta}\right]. \tag{A.8}$$

Moreover, the arguments above imply that the cutoff $\chi(t)$ yields a valid description of firm innovation decisions: For any firm innovating for technology A at $t = 0$, the ratio $\frac{k_B(t)}{k_A(t)}$ is strictly above $\chi(t)$ until it satisfies the smooth-pasting condition (A.7), at which time the firm transitions to B.

Step 3: $K(t)$ evolution. To conclude the proof, I provide a self-contained version of the evolution equations (A.5, A.6) for the knowledge stocks. Given $\chi(0)$, let $z_{A0} \geq 1$ denote the initial relative quality of the firm with initial knowledge ratio $\frac{k_B(0)}{k_A(0)}$ equal to $\chi(0)$:

$$\frac{\lambda\sigma_E + \sigma_I}{\lambda\sigma_E z_{A0} + \sigma_I}\frac{K_B(0)}{K_A(0)} = \chi(0). \tag{A.9}$$

Let T_0 denote the time at which this firm transitions to technology B. For $t \in [0, T_0)$, the knowledge stocks evolve according to

$$\begin{aligned}\ddot{K}_A(t) &= -\left[\delta - (\lambda + \sigma_I N_A(t))\,\eta_A s\right]\dot{K}_A(t), \\ \ddot{K}_B(t) &= -\left[\delta - (\lambda + \sigma_I N_B(t))\,\eta_B s\right]\dot{K}_B(t) + \delta N\left(\lambda\sigma_E + \sigma_I\right)\eta_B s K_B(t),\end{aligned}$$

where $N_A(t) = \exp(-\delta t)[N - H_{A0}(z_{A0})]$ and $N_B(t) = N - N_A(t)$. The initial conditions for these differential equations are

$$\begin{aligned} &K_A(0), \quad K_B(0), \\ &\frac{\dot{K}_A(0)}{K_A(0)} = \int_{z_{A0}}^{\infty} [\lambda\sigma_E z_A + \sigma_I]\,\eta_A s dH_{A0}(z_A), \\ &\frac{\dot{K}_B(0)}{K_B(0)} = [\lambda\sigma_E + \sigma_I]\,\eta_B s H_{A0}(z_{A0}). \end{aligned} \tag{A.10}$$

To determine the time T_0, for arbitrary $z_A \geq z_{A0}$ let $k_A(t; z_A)$ solve the differential equation

$$\dot{k}_A(t; z_A) = \lambda\eta_A s k_A(t; z_A) + \sigma_I \dot{K}_A(t)$$

with initial condition $k_A(0; z_A) = (\lambda\sigma_E z_A + \sigma_I) K_A(0)$. Then setting $z_A = z_{A0}$, if $\lambda\eta_A s < g_B^*$ the time T_0 is the unique time at which the ratio

$$\frac{\lambda\sigma_E K_B(0) + \sigma_I K_B(t)}{k_A(t; z_{A0})}$$

is increasing and satisfies the smooth-pasting condition (A.3). The inequality $t \geq T_0$ then holds if and only if

$$\begin{aligned} &\frac{\lambda\sigma_E K_B(0) + \sigma_I K_B(t)}{k_A(t; z_{A0})} \geq \frac{\eta_A}{\eta_B} \frac{\rho + \delta - \lambda\eta_B s}{\rho + \delta} \quad \text{and} \\ &\frac{\sigma_I \dot{K}_B(t)}{\lambda\sigma_E K_B(0) + \sigma_I K_B(t)} \geq \lambda\eta_A s + \frac{\sigma_I \dot{K}_A(t)}{k_A(t; z_{A0})}. \end{aligned} \tag{A.11}$$

If $\lambda\eta_A s > g_B^*$, the second inequality in this condition cannot be satisfied; hence $T_0 = \infty$.

For $t \geq T_0$, we can explicitly identify the initial relative quality $z_A(0)$ of firms transitioning to technology B. This function is simply $\tilde{z}_A(t)$ defined above, and for $t \geq T_0$ it satisfies the smooth-pasting condition

$$\frac{\lambda\sigma_E K_B(0) + \sigma_I K_B(t)}{k_A(t; \tilde{z}_A(t))} = \frac{\eta_A}{\eta_B} \frac{\rho + \delta - \lambda\eta_B s}{\rho + \delta}.$$

Integrating the evolution equation for $k_A(t; z_A)$ reveals that this function is linear in z_A:

$$k_A(t; z_A) = k_A(t; z_{A0}) + (z_A - z_{A0}) \exp(\lambda\eta_A s t)\, \lambda\sigma_E K_A(0).$$

Substituting into the previous equation yields an implicit equation for $\tilde{z}_A(t)$:

$$\frac{\lambda\sigma_E K_B(0)+\sigma_I K_B(t)}{k_A(t;z_{A0})+(\tilde{z}_A(t)-z_{A0})\exp(\lambda\eta_A s t)\,\lambda\sigma_E K_A(0)}=\frac{\eta_A}{\eta_B}\frac{\rho+\delta-\lambda\eta_B s}{\rho+\delta}.$$

Solving for $\tilde{z}_A(t)$ and recalling the characterization (A.11) of T_0, we can write

$$\begin{aligned}&\tilde{z}_A(t) \qquad\qquad (A.12)\\ &=z_{A0}+\frac{\exp(-\lambda\eta_A s t)}{\lambda\sigma_E K_A(0)}\left\{[\lambda\sigma_E K_B(0)+\sigma_I K_B(t)]\frac{\eta_B}{\eta_A}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_B s}-k_A(t;z_{A0})\right\}\mathbb{1}[(A.11)].\end{aligned}$$

To summarize, I find that the knowledge stocks satisfy the dynamical system

$$\begin{cases}\ddot{K}_A(t) &= -[\delta-(\lambda+\sigma_I N_A(t))\eta_A s]\dot{K}_A(t)\\ &\quad -\frac{\rho+\delta}{\rho+\delta-\lambda\eta_B s}[\lambda\sigma_E K_B(0)+\sigma_I K_B(t)]\eta_B s\exp(-\delta t)h_{A0}(\tilde{z}_A(t))\dot{\tilde{z}}_A(t),\\ \ddot{K}_B(t) &= -[\delta-(\lambda+\sigma_I N_B(t))\eta_B s]\dot{K}_B(t)+\delta N(\lambda\sigma_E+\sigma_I)\eta_B s K_B(t)\\ &\quad +[\lambda\sigma_E K_B(0)+\sigma_I K_B(t)]\eta_B s\exp(-\delta t)h_{A0}(\tilde{z}_A(t))\dot{\tilde{z}}_A(t),\\ \dot{k}_A(t;z_{A0}) &= \lambda\eta_A s k_A(t;z_{A0})+\sigma_I\dot{K}_A(t),\end{cases} \qquad (A.13)$$

where $\tilde{z}_A$ is given by (A.12), z_{A0} is given by (A.9), and $N_A(t)$ and $N_B(t)$ are given by (A.4). The initial conditions are (A.10) and $k_A(0;z_{A0})=(\lambda\sigma_E z_{A0}+\sigma_I)K_A(0)$.

Step 4: H_{A0} comparative static. The comparative static with respect to H_{A0} is monotone: If $\sigma_I=0$, firm innovation decisions are independent of all other firms after entry, so the cutoff $\chi(t)$ is invariant to H_{A0}.[3] If $\sigma_I>0$, initial incumbents are more inclined to start innovating for technology A when the trajectory $K_A(t)$ is larger and the trajectory $K_B(t)$ is lower. With innovation decisions held fixed, a first-order stochastically increasing shift in the distribution H_{A0} implies a larger growth rate for K_A at fixed innovation decisions, and hence a larger level since $K_A(0)$ is fixed. This induces more incumbents to innovate for technology A, producing an upward shift in the trajectory of $K_A(t)$ and a downward shift in the trajectory of $K_B(t)$. These effects are mutually reinforcing, so $\chi(t)$ must be strictly increasing in H_{A0}. ■

Equilibrium Converging to A

The following proposition characterizes all monotone equilibria converging to the BGP for technology A, analogous to Proposition 3 in Section 1.4.2:

[3] See also the explicit solution for this case in Section 1.4.3 or Appendix A.3.

Proposition A.2.1. *In any monotone equilibrium converging to the BGP for technology A, there exists a cutoff $\chi(t)$ such that a firm innovates for B if and only if*

$$\frac{k_B(t)}{k_A(t)} \geq \chi(t).$$

There exists a time $T_E \geq 0$ such that $\chi(t) = \min\{\chi(0), \kappa(t)\}$ for $t \in [0, T_E]$, with $\chi(T_E) = \kappa(T_E)$. For $t \in (T_E, \infty)$, the cutoff solves the differential equation

$$\frac{\dot{\chi}(t)}{\chi(t)} = \left(\lambda \eta_B s + \sigma_I \frac{\frac{1}{\chi(t)} \dot{K}_B(t) - \dot{K}_A(t)}{\lambda \sigma_E K_A(T_E) + \sigma_I K_A(t)} \right) \mathbb{1}\left[\chi(t) \leq \frac{\eta_A}{\eta_B} \frac{\rho + \delta}{\rho + \delta - \lambda \eta_A s} \right]. \tag{A.14}$$

The equilibrium is unique up to the parameters $\chi(0)$ and $T_E \geq 0$. The knowledge stocks $K(t)$ are the solutions to the dynamical systems (A.22, A.24).

Note several differences from Proposition 3: First, firms choose B provided that the weak inequality $\frac{k_B(t)}{k_A(t)} \geq \chi(t)$ holds, instead of the strong inequality $\frac{k_B(t)}{k_A(t)} > \chi(t)$. This ensures that entrants begin innovating for technology B for $t \leq T_E$. The evolution equation (A.14) ensures that $\chi(t)$ equals the knowledge ratio $\frac{k_B(t)}{k_A(t)}$ for the last entrant to innovate first for technology B, until the following smooth-pasting condition is satisfied:

$$\frac{k_B(t)}{k_A(t)} = \frac{\eta_A}{\eta_B} \frac{\rho + \delta}{\rho + \delta - \lambda \eta_A s}. \tag{A.15}$$

After this time, $\chi(t)$ remains equal to the expression on the right side.

Sketch of Proposition A.2.1. I only show how to derive the dynamical system for the knowledge stocks $K(t)$; the remaining parts of the proof are analogous to that for Proposition 3.

Note first that at each time $t > T_E$, there are potentially two groups of firms that began innovating for technology B but are now transitioning back to technology A: (i) initial incumbents with $z_A(0) < z_{A0}$, where z_{A0} satisfies

$$\frac{\lambda \sigma_E + \sigma_I}{\lambda \sigma_E z_{A0} + \sigma_I} \frac{K_B(0)}{K_A(0)}; \tag{A.16}$$

and (ii) entrants at each time $t_0 \in [0, T_E)$. We can track the transitions by initial incumbents just as in equilibria converging to technology B. Let $k_B(t)$ solve the differential equation

$$\dot{k}_B(t) = \lambda \eta_B s k_B(t) + \sigma_I \dot{K}_B(t),$$

with initial condition $k_B(0) = (\lambda\sigma_E + \sigma_I)K_B(0)$. The initial incumbent with $z_A = z_{A0}$ must be the first to transition back to technology A (if at all). If $\lambda\eta_B s < g_A^*$, this transition time is the unique time at which the ratio

$$\frac{k_B(t)}{\lambda z_{A0}\sigma_E K_A(0) + \sigma_I K_A(t)}$$

is decreasing and satisfies the smooth-pasting condition (A.15). Hence this time has passed if and only if

$$\begin{aligned} \frac{k_B(t)}{\lambda z_{A0}\sigma_E K_A(0) + \sigma_I K_A(t)} &\leq \frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s} \quad \text{and} \\ \lambda\eta_B s + \frac{\sigma_I \dot{K}_B(t)}{k_B(t)} &\leq \frac{\sigma_I \dot{K}_A(t)}{\lambda z_{A0}\sigma_E K_A(0) + \sigma_I K_A(t)}. \end{aligned} \tag{A.17}$$

For each time t after these conditions are satisfied, let $\tilde{z}_A(t)$ denote the initial relative quality of the incumbent transitioning back to technology A. Then $\tilde{z}_A(t)$ must solve the smooth-pasting condition

$$\frac{k_B(t)}{\lambda\sigma_E\tilde{z}_A(t)K_A(0) + \sigma_I K_A(t)} = \frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s}.$$

Setting $\tilde{z}(t) = z_{A0}$ before incumbents begin transitioning back to A, we can solve to find

$$\tilde{z}_A(t) = \begin{cases} z_{A0} & \text{if (A.17) not satisfied,} \\ \frac{1}{\lambda\sigma_E K_A(0)}\left[\frac{\eta_B}{\eta_A}\frac{\rho+\delta-\lambda\eta_A s}{\rho+\delta}k_B(t) - \sigma_I K_A(t)\right] & \text{if (A.17) satisfied.} \end{cases} \tag{A.18}$$

Now consider entrants at time $t_0 \in [0, T_E)$ that potentially transition back to technology A. Let $k_B(t; t_0)$ solve the differential equation

$$\dot{k}_B(t; t_0) = \lambda\eta_B s k_B(t; t_0) + \sigma_I \dot{K}_B(t),$$

with initial condition $k_B(t; t_0) = (\lambda\sigma_E + \sigma_I)K_B(t_0)$. The entrant with $t_0 = T_E$ must be the first to transition back (if at all). If $\lambda\eta_B s < g_A^*$, this is the unique time at which the ratio

$$\frac{k_B(t; T_E)}{\lambda\sigma_E K_A(T_E) + \sigma_I K_B(t)}$$

is decreasing and satisfies the smooth-pasting condition (A.15). Let $\bar{T}$ denote this time. We

have $t \geq \bar{T}$ if and only if

$$\frac{k_B(t;T_E)}{\lambda\sigma_E K_A(T_E)+\sigma_I K_A(t)} \leq \frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s} \quad \text{and} \tag{A.19}$$
$$\lambda\eta_B s + \frac{\sigma_I \dot{K}_B(t)}{k_B(t;T_E)} \leq \frac{\sigma_I \dot{K}_A(t)}{\lambda\sigma_E K_A(T_E)+\sigma_I K_A(t)}.$$

For $t \geq \bar{T}$, let $\tilde{t}_0(t)$ denote the entry time of the firms transitioning back to technology A. Then $\tilde{t}_0(t)$ must solve the smooth-pasting condition

$$\frac{k_B(t;\tilde{t}_0(t))}{\lambda\sigma_E K_A(\tilde{t}_0(t))+\sigma_I K_A(t)} = \frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s}. \tag{A.20}$$

To derive a differential equation for $\tilde{t}_0(t)$, note first that we can integrate the evolution equation for $k_B(t;t_0)$ to write

$$k_B(t;t_0) = \exp\left(\lambda\eta_B s(t-t_0)\right)(\lambda\sigma_E+\sigma_I)K_B(t_0) + \int_{t_0}^{t} \exp\left(\lambda\eta_B s(t-\tau)\right)\sigma_I \dot{K}_B(\tau)d\tau.$$

Differentiating in t_0 yields

$$\frac{\partial k_B(t;t_0)}{\partial t_0} = \exp\left(\lambda\eta_B s(t-t_0)\right)\left[-\lambda\eta_B s\left(\lambda\sigma_E+\sigma_I\right)K_B(t_0) + \lambda\sigma_E \dot{K}_B(t_0)\right].$$

We can then differentiate the smooth-pasting condition (A.20) to find that for $t \geq \bar{T}$,

$$\dot{\tilde{t}}_0(t) \tag{A.21}$$
$$= \frac{\lambda\eta_B s \frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s}\left[\lambda\sigma_E K_A(\tilde{t}_0(t))+\sigma_I K_A(t)\right] + \sigma_I \dot{K}_B(t) - \frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s}\sigma_I \dot{K}_A(t)}{\frac{\eta_A}{\eta_B}\frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s}\lambda\sigma_E \dot{K}_A(\tilde{t}_0(t)) - \exp\left(\lambda\eta_B s(t-\tilde{t}_0(t))\right)\left[\lambda\sigma_E \dot{K}_B(\tilde{t}_0(t)) - \lambda\eta_B s\left(\lambda\sigma_E+\sigma_I\right)K_B(\tilde{t}_0(t))\right]}.$$

By the same calculation as in the proof of Proposition 3, for $t < \bar{T}$ the knowledge stocks satisfy

the dynamical system

$$\begin{cases} \ddot{K}_A(t) &= -[\delta - (\lambda + \sigma_I N_A(t))\,\eta_A s]\dot{K}_A(t) + \delta N(\lambda\sigma_E + \sigma_I)\eta_A s K_A(t)\mathbb{1}[t \geq T_E] \\ &\quad + \frac{\rho+\delta-\lambda\eta_A s}{\rho+\delta} k_B(t)\eta_B s \exp(-\delta t)\, h_A^*(\tilde{z}_A(t))\left(-\dot{\tilde{z}}_A(t)\right), \\ \ddot{K}_B(t) &= -[\delta - (\lambda + \sigma_I N_B(t))\,\eta_B s]\dot{K}_B(t) + \delta N(\lambda\sigma_E + \sigma_I)\eta_B s K_B(t)\mathbb{1}[t < T_E] \\ &\quad - k_B(t)\eta_B s \exp(-\delta t)\, h_A^*(\tilde{z}_A(t))\left(-\dot{\tilde{z}}_A(t)\right), \\ \dot{k}_B(t) &= \lambda\eta_B s k_B(t) + \sigma_I \dot{K}_B(t), \\ \dot{N}_A(t) &= \delta[N\mathbb{1}[t \geq T_E] - N_A(t)] + \exp(-\delta t)\, h_A^*(\tilde{z}_A(t))\left(-\dot{\tilde{z}}_A(t)\right), \end{cases} \tag{A.22}$$

where $\tilde{z}_A$ is given by (A.18) and z_{A0} is given by (A.16). The initial conditions are

$$\begin{aligned} & K_A(0), \quad K_B(0), \\ & \frac{\dot{K}_A(0)}{K_A(0)} = \int_{z_{A0}}^{\infty} [\lambda\sigma_E z_A + \sigma_I]\,\eta_A s\, dH_A^*(z_A), \\ & \frac{\dot{K}_B(0)}{K_B(0)} = [\lambda\sigma_E + \sigma_I]\,\eta_B s H_A^*(z_{A0}), \end{aligned} \tag{A.23}$$

as well as $k_B(0) = (\lambda\sigma_E + \sigma_I)K_B(0)$ and $N_A(0) = N - H_A^*(z_{A0})$.

Solving the dynamical system (A.22) forward, the time $\bar{T}$ is the first time that (A.19) is satisfied. For $t \geq \bar{T}$, the knowledge stocks instead satisfy

$$\begin{cases} \ddot{K}_A(t) &= -[\delta - (\lambda + \sigma_I N_A(t))\,\eta_A s]\dot{K}_A(t) + \delta N(\lambda\sigma_E + \sigma_I)\eta_A s K_A(t) \\ &\quad + \frac{\rho+\delta-\lambda\eta_A s}{\rho+\delta} k_B(t)\eta_B s \exp(-\delta t)\, h_A^*(\tilde{z}_A(t))\left(-\dot{\tilde{z}}_A(t)\right) \\ &\quad + [\lambda\sigma_E K_A(\tilde{t}_0(t)) + \sigma_I K_A(t)]\,\eta_A s \exp(-\delta(t - \tilde{t}_0(t)))\,\delta N\left(-\dot{\tilde{t}}_0(t)\right), \\ \ddot{K}_B(t) &= -[\delta - (\lambda + \sigma_I N_B(t))\,\eta_B s]\dot{K}_B(t) \\ &\quad - k_B(t)\eta_B s \exp(-\delta t)\, h_A^*(\tilde{z}_A(t))\left(-\dot{\tilde{z}}_A(t)\right) \\ &\quad - \frac{\rho+\delta}{\rho+\delta-\lambda\eta_A s}[\lambda\sigma_E K_A(\tilde{t}_0(t)) + \sigma_I K_A(t)]\,\eta_A s \exp(-\delta(t - \tilde{t}_0(t)))\,\delta N\left(-\dot{\tilde{t}}_0(t)\right), \\ \dot{k}_B(t) &= \lambda\eta_B s k_B(t) + \sigma_I \dot{K}_B(t), \\ \dot{N}_A(t) &= \delta[N - N_A(t)] + \exp(-\delta t)\, h_A^*(\tilde{z}_A(t))\left(-\dot{\tilde{z}}_A(t)\right) \\ &\quad + \exp(-\delta(t - \tilde{t}_0(t)))\,\delta N\left(-\dot{\tilde{t}}_0(t)\right), \\ \dot{\tilde{t}}_0(t) &= \text{(A.21)}, \end{cases} \tag{A.24}$$

where the initial condition for $\tilde{t}_0$ is $\tilde{t}_0(\bar{T}) = T_E$. ■

A.2.3 Equilibrium Existence: Proof of Proposition 4

In this section, I provide a full proof of the existence of monotone equilibria converging to the BGP for technology B, given the assumptions of Proposition 4. I sketch the analogous proof for equilibria converging to the BGP for A, which is almost identical. Finally, I show that multiple equilibria arise whenever $\lambda, \sigma_I > 0$, and I prove comparative statics for the thresholds κ_A^* and κ_B^* in the initial distribution H_{A0}.

Equilibrium Converging to *B*

Each monotone equilibrium is uniquely determined by the initial cutoff $\chi(0) \geq \frac{K_B(0)}{K_A(0)}$. I let z_{A0} denote the initial relative quality for an initial incumbent at the cutoff:

$$\frac{\lambda\sigma_E + \sigma_I}{\lambda\sigma_E z_{A0} + \sigma_I}\kappa(0) = \chi(0).$$

The following proposition proves the existence of a monotone equilibrium converging to the BGP for technology B under weaker conditions than assumed in Proposition 4:

Proposition A.2.2. *There exists a threshold κ_B^* such that a monotone equilibrium converging to the BGP for technology B exists if the following hold:*

1. *For each $t \geq 0$*

$$\lambda\eta_B s + \frac{\sigma_I \dot{K}_B(t)}{k_B(t)} \geq \frac{\sigma_I \dot{K}_A(t)}{k_A(t)}, \tag{A.25}$$

   where $\dot{k}_A(t) = \sigma_I \dot{K}_A(t)$ and $\dot{k}_B(t) = \lambda\eta_B s k_B(t) + \sigma_I \dot{K}_B(t)$ with $k_\theta(0) = (\lambda\sigma_E + \sigma_I)K_\theta(0)$;

2. *$\kappa(0) \geq \kappa_B^*$.*

The initial cutoff $\chi(0)$ is such that z_{A0} is a stable solution to the fixed-point equation (A.26). The threshold κ_B^* is strictly increasing in H_{A0} when $\lambda > 0$.

Proof of Proposition A.2.2: Fix an initial cutoff $\chi(0)$ and the corresponding initial relative quality z_{A0}. The purpose of the proof is to demonstrate that the innovation decisions implied by the resulting cutoff $\chi(t)$ are privately optimal for all firms.

First consider an initial incumbent with relative quality $z_A \geq z_{A0}$. This firm initially innovates for technology A, and the arguments in Step 3 of the proof of Proposition 3 demonstrate that the innovation decisions implied by the cutoff $\chi(t)$ are optimal.

Now consider an initial incumbent with relative quality $z_A < z_{A0}$. This firm initially innovates for technology B, and the innovation decisions implied by the cutoff require that it never innovates for technology A. The firm's research productivities then evolve according to

$$\frac{\dot{k}_A(t)}{k_A(t)} = \frac{\sigma_I \dot{K}_A(t)}{k_A(t)} \quad \text{and} \quad \frac{\dot{k}_B(t)}{k_B(t)} = \lambda \eta_B s + \frac{\sigma_I \dot{K}_B(t)}{k_B(t)}.$$

If the incumbent reverted back to technology A at some time T, Lemma A.2.2 would require that k_A and k_B satisfy the smooth-pasting condition

$$k_B(T)\eta_B \Psi_B(T) = k_A(T)\eta_A \Psi_A(T),$$

with the corresponding second-order necessary condition

$$\lambda \eta_B s + \frac{\sigma_I \dot{K}_B(t)}{k_B(t)} + \rho + \delta - \frac{1}{\Psi_B(T)} \leq \frac{\sigma_I \dot{K}_A(t)}{k_A(t)} + \rho + \delta - \lambda \eta_A s - \frac{1}{\Psi_A(T)}.$$

Given the bounds $\Psi_\theta(T) \in \left[\frac{1}{\rho+\delta}, \frac{1}{\rho+\delta-\lambda\eta_\theta s}\right]$, this inequality is ruled out precisely by (A.25).

Finally, note that entrants at $t_0 \geq 0$ innovate for technology B permanently if and only if this yields higher value than innovating initially for technology A and transitioning to technology B at a later stopping time $T(t_0) \geq t_0$. This stopping time is either ∞ (if $\sigma_I = 0$ or $\lambda \eta_A s \geq g_B^*$) or is given by the unique solution to the smooth-pasting condition

$$\frac{k_B(T(t_0); t_0)}{k_A(T(t_0); t_0)} = \frac{\eta_A}{\eta_B} \frac{\rho + \delta - \lambda \eta_B s}{\rho + \delta},$$

where k_A and k_B satisfy the evolution equations

$$\frac{\dot{k}_A(t; t_0)}{k_A(t; t_0)} = \lambda \eta_A s + \frac{\sigma_I \dot{K}_A(t)}{k_A(t; t_0)} \quad \text{and} \quad \frac{\dot{k}_B(t; t_0)}{k_B(t; t_0)} = \frac{\sigma_I \dot{K}_B(t)}{k_B(t; t_0)},$$

with initial conditions $k_\theta(t_0; t_0) = (\lambda \sigma_E + \sigma_I) K_\theta(t_0)$.

If an entrant at t_0 permanently innovates for technology B, let $q_\theta^B(t; t_0)$ denote the corresponding quality for technology θ at time t. These qualities satisfy the evolution equations

$$\begin{aligned} \dot{q}_A^B(t; t_0) &= 0 \\ \dot{q}_B^B(t; t_0) &= \left[\lambda q_B^B(t; t_0) + \sigma_I K_B(t)\right] \eta_B s, \end{aligned}$$

with initial conditions $q_\theta^B(t_0, t_0) = \sigma_E K_\theta(t_0)$. If an entrant instead innovates initially for

technology A, let $q_\theta^A(t;t_0)$ denote the corresponding quality for technology θ. The evolution equations are now

$$\dot{q}_A^A(t;t_0) = \mathbb{1}_{t<T(t_0)}\left[\lambda q_A^A(t;t_0) + \sigma_I K_A(t)\right]\eta_A s$$
$$\dot{q}_B^B(t;t_0) = \mathbb{1}_{t\geq T(t_0)}\left[\lambda q_B^B(t;t_0) + \sigma_I K_B(t)\right]\eta_B s,$$

with the same initial conditions as above. The entrant's value when choosing to innovate technology θ initially is

$$V_E^\theta(t_0) \equiv \bar{\pi}\int_{t_0}^{\infty} \exp\left(-(\rho+\delta)(t-t_0)\right)\left[q_A^\theta(t;t_0) + q_B^\theta(t;t_0)\right]dt.$$

Entrants at t_0 innovate for B if and only if $V_E^B(t_0) \geq V_E^A(t_0)$.

To determine when this condition holds, note first that the dynamical system (A.13) that describes the evolution of the knowledge stocks is linearly homogeneous in $K(t)$ conditional on z_{A0}. As a result, there exists a function $\tilde{K}(t;z_{A0},\kappa(0))$ that depends on initial conditions only through z_{A0} and $\kappa(0)$ such that

$$K(t) = \tilde{K}(t;z_{A0},\kappa(0))K_A(0).$$

This normalization implies $\tilde{K}_A(0;z_{A0},\kappa(0)) = 1$ and $\tilde{K}_B(0;z_{A0},\kappa(0)) = \kappa(0)$. To understand how these normalized knowledge stocks depend on z_{A0} and $\kappa(0)$, hold $\kappa(0)$ fixed and consider an increase in z_{A0}, or equivalently a decrease in $\chi(0)$. This adjustment raises the initial growth rate of $\tilde{K}_B$ and lowers the initial growth rate of $\tilde{K}_A$, implying that initial incumbents transition more rapidly to B. These effects are mutually reinforcing, and they imply that $\tilde{K}_A$ and $\tilde{K}_B$ are respectively strictly decreasing and strictly increasing in z_{A0} at $t > 0$. Now hold z_{A0} fixed, and consider an increase in $\kappa(0)$. This leaves the initial growth rates of $\tilde{K}_A$ and $\tilde{K}_B$ unaffected, and if all initial incumbents with $z_A > z_{A0}$ never transition to B, it only raises $\tilde{K}_B$ for $t \geq 0$ while leaving $\tilde{K}_A$ unchanged. But if initial incumbents eventually transition to B, the increase in $\tilde{K}_B$ induces the incumbents to transition more rapidly. This leads to a further increase in $\tilde{K}_B$ and a decrease in $\tilde{K}_A$ for t sufficiently large.

These linear homogeneity and comparative dynamics observations carry over directly to the qualities $q_\theta^{\theta'}(t;t_0)$ and the values $V_E^\theta(t_0)$. In particular, the difference

$$\frac{V_E^B(t_0)}{K_A(0)} - \frac{V_E^A(t_0)}{K_A(0)}$$

depends on the initial conditions $K(0)$ only through $\kappa(0)$, and the Envelope Theorem implies that it is strictly increasing in $\kappa(0)$ and z_{A0} when $T(t_0) > 0$. With $g_B(t) \uparrow 1$ and $g_A(t) \downarrow 0$, there also exists a finite time t_0 after which the difference is (weakly) positive. The Intermediate Value Theorem implies a function $\kappa^E(t, z_{A0})$ such that $V_E^B(t_0) \geq V_E^A(t_0)$ if and only if

$$\kappa(0) \geq \kappa^E(t_0, z_{A0}),$$

where κ^E is strictly decreasing in z_{A0} and strictly decreasing in t_0 for t_0 sufficiently large.[4]

Now consider optimality for initial incumbents. Given an incumbent with relative quality z_A, denote the value of innovating initially for technology θ by

$$\tilde{V}^\theta(z_A; z_{A0}, \kappa(0))K_A(0),$$

where the same arguments as above imply linear homogeneity in $[K(t)]_t$. Similarly, the difference

$$\Delta(z_A; z_{A0}, \kappa(0)) \equiv \tilde{V}^B(z_A; z_{A0}, \kappa(0)) - \tilde{V}^A(z_A; z_{A0}, \kappa(0))$$

is also strictly increasing in z_{A0} and $\kappa(0)$ when the firm's transition time is positive, while the Envelope Theorem implies that it is strictly decreasing in z_A (again when the transition time is positive). In equilibrium, $z_{A0} \geq 1$ must be such that this difference is zero:[5]

$$\Delta(z_{A0}; z_{A0}, \kappa(0)) = 0. \tag{A.26}$$

Note first that a solution can only exist for $\kappa(0)$ sufficiently large: (A.26) is trivially violated

[4] In the case with $\sigma_I = 0$ or $\lambda\eta_A s \geq g_B$, we can derive an explicit formula:

$$\kappa^E(t, z_{A0}) = \frac{\frac{\lambda\eta_A s}{\rho+\delta-\lambda\eta_A s}\frac{\sigma_E \tilde{K}_A(t;z_{A0})}{\rho+\delta} + \int_t^\infty \exp\left(-(\rho+\delta)(\tau-t)\right)\frac{\eta_A s\sigma_I \tilde{K}_A(\tau;z_{A0})}{\rho+\delta-\lambda\eta_A s}d\tau}{\frac{\lambda\eta_B s}{\rho+\delta-\lambda\eta_B s}\frac{\sigma_E \tilde{K}_B(t;z_{A0})}{\rho+\delta} + \int_t^\infty \exp\left(-(\rho+\delta)(\tau-t)\right)\frac{\eta_B s\sigma_I \tilde{K}_B(t;z_{A0})}{\rho+\delta-\lambda\eta_B s}d\tau}.$$

Here $\tilde{K}_A(t; z_{A0}) \equiv K_A(t)/K_A(0)$ as in the text of the proof, while for this expression I instead define $\tilde{K}_B(t; z_{A0}) \equiv K_B(t)/K_B(0)$. I can normalize both knowledge stocks independently because the function $\tilde{z}_A(t)$ described in the dynamical system (A.13) is constant in this case, so the differential equations for K are not interdependent conditional on z_{A0}.

[5] Using the same assumptions and notation as in footnote 4, this equation can be written explicitly:

$$\begin{aligned}\frac{\lambda\eta_A s}{\rho+\delta-\lambda\eta_A s}\frac{\sigma_E}{\rho+\delta}z_{A0} = &\left[\frac{\lambda\eta_B s}{\rho+\delta-\lambda\eta_B s}\frac{\sigma_E}{\rho+\delta} + \int_0^\infty \exp\left(-(\rho+\delta)\tau\right)\frac{\eta_B s\sigma_I \tilde{K}_B(\tau;z_{A0})}{\rho+\delta-\lambda\eta_B s}d\tau\right]\kappa(0)\\ &- \int_0^\infty \exp\left(-(\rho+\delta)\tau\right)\frac{\eta_A s\sigma_I \tilde{K}_A(\tau;z_{A0})}{\rho+\delta-\lambda\eta_A s}d\tau.\end{aligned}$$

when $\kappa(0) = 0$. Fixing any $z_A \geq 1$, we have that $\Delta(z_A; z_A, \kappa(0)) \to \infty$ as $\kappa(0) \to \infty$, so we can take any $\kappa(0)$ such that $\Delta(z_A; z_A, \kappa(0)) > 0$. It is also straightforward to see that $\Delta(z_A; z_A, \kappa(0)) \to -\infty$ as $z_A \to \infty$: Even with all initial incumbents innovating for technology B, an incumbent's value from doing so remains uniformly bounded. By the Intermediate Value Theorem, these arguments imply that for $\kappa(0)$ sufficiently large, there exists a solution z_{A0} to the fixed-point equation (A.26). This solution must also be *stable* in the sense that $\Delta(z_A; z_A, \kappa(0))$ is strictly declining in z_A in a neighborhood of the solution z_{A0}. The Implicit Function Theorem then guarantees that this solution is strictly increasing in $\kappa(0)$.

To conclude, recall that entrant optimality requires $\kappa(0) \geq \kappa^E(t_0, z_{A0})$ for all $t_0 \geq 0$. With $z_{A0} \geq 1$ and $\kappa(0) > 0$ fixed, there exists a time $\bar{t} \geq 0$ such that this inequality is satisfied for $t_0 \geq \bar{t}$. With κ^E strictly decreasing in z_{A0} and z_{A0} strictly increasing in $\kappa(0)$ for any stable solution to (A.26), we also have that the entry condition can is satisfied for all $t_0 \geq 0$ provided that $\kappa(0)$ is sufficiently large. ■

The following lemma verifies that the condition (A.25) is implied by the assumptions of Proposition 4.

Lemma A.2.2. *If $\lambda \eta_B s \geq \frac{\sigma_I}{\lambda \sigma_E + \sigma_I} g_A^*$, then (A.25) holds in any equilibrium with $\ddot{K}_A(t) \leq 0$ for all $t \geq 0$ and $g_A(0) \leq g_A^*$.*

Proof of Lemma A.2.2: Differentiating the right side of (A.25) in t yields

$$\frac{\partial}{\partial t} \frac{\sigma_I \dot{K}_A(t)}{k_A(t)} = \frac{\sigma_I \ddot{K}_A(t)}{k_A(t)} - \left(\frac{\sigma_I \dot{K}_A(t)}{k_A(t)} \right)^2 \leq 0.$$

Hence the right side of (A.25) is bounded above by

$$\frac{\sigma_I K_A(0)}{k_A(0)} g_A(0) = \frac{\sigma_I}{\lambda \sigma_E + \sigma_I} g_A(0) \leq \frac{\sigma_I}{\lambda \sigma_E + \sigma_I} g_A^*.$$

■

Equilibrium Converging to *A*

The following proposition states the existence result for equilibria converging to the BGP for technology *A*, analogous to Proposition A.2.2:

Proposition A.2.3. *There exists a threshold $\kappa_A^* \geq \kappa_B^*$ such that a monotone equilibrium converging to the BGP for technology A exists if $\kappa(0) \leq \kappa_A^*$.*

Note the absence of a condition on the growth rates of the knowledge stocks K. In any monotone equilibrium converging to A, at some time T_E entrants must start by innovating for technology A. This only holds if the knowledge stock K_A always grows relative to the knowledge stock K_B. As a result, any firm that initially innovates for technology A never finds it optimal to begin innovating for technology B.

Proof of Proposition A.2.3: Many proof details are similar to those of Proposition A.2.2, so I only highlight the distinctions. As Proposition A.2.1 indicates, a monotone equilibrium converging to the BGP for technology A is fully characterized by $z_{A0} \geq 1$ and $T_E \geq 0$: All initial incumbents with relative qualities $z_A(0) \geq z_{A0}$ permanently innovate for technology A, while those with $z_A(0) < z_{A0}$ initially innovate for technology B. Similarly, all entrants at time $t_0 \geq T_E$ permanently innovate for technology A, while entrants at time $t_0 < T_E$ initially innovate for technology B. An equilibrium consists of a pair (z_{A0}, T_E) such that these decisions are optimal.

For each $z_{A0} \geq 0$, define the function $T_E(z_{A0})$ such that $T_E(z_{A0}) \in [0, \infty]$ gives the smallest time at which the entrant decisions described above are optimal. To see that such a time exists, suppose first $T_E = 0$. If entrants at time $t_0 = 0$ find it optimal to innovate initially for A, then so must every subsequent entrant given that K_A grows relative to K_B. We can then set $T_E(z_{A0}) = 0$. Otherwise, we can raise T_E until we find a time $T_E(z_{A0}) > 0$ at which entrants are indifferent to their initial innovation directions. Necessarily each entrant before $T_E(z_{A0})$ must strictly prefer to innovate initially for B, while each subsequent entrant strictly prefers to innovate for A. If such an indifference time does not exist, then we can set $T_E(z_{A0}) = \infty$, and the economy instead converges asymptotically to the BGP for technology B. With positive spillovers to entrants and (potentially) across incumbents, we must have $T_E(z_{A0})$ weakly increasing in z_{A0}, and strictly so whenever $T_E(z_{A0})$ is interior.

Given the function $T_E(z_{A0})$, we must now find an equilibrium relative quality cutoff z_{A0}. First set $z_{A0} = 1$. If, given $T_E(1)$, an initial incumbent with relative quality $z_A(0) = 1$ prefers to innovate for technology A, then we must have $T_E(1) = 0$. We have then found a "corner" equilibrium with $z_{A0} = 1$ and no innovation for technology B. Otherwise, we can raise z_{A0} until we find a value at which an initial incumbent with relative quality z_{A0} is indifferent its initial innovation direction. Necessarily each initial incumbent with $z_A(0) < z_{A0}$ strictly prefers to innovate initially for B, while each initial incumbent with $z_A(0) > z_{A0}$ strictly prefers to innovate for technology A. To see that such a value z_{A0} exists, note first that for $z_A(0)$ sufficiently high an initial incumbent strictly prefers to innovate initially for technology A regardless of other firms' innovation decisions. But we cannot immediately apply the Intermediate Value Theorem: There exists a minimum value

$\underline{z} \in [1, \infty]$ such that $T_E(z) = \infty$ for $z \geq \underline{z}$. If $\underline{z} < \infty$ and $\sigma_I > 0$, then each incumbent's net value from innovating initially for technology B jumps upward at $z_{A0} = \underline{z}$. This holds because a positive mass of future firms have altered their innovation directions, and with $\sigma_I > 0$ initial incumbents find it valuable to align with them. But this poses no issues for the argument: If the indifference cutoff z_{A0} appears before $\underline{z}$, then we have found a monotone equilibrium converging to the BGP for technology A. If the indifference cutoff z_{A0} appears after (or at) $\underline{z}$, then the Intermediate Value Theorem implies a monotone equilibrium converging to the BGP for technology B.

This argument establishes the existence of a monotone equilibrium regardless of the initial knowledge stock ratio $\kappa(0)$. But clearly $T_E(z_{A0})$ is strictly increasing in $\kappa(0)$: By the same argument as in the proof of Proposition A.2.2, entrants always have greater incentives to start innovating for technology B when its initial knowledge stock $K_B(0)$ increases relative to technology A's. Hence the minimum value $\underline{z}$ defined above is decreasing in $\kappa(0)$. Since clearly $T_E(z_{A0}) = 0$ when $\kappa(0) = 0$, there must exist a threshold κ_A^* such that the indifference cutoff z_{A0} is reached before $\underline{z}$ if and only if $\kappa(0) \leq \kappa_A^*$. This provides the threshold characterization for the existence of a monotone equilibrium converging to A stated in the Proposition.

Finally, since we have established that a monotone equilibrium always exists, we must have $\kappa_A^* \geq \kappa_B^*$. ■

Equilibrium Multiplicity

It remains to show that the economy can feature multiple monotone equilibria when $\lambda > 0$ and $\sigma_I > 0$. To see this, fix $\kappa(0) = \kappa_B^*$, and let z_{A0}^* denote the corresponding equilibrium cutoff for incumbent relative qualities. Note that we must have $z_{A0}^* > 1$ to sustain a transition to technology B.

Using the function $T_E(z_{A0})$ defined in the proof of Proposition A.2.3, observe that we must have $T_E(z_{A0}) < \infty$ for any $z_{A0} < z_{A0}^*$. Otherwise, we could reduce both $\kappa(0)$ and z_{A0}^* to maintain a monotone equilibrium that converges to the BGP for technology B, a contradiction to the definition of κ_B^*. Hence $T_E(z_{A0})$ jumps down discretely as z_{A0} is reduced from z_{A0}^*. But with $\sigma_I > 0$, this yields a jump upward in each incumbent's net value from innovating initially for technology A instead of technology B. Since the marginal incumbent with $z_A(0) = z_{A0}^*$ was before indifferent, the incumbent with $z_A(0) = z_{A0}^* - \varepsilon$ must strictly prefer technology A for $\varepsilon > 0$ sufficiently small. Hence by decreasing z_{A0} from z_{A0}^* we can find a new indifference point (or a corner solution with $z_{A0} = 1$) that yields a monotone equilibrium converging to the BGP for technology A.

This argument establishes multiplicity of equilibria at $\kappa(0) = \kappa_B^*$. But it applies almost unchanged to $\kappa(0) = \kappa_B^* + \varepsilon$ for $\varepsilon > 0$ small, because z_{A0}^* is continuous in $\kappa(0)$ for $\kappa(0) > \kappa_B^*$. They key observation is again that $\sigma_I > 0$ implies a discrete jump in incentives at $z_{A0} = \underline{z}$. Hence $\kappa_B^* < \kappa_A^*$ when $\sigma_I > 0$. ■

A.3 Equilibrium with Technology Choice: Benchmarks

In this appendix, I provide a full characterization of the economy's equilibria in two benchmark cases. In Appendix A.3.1, firms build exclusively on the aggregate knowledge stock when innovating ($\lambda = 0$), reflecting the most common way of modeling innovation in existing work on directed innovation and technological transitions. I show that this assumption leaves no role for firm heterogeneity or the concentraton of R&D to affect the economy's equilibrium. In Appendix A.3.2, I address the other extreme in which firms build exclusively on their own past advances ($\sigma_I = 0$), so that all knowledge spillovers are confined to entrants. As discussed in Section 1.4.3, in this case the economy's equilibrium depends richly on the initial distribution of incumbents and the concentration of scientists. Both cases feature a unique equilibrium, which is also monotone in the sense of Definition 3.

A.3.1 External Knowledge Accumulation: $\lambda = 0$

Equilibrium

Suppose firms build exclusively on external knowledge when innovating ($\lambda = 0 < \sigma_I$), and consider the problem of a firm with qualities $q(t_0)$ at time $t_0 \geq 0$ choosing its innovation direction at each time $t \geq t_0$. We can directly integrate the quality evolution equation (1.8) to find

$$q_\theta(t) = q_\theta(t_0) + \sigma_I \eta_\theta \int_{t_0}^{t} K_\theta(\tau) s_\theta(q(\tau), \tau) d\tau.$$

The firm's objective (1.11) can then be written

$$\frac{\bar{\pi}(q_A(t_0) + q_B(t_0))}{\rho + \delta} + \bar{\pi} \int_{t_0}^{\infty} \exp\left(-(\rho + \delta)(t - t_0)\right) \sum_{\theta \in \{A,B\}} \frac{\sigma_I K_\theta(t) \eta_\theta s_\theta(q(t), t)}{\rho + \delta} dt.$$

The first term gives the discounted value of profits given the initial qualities $q(t_0)$, and the second term incorporates the additional value generated by innovation at each time $t \geq t_0$. This objective function is linear in the allocation of scientists at each time $t \geq t_0$, so the firm's solution is simply to allocate all scientists toward the technology θ with the largest marginal productivity of research $K_\theta(t)\eta_\theta$ at each time $t \geq t_0$. Equivalently, a firm exclusively innovates for technology B at time t if and only if the knowledge ratio satisfies $\kappa(t) \geq \eta_A/\eta_B$, and otherwise exclusively innovating for technology A. Comparing to the more general analysis of Section 1.4.1, here we also obtain a corner solution for the allocation of scientists because each technology's innovation rate is linear in the mass of scientists. This case additionally implies that each firm's allocation

is identical, because $\lambda = 0$ and flow profits are linear in qualities.

This threshold characterization of firm innovation decisions carries over to the economy's equilibrium. When $\kappa(t) \geq \eta_A/\eta_B$, all innovation is directed toward technology B, so that the knowledge stock for B grows relative to that for A. The ratio $\kappa(t)$ is then strictly increasing, so that innovation is permanently directed toward B. The opposite holds when $\kappa(t) < \eta_A/\eta_B$. Thus the ratio of research productivities η_A/η_B functions as a threshold for $\kappa(0)$ that determines the equilibrium direction of technological change:

Proposition A.3.1. *The economy with $\lambda = 0$ has a unique equilibrium. All firms exclusively innovate for technology B at $t \geq 0$ if and only if*

$$\kappa(0) \geq \frac{\eta_A}{\eta_B},$$

with all innovation directed toward A otherwise. If innovation is directed toward technology B, the knowledge stock K_B immediately grows at the BGP rate $g_B^* = \sigma_I \eta_B S$, and $Q_B(t)/K_B(t)$ increases monotonically to its BGP value.

Proof of Proposition A.3.1: The threshold characterization of equilibrium research follows immediately from the discussion preceding the proposition. If research is directed toward technology B, the evolution equation (1.9) for K_B immediately implies $\dot{K}_B(t) = \sigma_I \eta_B S K_B(t)$. Integrating the evolution equation (1.15) for Q_B then yields

$$Q_B(t) = \frac{g_B + \delta N \sigma_E}{g_B + \delta} K_B(t) - \frac{(1 - N\sigma_E) g_B}{g_B + \delta} K_B(0) \exp(-\delta t).$$

Hence

$$\frac{Q_B(t)}{K_B(t)} = \frac{g_B + \delta N \sigma_E}{g_B + \delta} - \frac{(1 - N\sigma_E) g_B}{g_B + \delta} \exp(-(\delta + g_B)t)$$

By (1.19), this ratio is strictly increasing in t and limits to the BGP value from Proposition 1. ■

This result is a special case of Propositions 3 and 4. The economy's unique equilibrium features a transition to technology B if and only if the initial ratio of knowledge stocks $\kappa(0)$ is sufficiently high. But with $\eta_B > \eta_A$, the economy may transition even if technology B is initially inferior to technology A ($\kappa(0) < 1$). Along a transition, the knowledge stock K_B features no transitional dynamics, while the aggregate quality Q_B increases relative to K_B as incumbents improve relative to entrants before exit.

Efficiency

A key implication of this analysis is that the equilibrium direction of technological change is essentially myopic: At each time, all firms research the technology with the higher research productivity $\sigma_I K_\theta(t)\eta_\theta$. This holds because firms do not internalize the value of their knowledge spillovers for future innovating firms, and with $\lambda = 0$ they do not benefit directly from their past knowledge production. To understand the resulting inefficiency, consider the problem of a social planner who can choose each firm's allocation of scientists to maximize the consumer's discounted utility (1.1), but cannot otherwise modify the equilibrium. This social planner solves

$$\max_{[s_\theta(q,t)]_{\theta,q,t}} \int_0^\infty \exp(-\rho t)\, C(t) dt, \tag{A.27}$$

subject to the resource constraint $s_A(q,t) + s_B(q,t) \leq s$ and with all remaining quantities determined in equilibrium. The following proposition solves this problem for the case with $\lambda = 0$, demonstrating that the equilibrium transitions to technology B insufficiently often:

Proposition A.3.2. *With $\lambda = 0$, the social planner allocates all scientists to technology B at $t \geq 0$ if and only if*

$$\kappa(0) \geq \frac{\eta_A}{\eta_B} J,$$

with all innovation directed toward A otherwise. The adjustment factor $J \in (0,1)$ is independent of initial conditions.

Proof of Proposition A.3.2: Since incumbents' qualities do not affect their research productivities, it is optimal for the social planner to permanently allocate all scientists to one technology. If the social planner chooses technology A, then Q_A continues to grow at the BGP rate $g_A^* = \sigma_I \eta_A S$, while Q_B remains constant at its initial value $Q_B(0) = N\sigma_E K_B(0)$. With flow consumption equal to $\bar{C}(Q_A + Q_B)$ for a constant $\bar{C} > 0$, this yields social value

$$\begin{aligned}\frac{U^A}{\bar{C}} &= \frac{Q_A(0)}{\rho - g_A^*} + \frac{Q_B(0)}{\rho} \\ &= \frac{1}{\rho - g_A^*} \frac{g_A^* + \delta N \sigma_E}{g_A^* + \delta} K_A(0) + \frac{1}{\rho} N \sigma_E K_B(0).\end{aligned}$$

If instead the social planner permanently allocates all scientists to technology B, then Q_A and Q_B evolve according to

$$\dot{Q}_A(t) = \delta\left[N\sigma_E K_A(0) - Q_A(t)\right],$$
$$\dot{Q}_B(t) = g_B^* K_B(0)\exp\left(g_B^* t\right) + \delta\left[N\sigma_E K_B(0)\exp\left(g_B^* t\right) - Q_B(t)\right].$$

In the second equation I make use of the relation $K_B(t) = K_B(0)\exp\left(g_B^* t\right)$. Integrating these equations yields

$$Q_A(t) = N\sigma_E K_A(0) + \frac{(1-N\sigma_E)\,g_A^*}{g_A^*+\delta}K_A(0)\exp\left(-\delta t\right),$$
$$Q_B(t) = \frac{g_B^*+\delta N\sigma_E}{g_B^*+\delta}K_B(0)\exp\left(g_B^* t\right) - \frac{(1-N\sigma_E)\,g_B^*}{g_B^*+\delta}K_B(0)\exp\left(-\delta t\right).$$

The social welfare from researching technology B permanently is then

$$\begin{aligned}\frac{U^B}{\bar{C}} &= \int_0^\infty \exp\left(-\rho t\right)\left[Q_A(t)+Q_B(t)\right]dt \\ &= \left[\frac{1}{\rho}N\sigma_E + \frac{1}{\rho+\delta}\frac{(1-N\sigma_E)\,g_A^*}{g_A^*+\delta}\right]K_A(0) \\ &\quad + \left[\frac{1}{\rho-g_B^*}\frac{g_B^*+\delta N\sigma_E}{g_B^*+\delta} - \frac{1}{\rho+\delta}\frac{(1-N\sigma_E)\,g_B^*}{g_B^*+\delta}\right]K_B(0).\end{aligned}$$

Researching technology B is socially optimal if and only if $U^B \geq U^A$, or equivalently

$$j(g_B^*)g_B^*K_B(0) \geq j(g_A^*)g_A^*K_A(0),$$

where I define

$$j(g) \equiv \left[\frac{1}{\rho-g}\frac{g+\delta N\sigma_E}{g+\delta} - \frac{1}{\rho+\delta}\frac{(1-N\sigma_E)\,g}{g+\delta} - \frac{1}{\rho}N\sigma_E\right]\frac{1}{g}.$$

This function can equivalently be written

$$j(g) = \frac{\frac{g+\delta N\sigma_E}{\rho-g} + \frac{\delta}{\rho+\delta}(1-N\sigma_E)}{g+\delta}\frac{1}{\rho}.$$

Direct calculation implies that this function is strictly increasing in g:

$$j'(g) \propto (g+\delta)\left(1+\frac{g+\delta N\sigma_E}{\rho-g}\right) - \left[g+\delta N\sigma_E + (\rho-g)\frac{\delta}{\rho+\delta}(1-N\sigma_E)\right].$$

This value is positive if and only if

$$\delta(1-N\sigma_E)+(g+\delta)\frac{g+\delta N\sigma_E}{\rho-g}>-(\rho-g)\frac{\delta}{\rho+\delta}(1-N\sigma_E).$$

This necessarily holds, because the left side is positive and the right side is negative because of the assumed upper bound on entrant spillovers (1.19) and the assumed lower bound on the discount rate (1.20). This analysis demonstrates that the social planner exclusively and permanently researches technology B if and only if

$$\kappa(0)\geq\frac{\eta_A}{\eta_B}J,$$

where the adjustment factor $J\equiv j(g_A^*)/j(g_B^*)$ is positive but strictly smaller than one. ■

For intuition, consider the limit $\sigma_E N\uparrow 1$, so that the aggregate qualities Q are unaffected by entry and exit and hence display no transitional dynamics. The factor J then limits to

$$\lim_{\sigma_E N\uparrow 1}J=\frac{\rho-g_B^*}{\rho-g_A^*}.$$

This value is strictly smaller than one precisely because the *long-run* growth rate for technology A is smaller than that for technology B. The social planner internalizes the value of knowledge spillovers for future innovation, so her transition threshold takes into account not only the $t=0$ research productivities for each technology, but also the implied long-run growth rates g_A^* and g_B^*.

Finally, note the distinction between Proposition A.3.2 and Proposition 6, which is the corresponding efficiency result for the benchmark case with only internal knowledge ($\sigma_I=0$). With only external knowledge, the social planner prefers to transition to technology B more often than in equilibrium; but conditional on the direction of innovation, the equilibrium is efficient (Proposition A.3.2). With only internal knowledge, the social planner does not necessary prefer to transition to technology B more often; but conditional on the direction of innovation, the equilibrium converges too slowly to the limiting BGP.

A.3.2 Internal Knowledge Accumulation: $\sigma_I=0$

Equilibrium

Now consider the other extreme case in which firms build exclusively on internal knowledge when innovating ($\sigma_I=0<\lambda$), summarized in Section 1.4.3. With $\sigma_I=0$, we can again

integrate the quality evolution equation (1.8) to find

$$q_\theta(t) = q_\theta(t_0)\exp\left(\lambda\eta_\theta \int_{t_0}^{\tau} s_\theta(q(\tau),\tau)d\tau\right).$$

In contrast to the case with $\lambda = 0$, $q_\theta(t)$ is not additively separable between the initial quality $q_\theta(t_0)$ and the allocation of scientists. Since firms build on their past advances, an increase in the initial quality $q_\theta(t_0)$ raises the firm's research productivity for θ at each time $\tau \geq t_0$. With profits linear in quality, this mechanism leads to path dependence whereby a firm is more likely to research a technology that it has researched previously. The following lemma precisely characterizes equilibrium innovation decisions:

Lemma A.3.1. *With $\sigma_I = 0$, a firm with qualities q exclusively innovates for technology B if*

$$\frac{\eta_B}{\rho+\delta-\lambda\eta_B s}q_B \geq \frac{\eta_A}{\rho+\delta-\lambda\eta_A s}q_A,$$

with all innovation directed toward A otherwise.

Proof of Lemma A.3.1: It is clearly optimal for the firm to innovate permanently for one technology. If the firm innovates permanently for technology A, its value is

$$\frac{V^A}{\bar{\pi}} = \frac{q_A}{\rho+\delta-\lambda\eta_A s} + \frac{q_B}{\rho+\delta}$$

If the instead innovates permanently for technology B, its value is

$$\frac{V^B}{\bar{\pi}} = \frac{q_A}{\rho+\delta} + \frac{q_B}{\rho+\delta-\lambda\eta_B s}.$$

Innovating for technology B is optimal if and only if $V^B \geq V^A$. Substituting the expressions above into this inequality yields the inequality stated in the Lemma. This inequality is self-reinforcing, so an incumbent's innovation direction is perfectly persistent. ■

This result clarifies the two forces that determine the direction of a firm's innovation. First, as discussed above, a firm has a greater propensity to innovate for the technology for which it has a higher quality. Second, firms have a greater propensity to innovate for technology B because of its higher basic research productivity $\eta_B > \eta_A$. This force is stronger with $\lambda > 0$ because the firm's problem is genuinely foward-looking: Innovating for technology θ raises the quality q_θ, making research for θ more productive in the future. The firm internalizes this dynamic effect because its innovation decision has a non-negligible impact on the change in quality q_θ (in

contrast to the knowledge stock K_θ).

We can apply this result to derive a convenient characterization of entrant and incumbent innovation decisions. Substituting the entrant qualities $q_\theta^E(t) = \sigma_E K_\theta(t)$ into the inequality stated in the Lemma, we observe that an entrant at $t \geq 0$ permanently innovates for technology B if and only if the knowledge stock ratio $\kappa(t)$ is larger than the *entry threshold*

$$\kappa^E \equiv \frac{\eta_A}{\rho + \delta - \lambda \eta_A s} \left(\frac{\eta_B}{\rho + \delta - \lambda \eta_B s} \right)^{-1} \in (0,1).$$

An incumbent at $t = 0$ with initial qualities $q_A(0) = z_A(0) q_A^E(0)$ and $q_B(0) = q_B^E(0)$ permanently innovates for technology B if and only if its relative quality $z_A(0)$ is above the cutoff

$$z_{A0}^* \equiv \frac{\kappa(0)}{\kappa^E}$$

Note that path dependence implies that $t = 0$ incumbents always have a lower propensity to innovate for B than entrants at $t = 0$.

Using this characterization of innovation decisions, the following lemma derives a simple dynamical system for the knowledge stocks K, analogous to that of Proposition 3:

Lemma A.3.2. *With $\sigma_I = 0$, the knowledge stocks K evolve according to*

$$\ddot{K}_A(t) = -(\delta - \lambda \eta_A s) \dot{K}_A(t) + \delta N \lambda \eta_A s \sigma_E K_A(t) \mathbb{1}\left[\kappa(t) < \kappa^E\right], \tag{A.28}$$

$$\ddot{K}_B(t) = -(\delta - \lambda \eta_B s) \dot{K}_B(t) + \delta N \lambda \eta_B s \sigma_E K_B(t) \mathbb{1}\left[\kappa(t) \geq \kappa^E\right]. \tag{A.29}$$

The corresponding initial conditions are

$$\begin{aligned} & K_A(0), \quad K_B(0), \\ & \frac{\dot{K}_A(0)}{K_A(0)} = \lambda \eta_A s \sigma_E \int_{z_{A0}^*}^{\infty} z_A \, dH_A(z_A), \\ & \frac{\dot{K}_B(0)}{K_B(0)} = \lambda \eta_B s \sigma_E H_A(z_{A0}^*). \end{aligned} \tag{A.30}$$

Proof of Lemma A.3.2: For expositional purposes only, suppose $\kappa(0) > \kappa^E$, and fix t small enough that $\kappa(t') > \kappa^E$ for $t' \in [0,t]$. Consider technology A. Using Lemma A.3.1 and integrating the evolution equation (1.8) for q_A, we have that for any $t = 0$ incumbent with

$z_A(0) > z^*_{A0}$,

$$q_A(t) = z_A(0)\sigma_E K_A(0)\exp(\lambda\eta_A s t).$$

The density of these incumbents at $t = 0$ is $h_A(z_A(0))$. But with exit at rate $\delta > 0$, the density at time $t > 0$ falls to $\exp(-\delta t)h_A(z_A(0))$. Hence the evolution equation (1.9) for K_A can be written

$$\dot{K}_A(t) = \lambda\eta_A s\sigma_E K_A(0)\left(\int_{z_A > z^*_{A0}} z_A dH_A(z_A)\right)\exp(-(\delta - \lambda\eta_A s)t).$$

Differentiating in t yields (A.28).

Turning to technology B, the same argument as above implies that for any $t = 0$ incumbent with $z_A(0) \leq z^*_{A0}$,

$$q_B(t) = \sigma_E K_B(0)\exp(\lambda\eta_B s t).$$

By time t, the total mass of $t = 0$ incumbents with $z_A(0) \leq z^*_{A0}$ is $\exp(-\delta t)H_A(\hat{z}_A)$. Similarly, for any firm who entered at time $\tau \in [0, t]$,

$$q_B(t) = \sigma_E K_B(\tau)\exp(\lambda\eta_B s(t - \tau)).$$

By time $t \geq \tau$, the density of these firms is $\exp(-\delta(t - \tau))\,\delta N$. Hence the evolution equation (1.9) can be written

$$\begin{aligned}\dot{K}_B(t) = {} & \lambda\eta_B s\sigma_E K_B(0)H_A(z^*_{A0})\exp(-(\delta - \lambda\eta_B s)t) \\ & + \lambda\eta_B s\sigma_E \delta N\int_0^t K_B(\tau)\exp(-(\delta - \lambda\eta_B s)(t - \tau))\,d\tau.\end{aligned}$$

Differentiating in t yields (A.29).

The same arguments yield the evolution equations (A.28, A.29) regardless of $\kappa(0)$ or t. The initial conditions (A.30) follow directly from the evolution equation (1.9) for the knowledge stocks K, given the assumption that the economy is following the BGP for technology A at $t = 0$. ■

The benchmark case with $\sigma_I = 0$ is convenient precisely because the system (A.28, A.29) is autonomous and can be integrated in closed form. The following lemma demonstrates this for the case when the economy transitions to the BGP for technology B:

Lemma A.3.3. *Suppose entrants innovate for technology B at each time, $\kappa(t) \geq \kappa^E$ for $t \geq 0$.*

1. *The solution to the system (A.28, A.29) is*

$$\frac{K_A(t)}{K_A(0)} = c_{A1} - c_{A2}\exp\left(-\left(\delta - \lambda\eta_A s\right)t\right),$$
$$\frac{K_B(t)}{K_B(0)} = c_{B1}\exp\left(-\left(g_B^* + \delta - \lambda\eta_B s\right)t\right) + c_{B2}\exp\left(g_B^* t\right),$$

   where $c_{A1}, c_{A2}, c_{B1}, c_{B2} > 0$ satisfy the initial conditions (A.30).

2. *The knowledge ratio $\kappa(t) = K_B(t)/K_A(t)$ is strictly increasing and strictly convex in the initial condition $\kappa(0)$ for $t > 0$, and there exists a time $\hat{t} \geq 0$ such that $(t - \hat{t})\dot{\kappa}(t) > 0$.*

Proof of Lemma A.3.3: Integrating the evolution equation (A.28) for $\dot{K}_A$ yields

$$\frac{K_A(t)}{K_A(0)} = 1 + \lambda\eta_A s\sigma_E\left(\int_{z_{A0}^*}^{\infty} z_A dH_A(z_A)\right)\frac{1 - \exp\left(-\left(\delta - \lambda\eta_A s\right)t\right)}{\delta - \lambda\eta_A s}.$$

Hence the integration constants from the Lemma are

$$c_{A1} = 1 + c_{A2},$$
$$c_{A2} = \lambda\eta_A s\sigma_E\left(\int_{z_{A0}^*}^{\infty} z_A dH_A(z_A)\right)\frac{1}{\delta - \lambda\eta_A s}.$$

Integrating the evolution equation (A.29) for $\dot{K}_B$ and making use of the formula for g_B^* given in Proposition 1 yields the expression for $K_B(t)/K_B(t)$ given in the Lemma statement. The integration constants jointly satisfy the initial conditions

$$1 = c_{B1} + c_{B2},$$
$$\lambda\eta_B s\sigma_E H_A(z_{A0}^*) = -\left(g_B^* + \delta - \lambda\eta_B s\right)c_{B1} + g_B^* c_{B2}.$$

The solution to this system is

$$c_{B1} = \frac{g_B^* - \lambda\eta_B s\sigma_E H_A(z_{A0}^*)}{2g_B^* + \delta - \lambda\eta_B s},$$
$$c_{B2} = \frac{g_B^* + \delta - \lambda\eta_B s\left(1 - \sigma_E H_A(z_{A0}^*)\right)}{2g_B^* + \delta - \lambda\eta_B s}.$$

Clearly $c_{B2} > 0$ given the assumed lower bound (1.14) on the exit rate δ. We also have

$$\begin{aligned} g_B^* - \lambda\eta_B s\sigma_E H_A(z_{A0}^*) &\geq g_B^* - \lambda\eta_B S\sigma_E \\ &\geq g_B^* - \lambda\eta_B s \\ &= -\frac{\delta - \lambda\eta_B s}{2} - \lambda\eta_B S\sigma_E + \sqrt{\left(\frac{\delta - \lambda\eta_B s}{2}\right)^2 + \lambda\sigma_E\delta\eta_B S}. \end{aligned}$$

Direct calculation implies that the right side is positive given the assumed upper bound on spillovers to entrants (1.19), so $c_{B1} > 0$.

To prove the properties of κ stated in the Lemma, note that the above analysis shows that $K_A(t)/K_A(0)$ depends on the initial knowledge stocks $K(0)$ only through z_{A0}^*. This cutoff is strictly increasing in $\kappa(0)$ while $K_A(t)/K_A(0)$ is strictly decreasing in $\hat{z}_A$ for $t > 0$, so $K_A(t)/K_A(0)$ is strictly decreasing in $\kappa(0)$ for $t > 0$. A similar argument implies that $K_B(t)/K_B(0)$ is strictly increasing in $\kappa(0)$ for $t > 0$. Hence

$$\kappa(t) = \kappa(0)\frac{K_B(t)/K_B(0)}{K_A(t)/K_A(0)}$$

depends on $K(0)$ only through $\kappa(0)$, and it is strictly increasing and strictly convex in $\kappa(0)$ for $t > 0$. Finally, note that the growth rate of κ satisfies

$$\frac{\dot{\kappa}(t)}{\kappa(t)} = \frac{\dot{K}_B(t)}{K_B(t)} - \frac{\dot{K}_A(t)}{K_A(t)}.$$

We can directly calculate

$$\begin{aligned} \frac{\dot{K}_A(t)}{K_A(t)} &= \frac{(\delta - \lambda\eta_A s)\, c_{A2}}{c_{A1}\exp\left((\delta - \lambda\eta_A s)\, t\right) - c_{A2}}, \\ \frac{\dot{K}_B(t)}{K_B(t)} &= g_B^* \frac{-\left(1 + \frac{\delta - \lambda\eta_B s}{g_B^*}\right) c_{B1} + c_{B2}\exp\left(\left(2g_B^* + \delta - \lambda\eta_B s\right) t\right)}{c_{B1} + c_{B2}\exp\left(\left(2g_B^* + \delta - \lambda\eta_B s\right) t\right)}. \end{aligned}$$

Since all integration coefficients are positive, we immediately have that the growth rate of $K_A(t)$ is declining over time, while the growth rate of $K_B(t)$ is increasing over time. The previous equation then implies that $\dot{\kappa}(t)$ is single-crossing from below, so there exists a time $\hat{t} \geq 0$ such that $\dot{\kappa}(t) < 0$ for $t < \hat{t}$ and $\dot{\kappa}(t) > 0$ for $t > \hat{t}$. ■

The second part of the lemma leverages the solution to the system (A.28, A.29) to prove two properties of the ratio $\kappa(t)$: It is strictly increasing in its initial value $\kappa(0)$, and it is generally "U-shaped" over time. These properties are useful because the solution in the lemma only describes

the dynamics of the knowledge stocks $K(0)$ while entrants continue to research technology B, $\kappa(t) \geq \kappa^E$. If this condition is ever violated, the economy fails to transition to technology B and instead converges back to the BGP for technology A. The properties of $\kappa(t)$ described in Lemma A.3.3 ensure that a transition takes place exactly when $\kappa(0)$ is sufficiently large.

Proposition A.3.3. *With $\sigma_I = 0$, the economy has a unique equilibrium. There exists a threshold $\kappa^* > \kappa^E$ such that if $\kappa(0) \geq \kappa^*$, all firms innovate for B as $t \to \infty$, and the economy converges to the BGP for B. Otherwise, all firms innovate for A as $t \to \infty$, and the economy converges to the BGP for A. The economy displays transitional dynamics when $\kappa(0) > \kappa^E$.*

Proof of Proposition A.3.3: Clearly if $\kappa(0) \leq \kappa^E$, all innovation is initially directed toward technology A, so that $\kappa(t) < \kappa^E$ for all $t > 0$. Hence all incumbents innovate for technology A, so that the economy continues along the BGP for A.

If instead $\kappa(0) > \kappa^E$ but $\kappa(t) < \kappa^E$ for some time $t > 0$, then the economy again converges back to the BGP for A. To see this, let $\underline{t} = \inf\{t : \kappa(t) < \kappa^E\}$. For t in a neighborhood to the right of $\underline{t}$, Lemma A.3.2 implies that the knowledge stocks K evolve according to

$$\ddot{K}_A(t) = -(\delta - \lambda\eta_A s)\dot{K}_A(t) + \lambda\eta_A s\delta N\sigma_E K_A(t),$$
$$\ddot{K}_B(t) = -(\delta - \lambda\eta_B s)\dot{K}_B(t).$$

As in Lemma A.3.3, the solution implies that the growth rate $\dot{K}_A(t)/K_A(t)$ is increasing for t in a neighborhood to the right of $\underline{t}$, while the growth rate $\dot{K}_B(t)/K_B(t)$ is decreasing. Since κ must be strictly decreasing for t in a neighborhood to the left of $\underline{t}$, this implies $\dot{K}_A(t)/K_A(t) > \dot{K}_B(t)/K_B(t)$ in a neighborhood to the right of $\underline{t}$. But then the inequality $\kappa(t) < \kappa^E$ is self-reinforcing, and the economy converges back to the BGP for technology A.

The argument above implies that the economy transitions to the BGP for technology B asymptotically if and only if the trajectory of $\kappa(t)$ implied by (A.28, A.29) with initial conditions (A.30) satisfies $\kappa(t) \geq \kappa^E$ for all $t \geq 0$. Lemma A.3.3 implies that $\kappa(t)$ is asymptotically increasing in t and strictly increasing in $\kappa(0)$, so there must exist a threshold κ^* such that $\kappa(t) \geq \kappa^E$ for all $t \geq 0$ if and only if $\kappa(0) \geq \kappa^*$. Note that this threshold must satisfy $\kappa^* > \kappa^E$, because otherwise $\kappa(t)$ initially falls below κ^E.

The transitional dynamics of the knowledge stocks K and the aggregate qualities Q are described by the system (1.15, A.28, A.29). This system is conveniently block diagonal, so that K can be recovered without integrating the evolution equations for Q. ■

Figure A.3.1 displays the trajectories of the knowledge stocks K, the growth rates $\dot{K}_\theta/K_\theta$, the knowledge ratio κ, and the ratios Q_θ/K_θ for an example transition. Note that the trajec-

Figure A.3.1: Example Transition with $\sigma_I = 0$

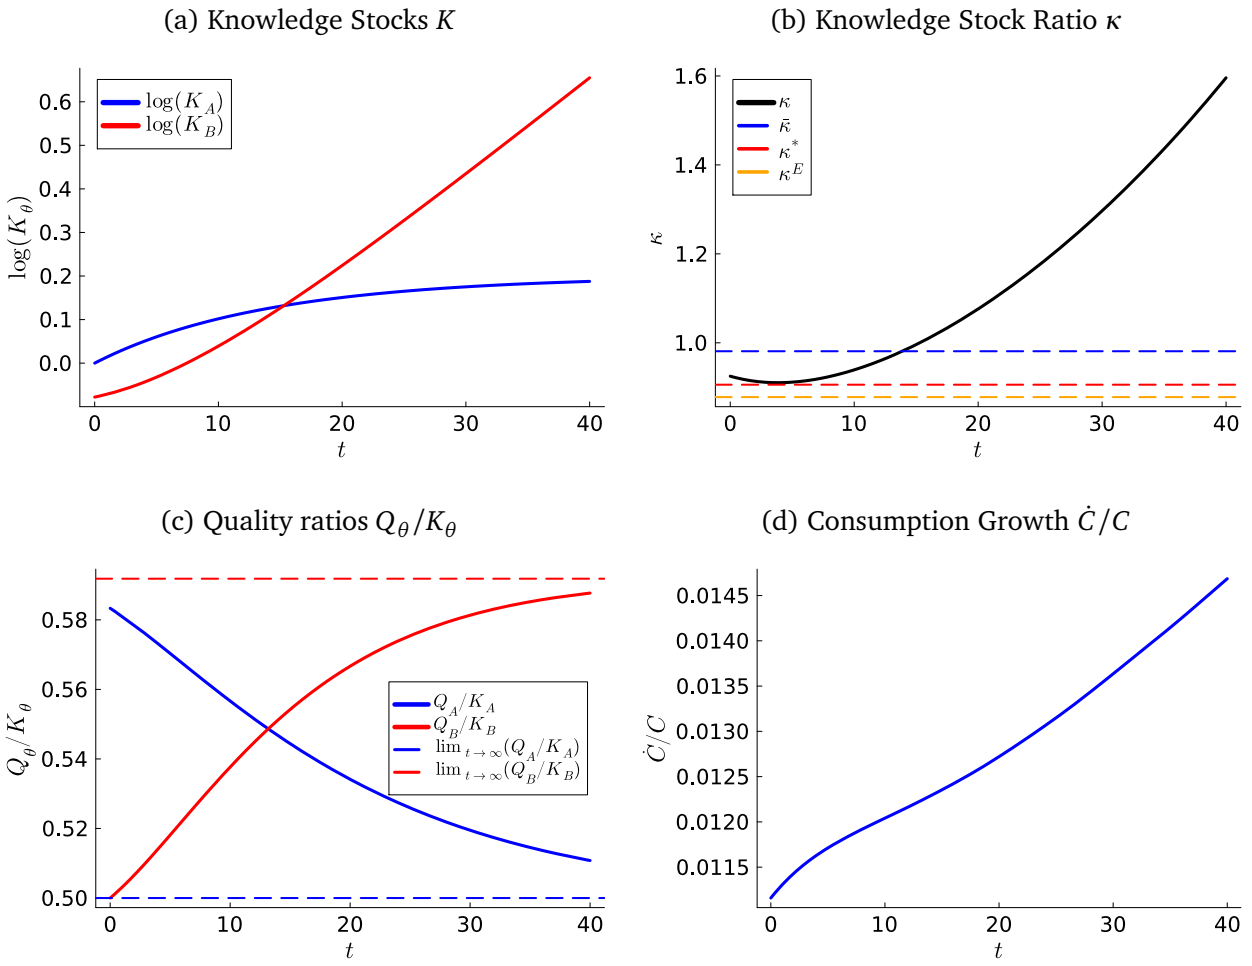

Notes: To calibrate, I set $\lambda = 1$, $\sigma_E = 0.5$, and $\sigma_I = 0$. I also set $\delta = 0.1$, $S = 1$, and $s = 1$, and I choose η_A and η_B to deliver BGP growth rates $g_A^* = 0.02$ and $g_B^* = 0.0225$. Finally, I set $\rho = 0.075$, and I specify the initial conditions $K_A(0) = 1$ and $K_B(0) = 0.925$ so that $\kappa(0) > \kappa^* \approx 0.91$.

tory of the knowledge stock ratio $\kappa(t)$ is "U-shaped," initially decreasing before increasing asymptotically.

The threshold κ^* determines the economy's propensity to transition in equilibrium, and it depends richly on model primitives. To gain intuition for the key forces, note that Lemma A.3.3 implies that a simple sufficient condition for the economy to transition is for the initial growth rate of K_B to dominate the initial growth rate of K_A. Given the initial conditions (A.30), this holds if and only if

$$\eta_B H_A\left(z_{A0}^*\right) \geq \eta_A \int_{z_{A0}^*}^{\infty} z_A dH_A(z_A). \tag{A.31}$$

This inequality depends on $\kappa(0)$ only through the relative quality cutoff $\hat{z}_A$, which is strictly increasing in $\kappa(0)$. As $\kappa(0)$ rises, the left side increases as a larger mass of $t=0$ incumbents transition to B, raising the initial growth rate of K_B. The right side instead decreases as fewer $t=0$ incumbents innovate for A, lowering the initial growth rate of K_A. There exists a unique value $\bar{\kappa}$ at which the inequality (A.31) binds:

$$1 = \frac{\eta_A}{\eta_B} \frac{1}{1-\xi_A^*} \frac{\frac{\bar{\kappa}}{\kappa^E}}{\left(\frac{\bar{\kappa}}{\kappa^E}\right)^{1/\xi_A^*} - 1}. \tag{A.32}$$

Here I have used the expression for z_{A0}^* from Lemma A.3.1 and the BGP relative quality distribution from Proposition 1, which is a standard Pareto distribution when $\sigma_I = 0$.

The right side of equation (A.32) is strictly increasing in ξ_A^* and κ^E and strictly decreasing in $\bar{\kappa}$, and it delivers a simple but powerful intuition about the drivers of a technological transition: Any change that thickens the tail of the old technology's firm-quality distribution slows the transition, because it raises both the relative mass of incumbents who choose not to transition and their initial innovation rates. Both effects increase their collective influence over the aggregate direction of innovation, which may be decisive if it induces entrants to switch back to innovating for technology A. However, any change that raises incentives for new firms to innovate for the new technology instead accelerates the transition by raising the relative mass of incumbents who choose to transition. The tail parameter ξ_A^* and the entry threshold κ^E respectively capture these two forces, but they depend on many of the same model primitives. The following proposition provides explicit comparative statics for $\bar{\kappa}$:

Proposition A.3.4. *With $\sigma_I = 0$, there exists a threshold $\bar{\kappa} \geq \kappa^*$ such that*

$$\frac{\dot{K}_B(0)}{K_B(0)} \geq \frac{\dot{K}_A(0)}{K_A(0)} \quad \Longleftrightarrow \quad \kappa(0) \geq \bar{\kappa}.$$

The threshold $\bar{\kappa}$ is strictly decreasing in σ_E, S, and η_B, and it is strictly increasing in ρ and η_A. For each variable $v \in \{\delta, \lambda, s\}$, there exists a discount rate $\rho^v \geq 0$ that depends on model primitives such that

1. *$\bar{\kappa}$ is strictly increasing (decreasing) in δ locally if and only if ρ is smaller (larger) than ρ^δ;*
2. *$\bar{\kappa}$ is strictly increasing (decreasing) in λ locally if and only if ρ is larger (smaller) than ρ^λ;*
3. *$\bar{\kappa}$ is strictly increasing (decreasing) in s locally if and only if ρ is larger (smaller) than ρ^s.*

Proof of Proposition A.3.4: The existence of the threshold $\bar{\kappa} \geq \kappa^*$ follows immediately from the discussion preceding the Proposition.

Throughout the remainder of the proof, let RHS denote the right side of (A.32), and let $\nu \equiv \bar{\kappa}/\kappa^E$. It is immediate that RHS is strictly decreasing in ν and strictly increasing in ξ_A^*:

$$\frac{\partial \text{RHS}}{\partial \nu} = -\frac{\eta_A}{\eta_B}\frac{1}{1-\xi_A^*}\frac{\left(\frac{1}{\xi_A^*}-1\right)\nu^{1/\xi_A^*}+1}{\left(\nu^{1/\xi_A^*}-1\right)^2},$$
$$\frac{\partial \text{RHS}}{\partial \xi_A^*} = \text{RHS}\left[\frac{1}{1-\xi_A^*}+\frac{1}{(\xi_A^*)^2}\frac{\nu^{1/\xi_A^*}\log(\nu)}{\nu^{1/\xi_A^*}-1}\right]$$

Several comparative statics follow immediately: $\bar{\kappa}$ is strictly decreasing in σ_E and S because these parameters only reduce ξ_A^*. Similarly, $\bar{\kappa}$ is strictly increasing in ρ because an increase in ρ only increases κ^E. Now RHS is directly decreasing in η_B and indirectly decreasing in η_B through κ^E, so $\bar{\kappa}$ is strictly decreasing in η_B. A symmetric argument shows that $\bar{\kappa}$ is strictly increasing in η_A.

For the comparative static with respect to $\nu \in \{\delta, \lambda, s\}$, we can differentiate to find

$$\frac{d\text{RHS}}{d\nu} = \frac{\partial \text{RHS}}{\partial \xi_A^*}\frac{\partial \xi_A^*}{\partial \nu} + \frac{\partial \text{RHS}}{\partial \nu}\nu\kappa^E\frac{\partial(\kappa^E)^{-1}}{\partial \nu}. \tag{A.33}$$

Now (A.32) determines ν independently of κ^E, and since RHS depends on ρ only through κ^E, this implies that ν is invariant to ρ. The first summand in (A.33) depends on κ^E only through ν, and similarly for the partial derivative $\partial\text{RHS}/\partial\nu$, so these terms are invariant to ρ. With Lemma A.3.1, we can directly calculate

$$\frac{\partial(\kappa^E)^{-1}}{\partial\delta} = -(\kappa^E)^{-1}\left[\frac{1}{\rho+\delta-\lambda\eta_B s}-\frac{1}{\rho+\delta-\lambda\eta_A s}\right],$$
$$\frac{\partial(\kappa^E)^{-1}}{\partial\lambda} = (\kappa^E)^{-1}\left[\frac{\eta_B}{\rho+\delta-\lambda\eta_B s}-\frac{\eta_A}{\rho+\delta-\lambda\eta_A s}\right]s,$$
$$\frac{\partial(\kappa^E)^{-1}}{\partial s} = (\kappa^E)^{-1}\left[\frac{\eta_B}{\rho+\delta-\lambda\eta_B s}-\frac{\eta_A}{\rho+\delta-\lambda\eta_A s}\right]\lambda.$$

With $\sigma_I = 0$, Corollary 1 yields

$$\frac{\partial\xi_A^*}{\partial\delta} < 0 \quad \text{and} \quad \frac{\partial\xi_A^*}{\partial\lambda}, \frac{\partial\xi_A^*}{\partial s} > 0.$$

Hence (A.33) implies that $d\text{RHS}/d\delta$ is strictly decreasing in ρ, becoming negative in the limit $\rho \to \infty$. There must then exist a value $\rho^\delta \geq 0$ such that $\bar{\kappa}$ is strictly increasing in δ locally if and only if $\rho < \rho^\delta$, while the opposite holds for $\rho > \rho^\delta$. The same argument implies the existence of the values $\rho^\lambda, \rho^s \geq 0$ and the corresponding comparative statics with respect to λ

Figure A.3.2: Comparative Statics for Thresholds $\bar{\kappa}$ and κ^*

(a) Sufficiency Threshold $\bar{\kappa}$ (b) Equilibrium Threshold κ^*

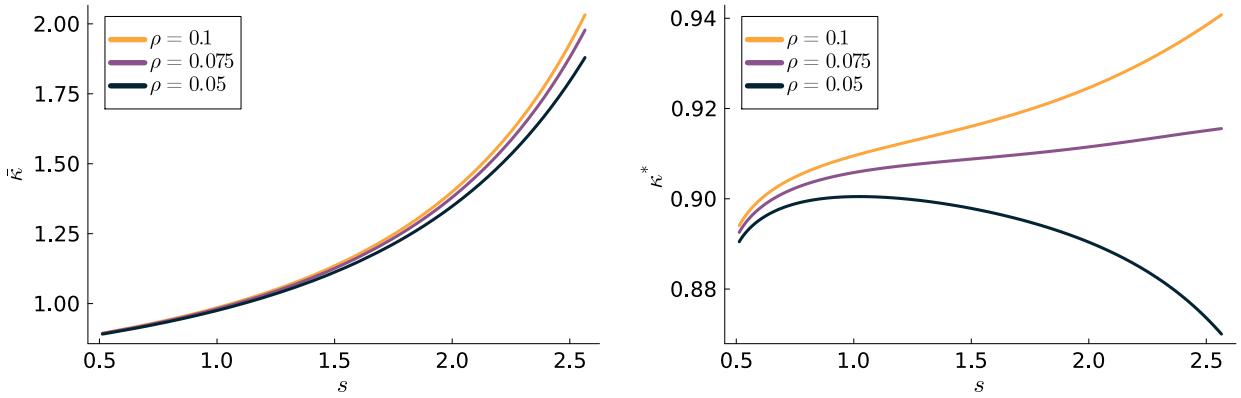

Notes: I vary the discount rate ρ around its baseline value $\rho = 0.075$. The remaining parameters are exactly as in Figure A.3.1.

and δ stated in the Proposition. ■

Comparative statics for the thresholds $\bar{\kappa}$ and κ^* with respect to s are illustrated in Figure A.3.2.

Efficiency

In the case with $\sigma_I = 0$, the solution to the planner's problem (A.27) is complex: The planner generally chooses different innovation directions for different firms and may reverse these directions over time. To develop intuition about equilibrium inefficiencies, I consider the simpler problem in which the planner can choose the initial relative quality cutoff $z_{A0} \geq 1$ describing initial incumbents' innovation decisions. The planner then solves

$$\max_{z_{A0} \geq 1} \int_0^\infty \exp(-\rho t) C(t) dt, \tag{A.34}$$

with entrant innovation decisions and all remaining quantities determined in equilibrium. The following proposition characterizes properties of the solution $\hat{z}_{A0}$:

Proposition A.3.5. *A solution $\hat{z}_{A0}$ to the social planner's problem (A.34) exists and depends on $K(0)$ only through $\kappa(0)$. There exists a threshold $\hat{\kappa}$ such that*

1. *the solution $\hat{z}_{A0}$ yields a transition to technology B if and only if $\kappa(0) \geq \hat{\kappa}$;*
2. *$\hat{z}_{A0} > z^*_{A0}$ if $\kappa(0) \geq \hat{\kappa}$; and*

3. $\hat{z}_{A0} \leq z^*_{A0}$ *if* $\kappa(0) < \hat{\kappa}$, *with equality only if* $z^*_{A0} = 1$.

Proof of Proposition A.3.5: Let U denote the $t = 0$ discounted value of future consumption. Lemma 1 implies that consumption at each time is proportional to output, which is in turn proportional to total flow profits earned by firms. Hence the discounted value of future consumption is proportional to the discounted value of future profits, or equivalently the discounted value of all firms:

$$\begin{aligned}\frac{\bar{\pi}}{\bar{C}}U = &\int_{z_{A0}}^{\infty}\left[\frac{1}{\rho+\delta-\lambda\eta_A s}z_A\sigma_E K_A(0)+\frac{1}{\rho+\delta}\sigma_E K_B(0)\right]dH_A^*(z_{A0})\\ &+\int_1^{z_{A0}}\left[\frac{1}{\rho+\delta}z_A\sigma_E K_A(0)+\frac{1}{\rho+\delta-\lambda\eta_B s}\sigma_E K_B(0)\right]dH_A^*(z_{A0})\\ &+\delta N\int_0^{\infty}\exp(-\rho t)\left[\frac{1}{\rho+\delta}\sigma_E K_A(t)+\frac{1}{\rho+\delta-\lambda\eta_B s}\sigma_E K_B(t)\right]\mathbb{1}\left[\kappa(t)\geq\kappa^E\right]dt\\ &+\delta N\int_0^{\infty}\exp(-\rho t)\left[\frac{1}{\rho+\delta-\lambda\eta_A s}\sigma_E K_A(t)+\frac{1}{\rho+\delta}\sigma_E K_B(t)\right]\mathbb{1}\left[\kappa(t)<\kappa^E\right]dt.\end{aligned}$$

This function is continuous in z_{A0} except at the critical value that separates convergence to the BGP for technology A from convergence to the BGP for technology B. To simplify, I assume that the social planner can choose the asymptotic direction of innovation at this value, so that the planner's problem (A.34) is guaranteed to have a solution for $z_{A0} \in [1, \infty]$. Clearly higher values of $K_B(0)$ raise the value of innovating for B relative to A, so the existence of the threshold $\hat{\kappa}$ is immediate given the linearity of the dynamical system (A.28, A.29) in K and the linearity of the initial conditions (A.30) in $K(0)$.

Suppose $\kappa(0) \geq \hat{\kappa}$, so that the social planner's solution $\hat{z}_{A0}$ is such that the economy converges to the BGP for B as $t \uparrow \infty$. If $\kappa(0) < \kappa^*$, so that in equilibrium the economy converges back to the BGP for A, then we immediately have $\hat{z}_{A0} > z^*_{A0}$. Suppose then that $\kappa(0) \geq \hat{\kappa}, \kappa^*$. We must have $\kappa(t) \geq \kappa^E$ for all $t \geq 0$ as entrants innovate for technology B, so given z_{A0} social welfare at $t = 0$ is

$$\begin{aligned}\frac{\bar{\pi}}{\bar{C}}U = &\int_{z_{A0}}^{\infty}\left[\frac{1}{\rho+\delta-\lambda\eta_A s}z_A\sigma_E K_A(0)+\frac{1}{\rho+\delta}\sigma_E K_B(0)\right]dH_A^*(z_{A0})\\ &+\int_1^{z_{A0}}\left[\frac{1}{\rho+\delta}z_A\sigma_E K_A(0)+\frac{1}{\rho+\delta-\lambda\eta_B s}\sigma_E K_B(0)\right]dH_A^*(z_{A0})\\ &+\delta N\int_0^{\infty}\exp(-\rho t)\left[\frac{1}{\rho+\delta}\sigma_E K_A(t)+\frac{1}{\rho+\delta-\lambda\eta_B s}\sigma_E K_B(t)\right]dt.\end{aligned}$$

Making use of Lemma A.3.3, we find

$$\int_0^\infty \exp(-\rho t) K_A(t) dt = K_A(0)\left(\frac{c_{A1}}{\rho} - \frac{c_{A2}}{\rho+\delta-\lambda\eta_A s}\right),$$
$$\int_0^\infty \exp(-\rho t) K_B(t) dt = K_B(0)\left(\frac{c_{B1}}{\rho+\delta+g_B^*-\lambda\eta_B s} + \frac{c_{B2}}{\rho-g_B^*}\right).$$

The integration constants c_{A1}, c_{A2}, c_{B1}, and c_{B2} depend on z_{A0} and are described in the proof of Lemma A.3.3. For our purposes, their derivatives satisfy

$$\begin{aligned}
\frac{\partial c_{A1}}{\partial z_{A0}} &= -\frac{\lambda\eta_A s\sigma_E}{\delta-\lambda\eta_A s} z_{A0} h_A^*(z_{A0}),\\
\frac{\partial c_{A2}}{\partial z_{A0}} &= \frac{\partial c_{A1}}{\partial z_{A0}},\\
\frac{\partial c_{B1}}{\partial z_{A0}} &= -\frac{\lambda\eta_B s\sigma_E}{2g_B^*+\delta-\lambda\eta_B s} h_A^*(z_{A0}),\\
\frac{\partial c_{B2}}{\partial z_{A0}} &= -\frac{\partial c_{B1}}{\partial z_{A0}}.
\end{aligned}$$

Differentiating social welfare in z_{A0}, we have

$$\begin{aligned}
\frac{\partial U}{\partial z_{A0}} \propto\ & \frac{1}{\rho+\delta} z_{A0}\sigma_E K_A(0) + \frac{1}{\rho+\delta-\lambda\eta_B s}\sigma_E K_B(0) - \left[\frac{1}{\rho+\delta-\lambda\eta_A s} z_{A0}\sigma_E K_A(0) + \frac{1}{\rho+\delta}\sigma_E K_B(0)\right]\\
&+ \frac{\delta N\sigma_E K_A(0)}{\rho+\delta}\left(\frac{1}{\rho} - \frac{1}{\rho+\delta-\lambda\eta_A s}\right)\frac{1}{h_A^*(z_{A0})}\frac{\partial c_{A1}}{\partial z_{A0}}\\
&+ \frac{\delta N\sigma_E K_B(0)}{\rho+\delta-\lambda\eta_B s}\left(\frac{1}{\rho-g_B^*} - \frac{1}{\rho+\delta+g_B^*-\lambda\eta_B s}\right)\frac{1}{h_A^*(z_{A0})}\frac{\partial c_{B2}}{\partial z_{A0}}.
\end{aligned}$$

Note that the right side is strictly decreasing in z_{A0}, so that the first-order condition $\frac{\partial U}{\partial z_{A0}} = 0$ is both necesasry and sufficient to characterize the planner's solution $\hat{z}_{A0}$. Evaluating the expression

above at z^*_{A0}, I observe that the first line collapses to zero by the definition of z^*_{A0}. Hence

$$\begin{aligned}
\left.\frac{\partial U}{\partial z_{A0}}\right|_{z^*_{A0}} &\propto \frac{1}{\rho+\delta}\left(\frac{1}{\rho}-\frac{1}{\rho+\delta-\lambda\eta_A s}\right)\frac{1}{h^*_A(z^*_{A0})}\left.\frac{\partial c_{A1}}{\partial z_{A0}}\right|_{z^*_{A0}} \\
&\quad+\frac{\kappa(0)}{\rho+\delta-\lambda\eta_B s}\left(\frac{1}{\rho-g^*_B}-\frac{1}{\rho+\delta+g^*_B-\lambda\eta_B s}\right)\frac{1}{h^*_A(z^*_{A0})}\left.\frac{\partial c_{B2}}{\partial z_{A0}}\right|_{z^*_{A0}} \\
&= \frac{1}{\rho+\delta}\left(\frac{1}{\rho+\delta-\lambda\eta_A s}-\frac{1}{\rho}\right)\frac{\lambda\eta_A s\sigma_E}{\delta-\lambda\eta_A s}z^*_{A0} \\
&\quad+\frac{\kappa(0)}{\rho+\delta-\lambda\eta_B s}\left(\frac{1}{\rho-g^*_B}-\frac{1}{\rho+\delta+g^*_B-\lambda\eta_B s}\right)\frac{\lambda\eta_B s\sigma_E}{2g^*_B+\delta-\lambda\eta_B s} \\
&\propto -\frac{1}{\rho+\delta}\frac{1}{\rho}\frac{\eta_A}{\rho+\delta-\lambda\eta_A s}z^*_{A0}+\frac{1}{\rho-g^*_B}\frac{1}{\rho+\delta+g^*_B-\lambda\eta_B s}\frac{\eta_B}{\rho+\delta-\lambda\eta_B s}\kappa(0).
\end{aligned}$$

Using the definition $z^*_{A0}=\frac{\kappa(0)}{\kappa^E}$, we have

$$\left.\frac{\partial U}{\partial z_{A0}}\right|_{z^*_{A0}} \propto \frac{1}{\rho-g^*_B}\frac{1}{\rho+\delta+g^*_B-\lambda\eta_B s}-\frac{1}{\rho+\delta}\frac{1}{\rho}.$$

This value is strictly positive if and only if

$$\rho(\rho+\delta)>\left(\rho-g^*_B\right)\left(\rho+\delta+g^*_B-\lambda\eta_B s\right).$$

Simplifying yields

$$g^*_B\left(g^*_B+\delta-\lambda\eta_B s\right)>-\rho\lambda\eta_B s.$$

This inequality always holds because of the assumed lower bound (1.14) on the exit rate δ. We can immediately conclude that $\hat{z}_{A0}>z^*_{A0}$ when $\kappa(0)\geq\hat{\kappa}$.

The remaining statement of the proposition follows by a symmetric argument: Whenever the social planner chooses $\hat{z}_{A0}$ so that the economy converges back to the BGP for technology A, knowledge spillovers lead the social planner to require more initial incumbents to innovate for technology A than in equilibrium. ■

In general, the transition thresholds for the social planner $\hat{\kappa}$ and the equilibrium κ^* cannot be ranked. This holds because the social planner internalizes knowledge spillovers on future entrants when choosing the long-run direction of innovation, but these spillovers are not necessarily always larger for a given technology: Technology B spillovers are larger in the long-run given $\eta_B>\eta_A$, but technology A spillovers may be larger in the short-run given incumbents'

Figure A.3.3: Social Welfare U as a Function of the Cutoff z_{A0}

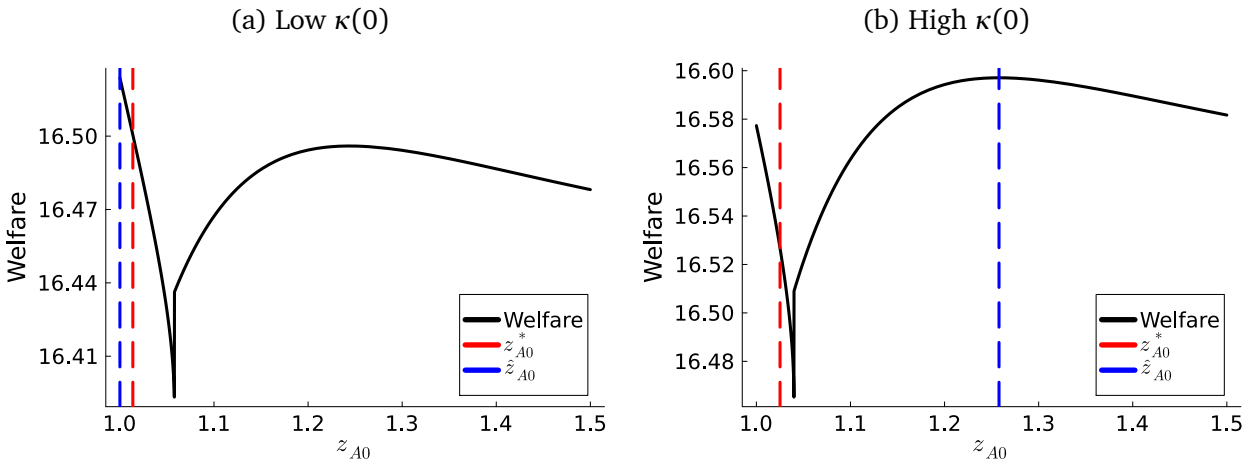

Notes: In each figure, I plot the social planner's objective (social welfare U) as a function of the initial cutoff z_{A0}. I also note the equilibrium cutoff z^*_{A0} by a red dashed line and the efficient cutoff $\hat{z}_{A0}$ by a blue dashed line. The "Low $\kappa(0)$" figure sets $\kappa(0) = 0.89$, while the "High $\kappa(0)$" figure sets $\kappa(0) = 0.90$. All remaining parameters are exactly as in Figure A.3.1.

initial expertise for technology A (i.e., the initial distribution H^*_A). However, Proposition A.3.5 shows that for a given long-run innovation direction, the social planner always prefers to direct greater initial innovation in that direction than in equilibrium.

To provide a sense of the basic non-convexity in the social planner's problem (A.34), Figure A.3.3 displays the objective U as a function of the cutoff z_{A0} for two different values of the initial knowledge stock ratio $\kappa(0)$. In both cases, social welfare attains a global minimum at the "critical mass" value of z_{A0} that divides convergence to the BGP for technology A from convergence to the BGP for technology B. On either side of that point, social welfare attains a global maximum, and the social planner's choice of the long-run innovation direction reduces to comparing the maxima from each side. In both cases, the equilibrium cutoff z^*_{A0} leads to convergence back to the BGP for technology A. When $\kappa(0)$ is low, the social planner also chooses to innovate for technology A in the long run, but consistent with Proposition A.3.5 the optimal cutoff $\hat{z}_{A0}$ is lower than that in equilibrium. When $\kappa(0)$ is high, the social planner instead chooses to innovate for technology B in the long run.

Figure A.3.4 plots the equilibrium and efficient cutoffs z^*_{A0} and $\hat{z}_{A0}$ for different values of $\kappa(0)$. In this example, the social planner prefers to transition to technology B more often than in equilibrium ($\hat{\kappa} < \kappa^*$).

Figure A.3.4: Equilibrium and Efficient Cutoffs z_{A0}

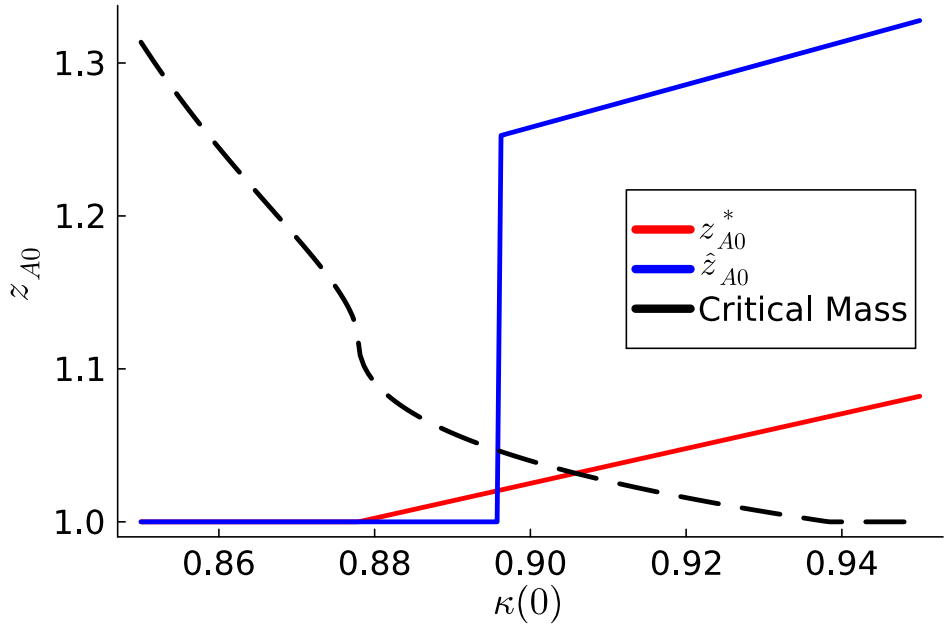

Notes: The dashed black line denotes the cutoff z_{A0} above which the economy converges to the BGP for technology B. The thresholds κ^* and $\hat{\kappa}$ are the values of $\kappa(0)$ where z^*_{A0} and $\hat{z}_{A0}$ intersect this line, respectively. The remaining parameters are exactly as in Figure A.3.1.

A.4 Data

This appendix describes the data for the analyses of Section 1.2 and Section 1.5.

A.4.1 Data Sources

Patents: PatentsView, USPTO-granted patents 1980-2023

- Download date: 08 Oct 2024
- Exclusions:
  - non-utility, withdrawn, reissued patents
  - patents with missing assignee, CPC, filing year, or grant year data
  - patents assigned to Ethicon, Inc. (Johnson & Johnson subsidiary with anomalous forward citation counts, second only to IBM)
- Title and abstract text are processed using the NLTK package in Python to remove standard stop words, punctuation, numbers, and extra white space

US Public Firm – Patent Match: DISCERN 2.0 (Arora et al., 2024)

- Download date: 29 Sept 2024

US Public Firm Financials: Compustat North America Fundamentals Annual

- Download date: 20 Sept 2024

New Technologies: Kalyani et al. (2023)

- Download date: 14 Sept 2024

Aggregate US R&D Expenditures: "Research and Development: US Trends and International Comparisons," National Science Board (NSB-2024-6)

- Download date: 20 Sept 2024

US GDP Deflator/GDP per Capita: Federal Reserve Economic Data (FRED)

A.4.2 Data Build: mRNA Case Study

mRNA Therapy Patents. I identify patents related to mRNA therapies by keyword search. I start with all patents whose title or abstract mention at least one term from each of the following lists:

mRNA terms: "mrna", " rna", "rna ", " ribonucleic"
therapy terms: "therap", "treat", "vaccin ", "innocul", "immun"

I then exclude patents that mention terms related to recombinant DNA/RNA or other types of RNA, which are involved in treatment technologies distinct from mRNA technology:

exclusion terms (RNA): "recombin", "rna interfer", "irna ", " rnai", "mirna", "sirna", "dsrna", "trna", "transfer rna", "double stranded rna", "small interfering rna", "double-stranded rna", "small-interfering rna", "micro rna", "micro-rna", "microrna", "reduce expression", "reducing expression", "inhibit expression", "inhibiting expression", "modulate expression", "modulating expression"

This procedure yields 3408 mRNA therapy patents granted between 1980 and 2023 to 1211 unique PatentsView assignees. Figure A.4.1(a) displays the number of patents filed and granted over 1980-2023.

Conventional Vaccine Patents. I identify patents related to conventional vaccines by keyword search. I start with all patents whose title or abstract mention at least one term from the following list:

vaccine terms: "vaccin", "innocul", "immuniz"

I then exclude patents that mention any of the exclusion terms for RNA technologies noted above. I also exclude any patents that mention terms related to cancer, as cancer vaccines constitute a distinct technology from conventional vaccines for infectious diseases:

exclusion terms (cancer): "cancer", "tumor", "tumour", "oncolog", "oncogen", "malign", "mestast", "neoplas"

Finally, I exclude patents found in the set of mRNA therapy patents constructed above. This procedure yields 9868 conventional vaccine patents granted between 1980 and 2023 to 2683 unique PatentsView assignees. Figure A.4.1(b) displays the number of patents filed and granted over 1980-2023.

Incumbent Patents. I determine the set of mRNA therapy and conventional vaccine patents assigned to any of the following top 20 pharmaceutical firms by keyword search on patent assignee names:

Top 20 Pharma: Johnson & Johnson, Sinopharm, Roche, Merck, Pfizer, AbbVie, Bayer, Sanofi, AstraZeneca, Novartis, Bristol-Myers Squibb, GSK, Eli Lilly Novo Nordisk, Shanghai Pharmaceuticals, Takeda, Amgen, Boehringer Ingelheim, Gilead Sciences, Siemens Healthineers

I search patent assignee names for keywords related to each of these firms, including the names of any major subsidiaries or recently acquired firms. This procedure identifies 243 mRNA therapy patents and 1385 conventional vaccine patents granted to the firms above.

Entrant Patents. I determine the set of mRNA therapy and conventional vaccine patents assigned to Moderna, BioNTech, CureVac, or RNARx by keyword search on patent assignees. To more fully capture these firms' expertise, I also include all patents that list one of their founders as an inventor. This procedure identifies 242 mRNA therapy patents and 47 conventional vaccine patents granted to the four entrant firms or their founders.

Forward Citations. I compute the number of forward citations received by each patent from any other patent in the PatentsView dataset. These counts suffer from truncation bias that becomes severe for recently-granted patents. To correct for this, I normalize each patent's forward citations by the average number of forward citations received by any USPTO patent granted in the same year. The normalized citation count of a patent p granted in year t is then

$$Count_p \equiv FCites_p \times \frac{|\{p' \text{ granted in } t\}|}{\sum_{p' \text{ granted in } t} FCites_{p'}}.$$

The average value of $Count_p$ across all mRNA therapy patents is 1.84, while the average across all conventional vaccine patents is 0.62.

Sample. On average, mRNA therapy and conventional vaccine patents each have a 3.5-year lag between filing and publication (see Figure A.4.2). Since I only observe granted patents, to avoid truncation bias I consider only patents filed through 2018 in Figure 1.2. Approximately 85% of each of mRNA therapy and conventional vaccine patents have a publication lag of five years or less.

Figure A.4.1: Patents Filed and Granted, 1980-2023

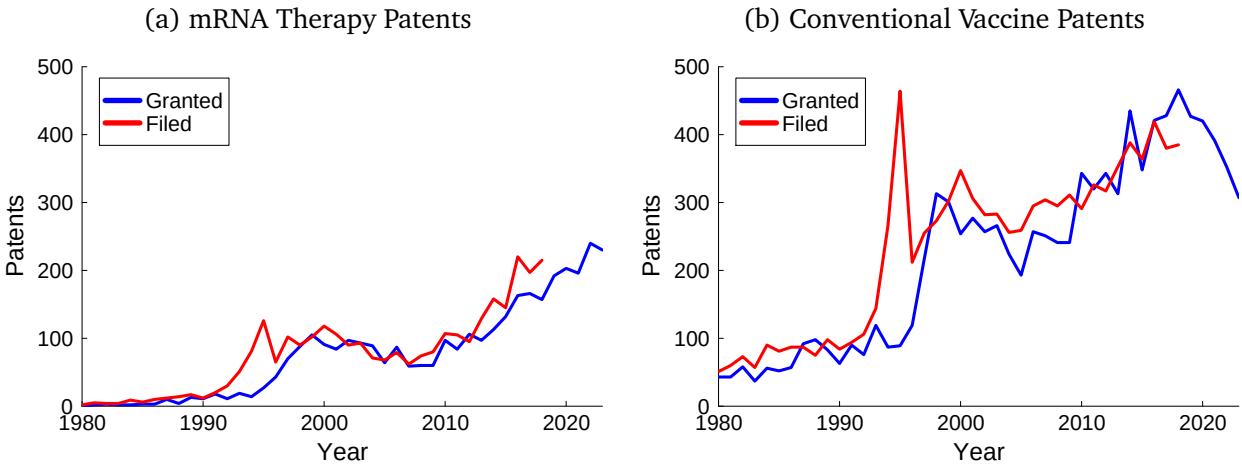

Notes: Given the lag between filing year and grant year documented in Figure A.4.2, I halt the "Filed" series in 2018. The spike in patent filings for conventional vaccines in 1995 coincides with implementation of the TRIPS agreement, which required WTO members to respect pharmaceutical patents.

Figure A.4.2: Lag from Patent Filing to Publication

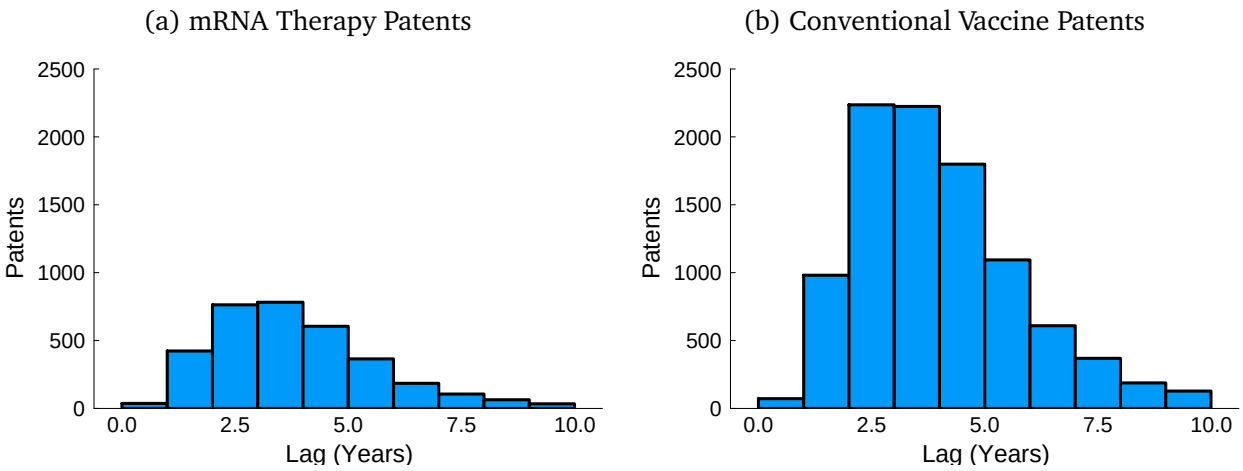

Notes: For both sets of patents, over 85% of patents have a lag from filing to publication of five years or less.

A.4.3 Data Build: Empirical Analysis

Firm-Patent Match. I use the DISCERN 2.0 dataset to recover all USPTO patents granted to US public firms from 1980 to 2021 (Arora et al., 2024). Each patent is matched to each of its original assignees; any firm with at least three consecutive years of coverage in Compustat North America is a potential match. I identify a firm with a unique *permno_adj* identifier defined in DISCERN 2.0, which is dynamically linked to *gvkey* identifiers in Compustat to allow consistent access to a firm's financial records over time (for example, accounting for name changes or acquisitions that could result in multiple *gvkey*s). The set of patents associated with a firm is the set of patents associated with the corresponding *permno_adj* identifier at the time of patent publication. I currently do not re-assign past patents associated with one *permno_adj* identifier to another after an acquisition, potentially biasing downward my estimates of firm knowledge stocks.

Forward Citations. I use yearly citation-weighted patent counts to measure a firm's innovative activity. The use of forward citations as a proxy for patent quality comes with two well-known limitations (e.g. Lerner and Seru, 2022): First, patents can only be cited after publication, so patents with more recent publication dates have mechanically fewer forward citations on average. Second, citation rates across different technology areas may be different independent of average patent quality (e.g., area-specific norms, differences in the combinatorial nature of innovation, etc.). As a result, raw citation counts may not yield a measure of quality comparable across technology areas.

I address these concerns with the following two-step normalization procedure. First, for each year t, I compute the mean number of forward citations $FCites_p$ across all patents p granted in year t:

$$\overline{FCites}_t \equiv \frac{\sum_{p \text{ granted in } t} FCites_p}{|\{p \text{ granted in } t\}|}.$$

This mean is plotted for each year t from 1980 to 2023 in Figure A.4.3(a). I define $FCites_{p1}$ for each patent p granted in year t as the normalized forward citation count

$$FCites_{p1} \equiv \frac{FCites_p}{\overline{FCites}_t}.$$

This step provides a simple correction for the truncation issue noted above. But I note that it can also introduce a bias of its own if citation counts are particularly high in a given year because a large number of high-quality patents were granted: This normalization effectively

equalizes the total measure of quality across all grant years.

Second, for each CPC subclass s, I compute the mean of $FCites_{p1}$ across all patents p assigned to class s:

$$\overline{FCites}_{s1} \equiv \frac{\sum_{p \in s} FCites_{p1}}{|\{p \in s\}|}.$$

The distribution of this mean across subclasses s is plotted in Figure A.4.3(b). I define $FCites_{p2}$ for each patent p as the normalized value of $\overline{FCites}_{p1}$, where the normalizing factor is the mean of $\overline{FCites}_{s1}$ across all subclasses s to which p is assigned:

$$FCites_{p2} \equiv FCites_{p1} \frac{|\{s \ni p\}|}{\sum_{s \ni p} \overline{FCites}_{s1}}.$$

This step corrects for differences in citation rates across CPC subclasses that are uniform over time. It can introduce bias if citation counts are particularly high in a given CPC subclass because a large number of high-quality patents were assigned to that subclass over the sample: This normalization effectively equalizes the total measure of quality across CPC subclasses over the sample.

Throughout the empirical analysis of Section 1.5, the "citation weight" of a patent is equal to $FCites_{p2}$. This normalization procedure is weaker than that used by, for example, Hall et al. (2001) and Kalyani et al. (2023), which normalizes citations within each class-year. My normalization allows the total measure of quality to vary across subclasses at each time and over time with a subclass, ensuring that it can reflect "innovation bursts" for particular subclasses over a period of several years (e.g., for personal computing-related classes in the 1980s and 1990s). It is the minimal normalization that plausibly allows comparability between patents granted at different times and in different technology areas.

Citation-Weighted Patents. The citation-weighted patent flow for firm i in year t is simply the sum of $FCites_{p2}$ across all patents filed in year t and assigned to firm i:

$$Pat_{it} \equiv \sum_{\{p \text{ filed in } t \text{ by } i\}} FCites_{p2}.$$

Figure A.4.4 displays raw and citation-weighted patent counts by year, both for the mean across firms in the sample of Section 1.5 and for a particular firm (Microsoft).

Generic Knowledge Stocks. As noted in Section 1.5, firm i's generic internal knowledge stock at time t is constructed by the perpetual inventory method from the patent flows Pat_{it}:

$$K_{it}^{\text{Firm}} \equiv (1-\nu)K_{it-1}^{\text{Firm}} + Pat_{it-1}.$$

If firm i becomes public in year t_0, I initialize by setting $K_{it_0} = 0$ because I do not observe the firm's previous patenting behavior. To help address this measurement error, each regression in Section 1.5 is restricted to firms that have been public for at least one year and includes fixed effects for the number of years for which the firm has been publicly listed.

To construct the aggregate knowledge stock K_{it}^{Agg}, I first compute the share of each firm i's patents assigned to each CPC subclass s over the sample:[6]

$$\omega_{is} \equiv \frac{|\{p \in i \cup s\}|}{|\{p \in i\}|}.$$

Since a single patent is generally assigned to many subclasses, the weights ω_{is} do not sum to one across s. I defined normalized weights

$$\tilde{\omega}_{is} \equiv \frac{\omega_{is}}{\sum_{s'} \omega_{is'}}.$$

I also compute the aggregate generic knowledge stock for each CPC subclass s from the yearly flows of patents Pat_{st} filed in subclass s:

$$K_{st}^{\text{Agg}} \equiv (1-\nu)K_{st-1}^{\text{Agg}} + Pat_{st-1},$$

with initial value $K_{s1980} = 0$. I then construct firm i's aggregate knowledge stock as the patent share-weighted average

$$K_{it}^{\text{Agg}} \equiv \sum_{s} \tilde{\omega}_{is} K_{st}^{\text{Agg}}.$$

Figure A.4.5 displays log generic firm and aggregate knowledge stocks, both for the mean across firms in the sample of Section 1.5 and for a particular firm (Microsoft). Note that for the sample means in Figure A.4.5(a), both series increase roughly linearly over the sample period 1985 to 2016, indicating exponential growth in firm and aggregate knowledge stocks. Figure A.4.6 displays the 1990 and 2010 distributions of the log firm and aggregate knowledge stocks across

[6] I perform this step of the construction at the subclass level to ensure a more precise characterization of the knowledge relevant for each firm's innovation.

firms in the sample of Section 1.5. The firm knowledge stock distribution displays substantially greater dispersion than the aggregate knowledge stock distribution.

Technology Knowledge Stocks. Firm i's internal knowledge stock for technology θ at time t is defined analogously to the generic knowledge stock K_{it}^{Firm}:

$$K_{i\theta t}^{\text{Firm}} \equiv (1-\nu)K_{i\theta t-1}^{\text{Firm}} + Pat_{i\theta t-1},$$

where

$$Pat_{i\theta t} \equiv \sum_{\{p \text{ filed in } t \text{ for } \theta \text{ by } i\}} FCites_{p2}.$$

If firm i becomes public in year t_0, I initialize by setting $K_{i\theta t_0} = 0$.

The aggregate knowledge stock for technology θ at time t is simply constructed by the perpetual inventory method from the yearly flows of patents $Pat_{\theta t}$ filed for technology θ:

$$K_{\theta t}^{\text{Agg}} \equiv (1-\nu)K_{\theta t-1}^{\text{Agg}} + Pat_{\theta t-1}.$$

Figure A.4.3: Citation Normalizing Factors

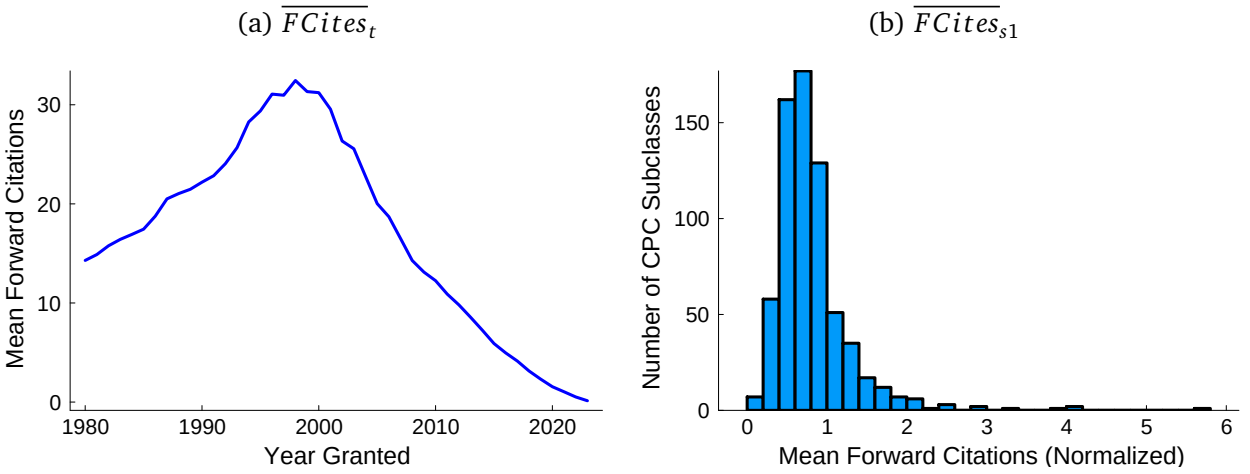

Notes: Figure A.4.3(a) displays the mean number of forward citations across all patents granted in each year t, $\overline{FCites}_t$. Figure A.4.3(b) displays the distribution (across subclasses) of the mean number of (year-normalized) forward citations for all patents in a subclass, $\overline{FCites}_{s1}$.

Figure A.4.4: Patent Counts, 1980-2021

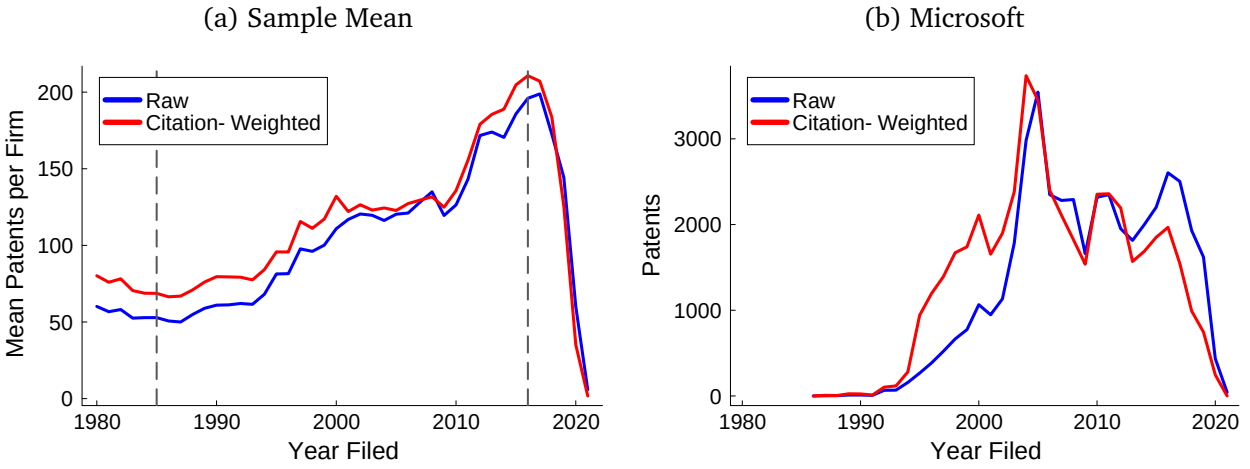

Notes: Figure A.4.4(a) displays the mean number of patents (raw and citation-weighted) filed each year across firms in the sample used in Section 1.5. Figure A.4.4(b) displays the corresponding patent counts for Microsoft. The vertical dashed lines in Figure A.4.4(a) denote the beginning (1985) and end (2016) of the sample period for the regressions in Section 1.5. Microsoft becomes publicly listed in 1986.

Figure A.4.5: Generic Knowledge Stocks, 1980-2021

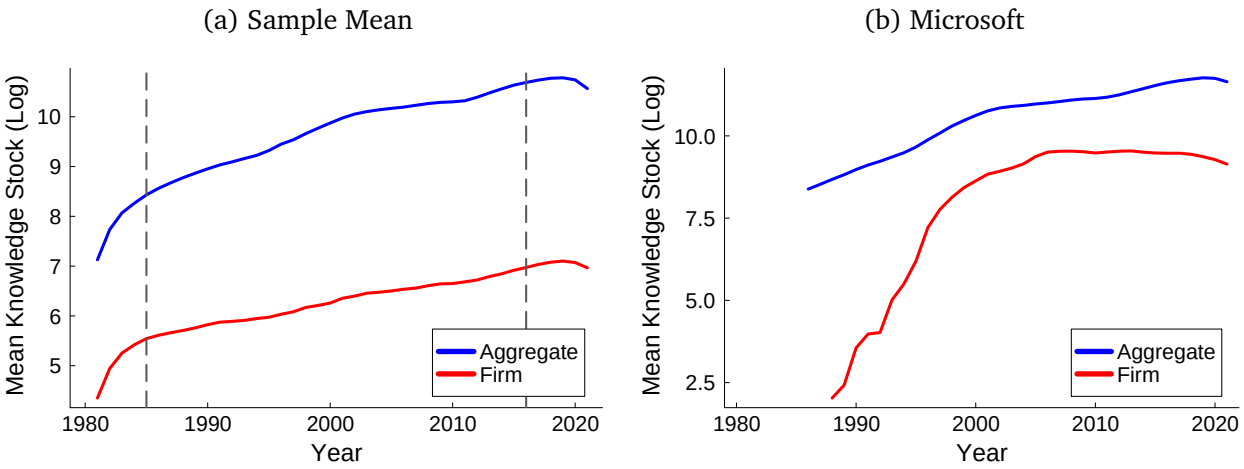

Notes: Figure A.4.5(a) displays the log of the mean firm and aggregate knowledge stocks each year across firms in the sample used in Section 1.5. Figure A.4.5(b) displays the corresponding log knowledge stocks for Microsoft. The vertical dashed lines in Figure A.4.5(a) denote the beginning (1985) and end (2016) of the sample period for the regressions in Section 1.5. Microsoft becomes publicly listed in 1986.

Figure A.4.6: Generic Knowledge Stock Distributions, 1990 and 2010

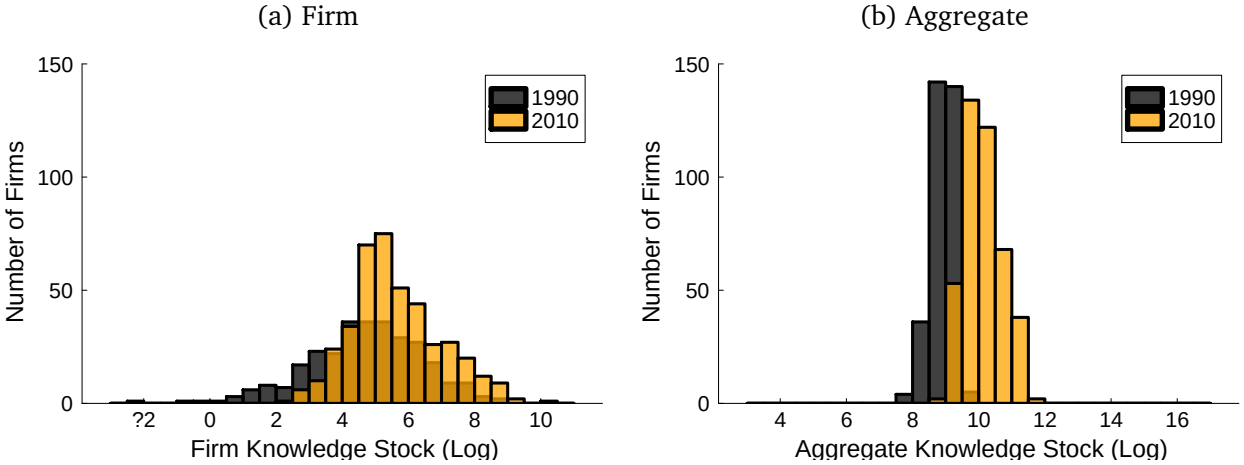

Notes: Figure A.4.6(b) displays the 1990 and 2010 distributions of the log firm knowledge stock across firms in the sample used in Section 1.5. Figure A.4.6(a) displays the corresponding distributions for the log aggregate knowledge stock.

A.5 Empirics: Additional Results and Robustness

This appendix provides additional results for the empirical analysis of Section 1.5.

A.5.1 Descriptive Statistics: New Technologies

Table A.5.1: Technology Summary Statistics

	p25	p50	p75	Mean
Overall				
Patents (Raw)	2733	6823	14658	13204
Patents (Weighted)	5429	15249	34763	27804
Emergence Year	1988	1991	1994	1991
Firm Sample				
Share of Patents (Raw)	0.29	0.37	0.44	0.37
Share of Patents (Weighted)	0.3	0.43	0.54	0.43
Number of Active Firms	66	128	194	143
Max Firm Share of Patents (Raw)	0.04	0.06	0.08	0.07
Max Firm Share of Patents (Weighted)	0.06	0.11	0.22	0.16

Notes: The top panel of Table A.5.1 displays general summary statistics for the set of new technologies described in Section 1.5. The bottom panel displays additional summary statistics after restricting to the sample of public firms described in Section 1.5. For example, "Share of Patents (Raw)" gives the share of all (unweighted) patents for a given technology that are filed by firms in the baseline sample. "Number of Active Firms" gives the number of firms with at least one patent for a given technology. "Max Firm Share of Patents (Raw)" gives the maximum share of (unweighted) patents filed by a single firm in the sample for a given technology.

Figure A.5.1: New Technology Emergence Years

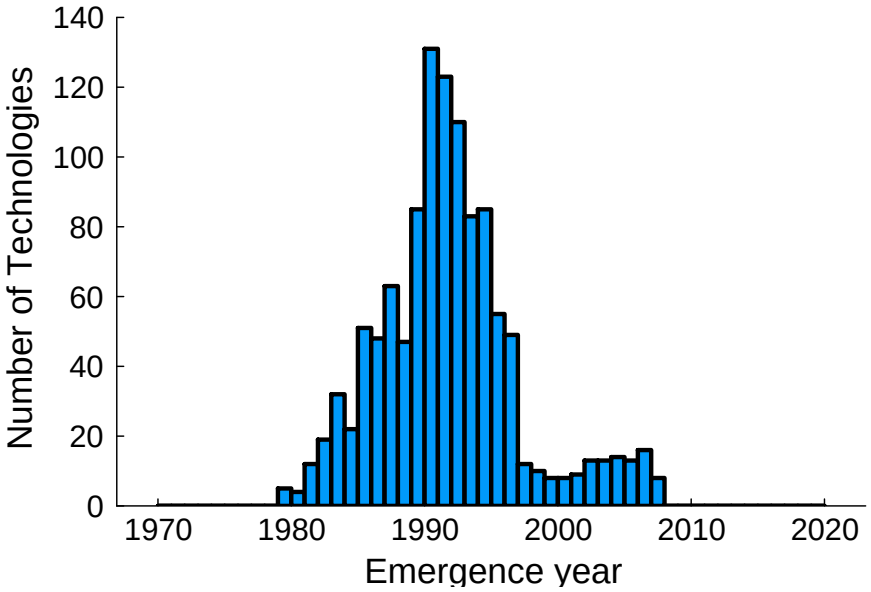

Notes: Figure A.5.1 displays the number of new technologies that emerged in each year from 1970 to 2020.

A.5.2 Patent Trajectories: New Technologies

To describe the average innovation activity for a new technology over time, I estimate several regressions of the form

$$Y_{\theta t} = \alpha_\theta + \sum_{k=-5}^{20} \beta_k \mathbb{1}_{t-T_\theta=k} + \varepsilon_{\theta t}. \tag{A.35}$$

I take the dependent variable $Y_{\theta t}$ to be the total number of citation-weighted patents $Pat_{\theta t}$ or the aggregate knowledge stock $K^{\text{Agg}}_{\theta t}$. Here α_θ is a technology fixed effect, while the coefficients β_k describe the (average) trajectory of total patenting or knowledge for a new technology around its emergence time T_θ. I normalize $\beta_{-5} \equiv 0$.

Figure A.5.2 displays the estimated fixed effects $\hat{\beta}_k$ for dependent variables $Y_{\theta t} = Pat_{\theta t}$ and $Y_{\theta t} = \log(Pat_{\theta t})$. Both figures indicate exponential growth in citation-weighted patents from five years before emergence to roughly five years after emergence, with a trend break to slower growth thereafter. Figure A.5.3 displays the estimated coefficients for dependent variables $Y_{\theta t} = K^{\text{Agg}}_{\theta t}$ and $Y_{\theta t} = \log\left(K^{\text{Agg}}_{\theta t}\right)$. The same pattern observed for patents in Figure A.5.2 is clear for the knowledge stock in Figure A.5.3.

The results in Figure A.5.2 have an important implication for the last empirical finding in Section 1.5: Figure 1.6 shows that experienced incumbents begin patenting a new technology at rates comparable to less experienced incumbents only with a substantial lag after emergence. Figure A.5.2 shows that this *cannot* be explained by falling patenting for the new technology in aggregate, which would also collapse the "level difference" between experienced and inexperienced incumbents. Rather, the data suggest growing patenting activity by experienced incumbents over time.

Figure A.5.2: Average Technology Patent Trajectory

(a) Patents (Level): $Pat_{\theta t}$ (b) Patents (Log): $\log(Pat_{\theta t})$

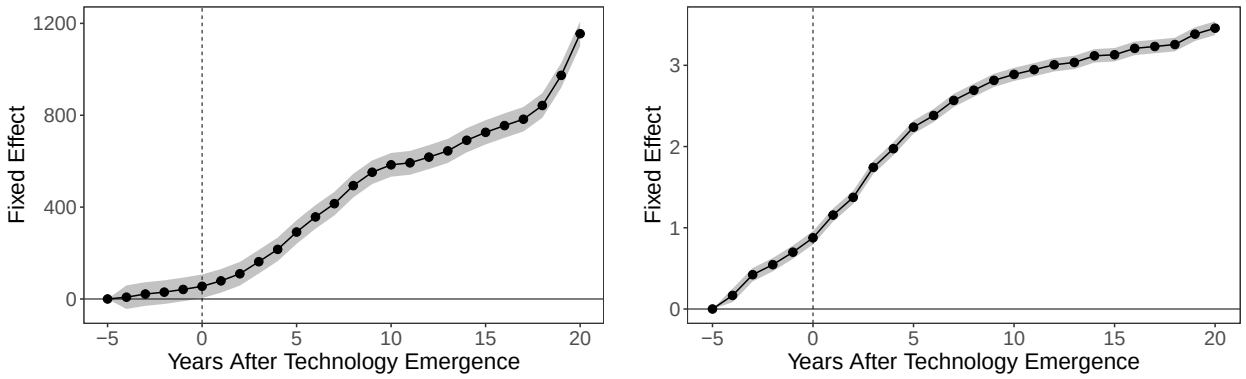

Notes: Figures A.5.2(a) and A.5.2(b) display the estimated fixed effects $\hat{\beta}_k$ corresponding to the regression (A.35) for dependent variables $Pat_{\theta t}$ and $\log(Pat_{\theta t})$, respectively. The gray area in each figure denotes the 95% confidence interval.

Figure A.5.3: Average Technology Knowledge Trajectory

(a) Knowledge Stock (Level): $K_{\theta t}^{\text{Agg}}$ (b) Knowledge Stock (Log): $\log\left(K_{\theta t}^{\text{Agg}}\right)$

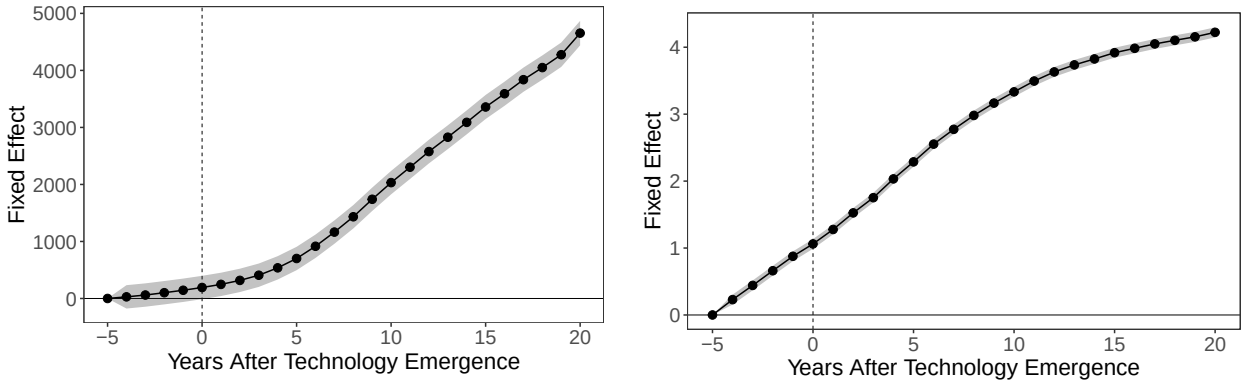

Notes: Figures A.5.3(a) and A.5.3(b) display the estimated fixed effects $\hat{\beta}_k$ corresponding to the regression (A.35) for dependent variables $K_{\theta t}^{\text{Agg}}$ and $\log(K_{\theta t}^{\text{Agg}})$, respectively. The gray area in each figure denotes the 95% confidence interval.

A.5.3 Additional Results: Knowledge and Patenting (Facts 1 and 2)

This section provides additional results about the relationship between past and current patenting, extending the analysis of Facts 1 and 2 in Section 1.5.

Spillover Specification. The aggregate knowledge stocks K_{it}^{Agg} and $K_{\theta t}^{\text{Agg}}$ used in Section 1.5 and defined in Appendix A.4 include each firm i's own past patents, consistent with the definition of the knowledge stocks in the model (1.9). To explore whether the empirical results are robust to excluding firm i's past patents from these aggregates, I define the *spillover knowledge stocks*

$$
\begin{aligned}
K_{it}^{\text{Spill}} &\equiv \sum_{s} \tilde{\omega}_{is} K_{ist}^{\text{Spill}}, \\
K_{ist}^{\text{Spill}} &\equiv (1-\nu) K_{ist-1}^{\text{Spill}} + (Pat_{st-1} - Pat_{ist-1}), \\
K_{i\theta t}^{\text{Spill}} &\equiv (1-\nu) K_{i\theta t-1}^{\text{Spill}} + (Pat_{\theta t-1} - Pat_{i\theta t-1}).
\end{aligned}
$$

The spillover knowledge stocks K_{it}^{Spill} and $K_{i\theta t}^{\text{Spill}}$ are defined analogously to the aggregate knowledge stocks K_{it}^{Agg} and $K_{\theta t}^{\text{Agg}}$, but excluding firm i's patents at each time t.

I estimate the regressions (1.33, 1.34) after substituting the spillover knowledge stocks for the aggregate knowledge stocks. The corresponding OLS and negative binomial estimates are found in Table A.5.2, which is analogous to Table 1.2. The estimates are generally unchanged from those in Table 1.2.

Similarity Robustness. To explore whether the cosine similarity cutoff plays an important role for the technology patenting results in Section 1.5.3, I re-estimate the regression (1.34) after imposing the stronger cutoff of 0.8. Table A.5.3 displays the baseline estimates with the weaker cutoff of 0.6 in columns (1) and (2), which correspond to columns (3) and (4) in Table 1.2. The estimates for the stronger cutoff are given in columns (3) and (4) of Table A.5.3. The estimates are generally unchanged, with some evidence of stronger correlations between past technology-specific patenting and current technology-specific patenting under the stronger cutoff. This is consistent with measurement error attenuating coefficient estimates under the weaker cutoff.

Table A.5.2: Firm Patenting Regressions (Spillover Specification)

	Overall Patenting		**Technology Patenting**	
	(1) Neg. Bin.	(2) OLS	(3) Neg. Bin.	(4) OLS
$\log(K_{it}^{\text{Firm}})$	0.7301***	0.7611***	0.3532***	0.0723***
	(0.0312)	(0.0308)	(0.0599)	(0.0147)
$\log(K_{it}^{\text{Spill}})$	0.1322	0.0862	-0.4137	-0.1490***
	(0.1048)	(0.1006)	(0.2626)	(0.0531)
$\log(s_{it})$	0.0657***	0.0787***	0.0880**	0.0209***
	(0.0171)	(0.0185)	(0.0431)	(0.0079)
$\log(K_{i\theta t}^{\text{Firm}})$			0.5184***	0.3342***
			(0.0142)	(0.0229)
$\log(K_{i\theta t}^{\text{Spill}})$			0.1855***	0.0362***
			(0.0313)	(0.0047)
Observations	12,237	12,237	562,080	564,062
Firms	555	555	435	453

*Signifance: ***: 0.01, **: 0.05, *: 0.1*

Notes: The overall patenting regressions include fixed effects by year, firm, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero internal knowledge stock K_{it}^{Firm}, zero spillover knowledge stock K_{it}^{Spill}, and zero R&D expenditures s_{it}. The technology-specific patenting regressions additionally include fixed effects by technology and the time since emergence. They also add dummy variables for zero internal and spillover technology-specific knowledge stocks $K_{i\theta t}^{\text{Firm}}$ and $K_{i\theta t}^{\text{Spill}}$. All standard errors are clustered at the firm level.

Table A.5.3: Technology-Specific Patenting Regressions (Similarity Robustness)

	Cosine Similarity $\geq$ 0.6		**Cosine Similarity $\geq$ 0.8**	
	(1) Neg. Bin.	(2) OLS	(3) Neg. Bin.	(4) OLS
$\log(K_{it}^{\text{Firm}})$	0.3557***	0.0714***	0.2965***	0.0805***
	(0.0605)	(0.0145)	(0.0586)	(0.0172)
$\log(K_{it}^{\text{Agg}})$	-0.4386*	-0.1386***	-0.1354	-0.1343*
	(0.2641)	(0.0530)	(0.3238)	(0.0717)
$\log(s_{it})$	0.0877**	0.0211***	0.1069*	0.0253**
	(0.0431)	(0.0079)	(0.0565)	(0.0126)
$\log(K_{i\theta t}^{\text{Firm}})$	0.5156***	0.3338***	0.5061***	0.3857***
	(0.0140)	(0.0228)	(0.0195)	(0.0295)
$\log(K_{\theta t}^{\text{Agg}})$	0.2093***	0.0416***	0.1918***	0.0623***
	(0.0330)	(0.0046)	(0.0385)	(0.0067)
Observations	562,080	564,062	203,433	204,292
Firms	435	453	298	305

*Significance: ***: 0.01, **: 0.05, *: 0.1*

Notes: All regressions have technology-specific patenting $Pat_{i\theta t}$ as the dependent variable, either in levels (negative binomial specifications) or logs (OLS specifications). All regressions include fixed effects by year, firm, the number of years for which the firm has been publicly listed, technology, and the time since emergence. They also include dummy variables for zero internal knowledge stock K_{it}^{Firm}, zero aggregate knowledge stock K_{it}^{Agg}, zero R&D expenditures s_{it}, and zero internal and aggregate technology-specific knowledge stocks $K_{i\theta t}^{\text{Firm}}$ and $K_{\theta t}^{\text{Agg}}$. All standard errors are clustered at the firm level.

A.5.4 Additional Results: Incumbents and New Technologies (Fact 3)

This section provides additional results about the relationship between a firm's past patenting experience and its future patenting of emerging technologies, extending the analysis of Fact 3 in Section 1.5.

Long Difference: Full Results. Table A.5.4 provides the full set of coefficient estimates for regression (1.35) in Section 1.5. The first row of Table A.5.4 corresponds exactly to Table 1.3.

Long Difference: Estimation Method. Table A.5.5 provides coefficient estimates for regression (1.35) from OLS and negative binomial regressions. The negative binomial regression is instead

$$Pat_{i\theta}^{10} = \exp(\beta_1 \log\left(K_{iT_\theta}^{\text{Firm}}\right) + \beta_2 \log\left(K_{iT_\theta}^{\text{Agg}}\right) + \beta_3 \log\left(s_{iT_\theta}\right) + \beta_4 \log\left(K_{i\theta T_\theta}^{\text{Firm}}\right) + \beta_5 \log\left(K_{\theta T_\theta}^{\text{Agg}}\right) + \alpha' X_{i\theta} + \varepsilon_{i\theta}). \tag{A.36}$$

Columns (1) and (2) of Table A.5.5 correspond exactly to columns (1) and (2) of Table A.5.4 and give the OLS estimates for regression (1.35), while columns (3) and (4) of Table A.5.5 give the negative binomial estimates for regression (A.36). The negative binomial estimates again demonstrate that experienced incumbent firms patent emerging technologies at lower rates than less-experienced firms.

Dynamic Difference: Balanced Panel. In the baseline estimation of the dynamic level difference regression (1.36), the firm-technology panel is unbalanced because of firm entry and exit. To explore whether this substantially influences the results, I subset the panel to firm-technology pairs such that each firm is continuously publicly listed from 6 years before to 18 years after the emergence of the new technology. The corresponding estimates $\hat{\beta}_{1w}$ for the marginal effect of initial firm patenting experience on its subsequent patenting for the new technology are plotted in Figure A.5.4. For both total patenting $\log(1 + Pat_{i\theta w})$ and the technology's share of firm patenting $TechShare_{i\theta w}$, Figure A.5.4 displays the same patterns observed in the baseline estimates of Figure A.5.4. With the balanced panel, the coefficient estimates are somewhat attenuated, and the standard errors are naturally larger given the smaller sample size.

Dynamic Difference: Full Interactions. In the dynamic level difference regression (1.36), only the marginal effect of $K_{iT_\theta}^{\text{Firm}}$ can vary across time. To ensure that this restriction is not

driving the results, I instead estimate a version of the regression with time window interactions for all covariates:

$$\begin{aligned}\log(1+Pat_{i\theta w}) = \sum_{\tilde{w}=-2}^{6} [&\beta_{1\tilde{w}} \log\left(K^{\text{Firm}}_{iT_\theta}\right) + \beta_{2\tilde{w}} \log\left(K^{\text{Agg}}_{iT_\theta}\right) + \beta_{3\tilde{w}} \log\left(s_{iT_\theta}\right) \\ &+ \beta_{4\tilde{w}} \log\left(K^{\text{Firm}}_{i\theta T_\theta}\right) + \beta_{5\tilde{w}} \log\left(K^{\text{Agg}}_{\theta T_\theta}\right)] \mathbb{1}_{w=\tilde{w}} \\ &+ \alpha' X_{i\theta w} + \varepsilon_{i\theta w}.\end{aligned} \tag{A.37}$$

The corresponding estimates $\hat{\beta}_{1w}$ are plotted in Figure A.5.5, and are broadly similar to the baseline estimates in Figure 1.6.

Dynamic Difference: No Early Patents. To ensure that firms with patents for a new technology before its emergence do not drive the results in Figure 1.6, I estimate the dynamic level difference regression (1.36) after subsetting to firm-technology pairs with no such "early patents." The corresponding estimates are plotted in Figure A.5.6, and are broadly similar to the baseline estimates in Figure 1.6.

Table A.5.4: Technology Patenting Level Difference (1.35)

	Technology Patents		**Technology Patent Share**	
	(1) Full Sample	(2) No Early Patents	(3) Full Sample	(4) No Early Patents
$\log(K^{\text{Firm}}_{iT_\theta})$	-0.0694**	-0.0848***	-0.0023***	-0.0016**
	(0.0345)	(0.0273)	(0.0007)	(0.0007)
$\log(K^{\text{Agg}}_{iT_\theta})$	-0.2995	-0.2236	-0.0142***	-0.0132***
	(0.2394)	(0.1946)	(0.0042)	(0.0036)
$\log(s_{iT_\theta})$	0.0028	-0.0071	-0.0001	-0.0001
	(0.0150)	(0.0133)	(0.0005)	(0.0003)
$\log(K^{\text{Firm}}_{i\theta T_\theta})$	0.3115***		0.0069***	
	(0.0187)		(0.0011)	
$\log(K^{\text{Agg}}_{\theta T_\theta})$	0.1686***	0.1462***	0.0036***	0.0030***
	(0.0112)	(0.0117)	(0.0004)	(0.0004)
R^2	0.57658	0.48188	0.22757	0.15679
Observations	13,674	10,974	13,624	10,935
Dep. Var. Mean	1.5817	1.2574	0.01432	0.01157
Dep. Var. SD	1.8125	1.6162	0.03968	0.03451

*Signifance: ***: 0.01, **: 0.05, *: 0.1*

Notes: All regressions include fixed effects by firm, emergence year, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. All standard errors are clustered at the firm level.

Table A.5.5: Technology Patenting Level Difference (1.35) (OLS vs Neg. Bin.)

	OLS		**Negative Binomial**	
	(1) Full Sample	(2) No Early Patents	(3) Full Sample	(4) No Early Patents
$\log(K^{\text{Firm}}_{iT_\theta})$	-0.0694**	-0.0848***	-0.0805	-0.1285**
	(0.0345)	(0.0273)	(0.0568)	(0.0622)
$\log(K^{\text{Agg}}_{iT_\theta})$	-0.2995	-0.2236	-0.0732	-0.3153
	(0.2394)	(0.1946)	(0.4532)	(0.4756)
$\log(s_{iT_\theta})$	0.0028	-0.0071	0.0415	0.0042
	(0.0150)	(0.0133)	(0.0299)	(0.0271)
$\log(K^{\text{Firm}}_{i\theta T_\theta})$	0.3115***		0.2659***	
	(0.0187)		(0.0353)	
$\log(K^{\text{Agg}}_{\theta T_\theta})$	0.1686***	0.1462***	0.2294***	0.2384***
	(0.0112)	(0.0117)	(0.0210)	(0.0254)
R^2	0.57658	0.48188		
Observations	13,674	10,974	13,539	10,764
Dep. Var. Mean	1.5817	1.2574	34.862	19.584
Dep. Var. SD	1.8125	1.6162	142.48	70.016

*Significance: ***: 0.01, **: 0.05, *: 0.1*

Notes: All regressions have technology-specific patenting $Pat^{10}_{i\theta}$ as the dependent variable, either in levels (negative binomial specifications) or logs (OLS specifications). All regressions include fixed effects by firm, emergence year, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. All standard errors are clustered at the firm level.

Figure A.5.4: Technology Patenting Level Difference, Dynamic (Balanced Panel)

(a) Patents: $\log(1 + Pat_{i\theta w})$ (b) Patent Share: $TechShare_{i\theta w}$

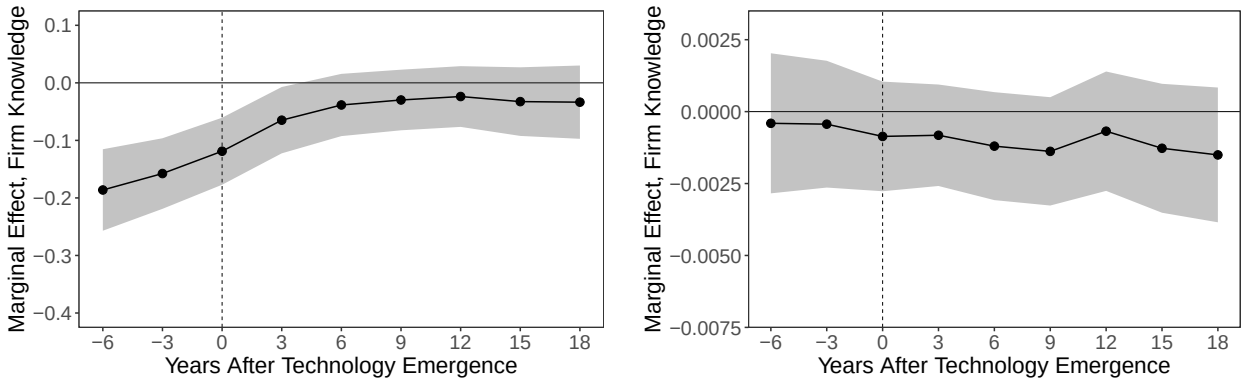

Notes: All regressions include fixed effects by firm-year, technology-window, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. All standard errors are clustered at the firm level.

Figure A.5.5: Technology Patenting Level Difference, Dynamic (Full Interaction)

(a) Patents: $\log(1 + Pat_{i\theta w})$ (b) Patent Share: $TechShare_{i\theta w}$

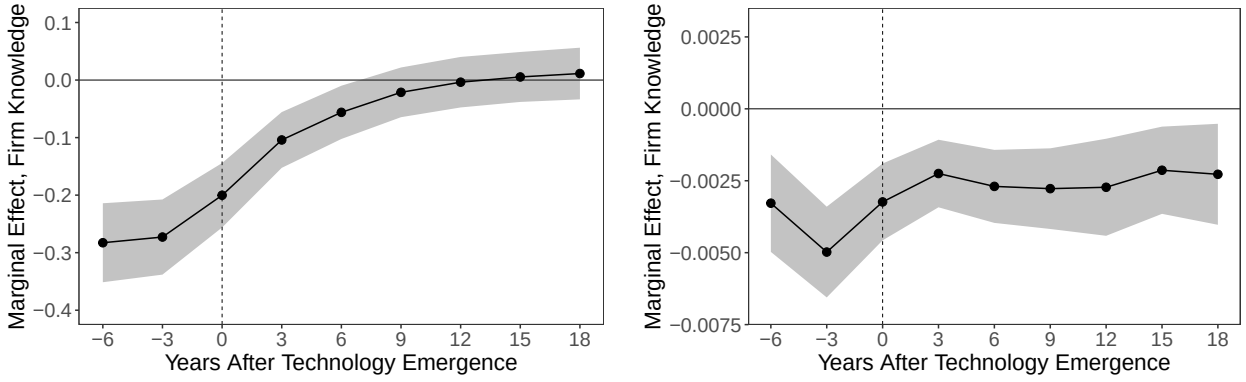

Notes: All regressions include fixed effects by firm-year, technology-window, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. All standard errors are clustered at the firm level.

Figure A.5.6: Technology Patenting Level Difference, Dynamic (No Early Patents)

(a) Patents: $\log(1 + Pat_{i\theta w})$ (b) Patent Share: $TechShare_{i\theta w}$

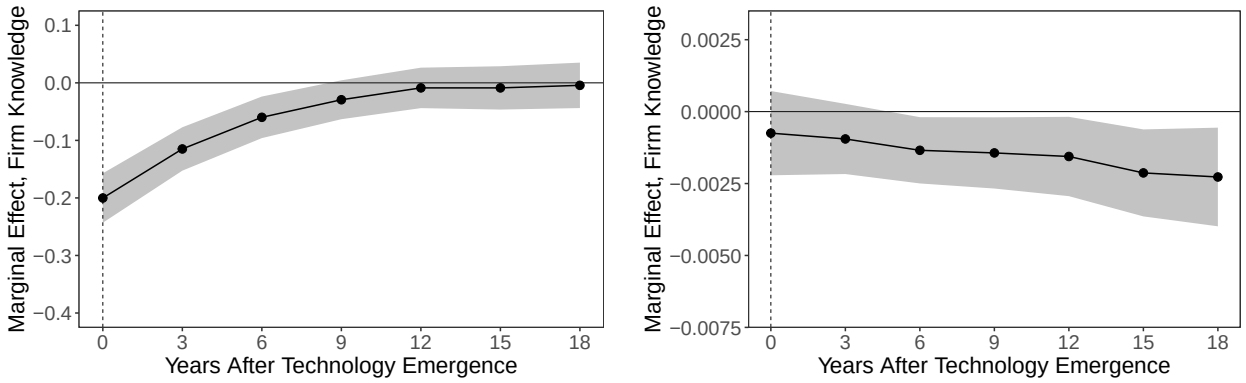

Notes: All regressions include fixed effects by firm-year, technology-window, and the number of years for which the firm has been publicly listed. They also include dummy variables for zero values of each of the knowledge stocks and R&D expenditures s_{it}. All standard errors are clustered at the firm level.

Appendix B

Appendix to Combining Complements: Theory and Evidence from Cancer Treatment Innovation

B.1 Data Construction Details

B.1.1 Cancer Drugs

We use three main sources of data about cancer drugs and characteristics.

GlobalData: We access GlobalData's Drug Database via MIT's institutional subscription.[1] This database is similar to other databases used in papers about pharmaceutical demand and supply, such as Pharmaprojects and Cortellis. We harmonize firm names across different drugs that they own, and we collapse ownership to the parent level (i.e., assign all drugs owned by subsidiaries to their parent company).

Drugs@FDA: We supplement the GlobalData drug database with marketing and other regulatory information from Drugs@FDA.[2] We count the number of generic competitors for each drug in a year as the total number of Abbreviated New Drug Application's (ANDA) filed cumulatively for that drug up to that year. More than one firm can be listed on each ANDA, but each is typically producing a different dosage.

[1]The database is described here: https://www.globaldata.com/marketplace/pharmaceuticals/pipeline-marketed-drugs/. We received an extract of all drugs classified as oncology drugs on October 31, 2023.

[2]The data is available here: https://www.fda.gov/drugs/drug-approvals-and-databases/drugsfda-data-files. We downloaded the full database on October 3, 2023.

SEER CanMED: We obtain HCPCS codes for marketed drugs from the CanMED: HCPCS database.[3] HCPCS codes often change overtime, so finding historical HCPCS codes is important to using historical Medicare and Marketscan data.

B.1.2 Clinical Trials

We obtain all metadata from all clinical trails on ClinicalTrials.gov.[4]

Identifying Cancer Clinical Trials

We use the list of oncology Medical Subject Headings (MeSH) terms and free text keywords in Califf et al. (2012) to subset to oncology clinical trials. We define an oncology clinical trial to be a trial with at least one MeSH term or condition that is in the list from Califf et al. (2012). We then manually make a crosswalk between the list of MeSH terms and cancer types in Chu and DeVita, 2019. The mapping is many-to-one (one MeSH term may map to many cancer types in Chu and DeVita (2019)), and each clinical trial may have more than one MeSH term. Appendix Table B.1.1 summarizes the number of trials we observe for each cancer type in Chu and DeVita (2019), where we count a trial for a specific cancer type if any of the MeSH codes map to that type of cancer.

[3]The data is available here: https://seer.cancer.gov/oncologytoolbox/canmed/hcpcs/. We downloaded the full database on October 4, 2023.

[4]The data is available here: https://classic.clinicaltrials.gov/api/gui/ref/download_all. Our extract is from January 12, 2023.

Table B.1.1: Clinical Trial Counts By Cancer

Cancer	# Trials	Fraction Combo	Fraction Industry
Hodgkin Lymphoma	2784	0.62	0.34
Breast Cancer	2624	0.65	0.31
Acute Lymphoblastic Leukemia	2469	0.6	0.32
Non-Small Cell Lung Cancer	2341	0.64	0.45
Prostate Cancer	1576	0.52	0.32
Brain Cancer	1532	0.48	0.2
Multiple Myeloma	1233	0.69	0.38
Malignant Melanoma	1201	0.64	0.34
Colorectal Cancer	991	0.7	0.31
Head and Neck Cancer	895	0.57	0.24
Pancreatic Cancer	760	0.73	0.26
Ovarian Epithelial Cancer	726	0.65	0.37
Renal Cell Carcinoma	642	0.55	0.41
Kaposi Sarcoma	562	0.53	0.22
Hepatocellular Carcinoma	456	0.49	0.37
Carcinoid and Neuroendocrine Tumors	434	0.55	0.21
Bladder Cancer	352	0.55	0.32
Gastric Cancer	304	0.76	0.33
Myelodysplastic Syndrome	263	0.5	0.41
Endometrial Cancer	258	0.57	0.22
Non-Hodgkin Lymphoma	240	0.64	0.15
Esophageal Cancer	179	0.73	0.18
Cervical Cancer	175	0.51	0.21
Biliary Tract Cancer	155	0.64	0.32
Malignant Mesothelioma	151	0.54	0.27
Thyroid Cancer	145	0.43	0.23
Osteogenic Sarcoma	89	0.43	0.24
Soft Tissue Sarcoma	75	0.49	0.33
Basal Cell Carcinoma	69	0.16	0.51
Gastrointestinal Stromal Tumor	67	0.34	0.55
Merkel Cell Carcinoma	66	0.48	0.47
Acute Myeloid Leukemia	60	0.63	0.4
Ovarian Germ Cell Tumor	59	0.69	0.03
Small Cell Lung Cancer	57	0.7	0.4
Anal Cancer	33	0.64	0.12
Waldenström's Macroglobulinemia	30	0.53	0.1
Chronic Lymphocytic Leukemia	27	0.81	0.37
Thymoma	25	0.44	0.08
Adrenocortical Carcinoma	24	0.42	0.21
Chronic Myeloid Leukemia	21	0.38	0.14
Cancer of Unknown Primary	13	0.85	0.08
Testicular Cancer	7	0.43	0.0
Hairy Cell Leukemia	2	0.5	0.5

Notes: Table shows the number of drug trials run by cancer. Fraction Combo is the fraction of trials for that cancer that are for combinations. Fraction Industry is the fraction of trials for that cancer that have an industry lead sponsor.

Extracting Drug Information from Clinical Trials

We use the OpenAI GPT-4o API to extract drugs used in the control and treatment arm(s) of each trial. The query we use is as follows:

```
request_all = "What drugs are used in the treatment versus control arms of
    the clinical trial described below? We provide the trial title and
    summary below. If a trial does not use drugs, just say 'none' for both
    the treatment and control arms.

You should return a list with 2 elements separated by semicolons: the word '
    treatment:' followed by the drug(s) used in the treatment arms of the
    trial (separated by commas), and the word 'control:' followed by the
    drug(s) (write none if none are used) used in the control arm (
    separated by commas).

Example outputs:
treatment: cisplatin; control: placebo
treatment: oxaliplatin, bevacizumab; control: none
treatment: none; control: none

####
The trial title is: '$(trial_title)'.
The trial brief summary is: '$(trial_brief_summary)'.
The trial full summary is: '$(trial_summary)'.
###"

output_instructions = "Desired output format is: treatment: drug1, drug2...;
     control: drug3, drug4..."

body = JSON.json(
    Dict(
    "model" => "gpt-4o",
    "messages" => [Dict("role" => "user", "content"=> request_all),
                   Dict("role" => "system", "content"=> output_instructions)
                        ],
    "temperature" => 0.0)
    )

url = "https://api.openai.com/v1/chat/completions"
```

Clinical Trial Sample Restrictions

After extracting the drugs used in a trial, we subset to trials that trial at least one drug (rather than only trialing e.g., a radiation or surgical treatment). We further subset to trials with at least one location in the US in order to focus on trials for treatments that are most likely to be used in the US, which our drug demand usage data covers.

Classifying Combination Clinical Trials

We define a combination trial to be a trial that tests at least two drugs together in a treatment arm of the trial. We do not count a single drug plus some other treatment (e.g., radiation or surgery) as a combination trial. We also do not count a trial as a combination trial if the control arm is a combination but the treatment arm is a monotherapy. Figure B.1.1 plots the CDF of the number of component drugs in combination therapies.

Figure B.1.1: Combination Trials by # of Drugs

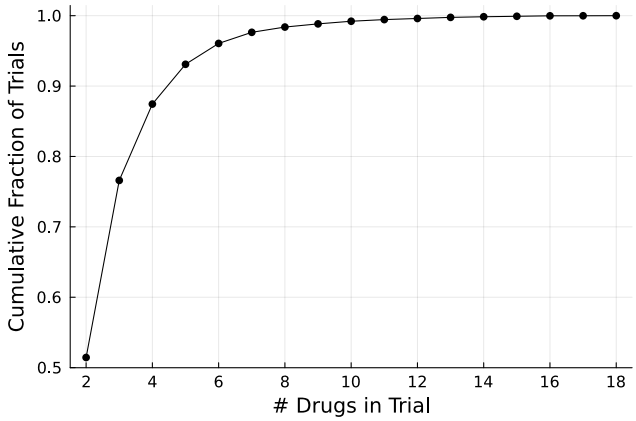

Notes: Figure shows the CDF of the number of drugs involved in combination trials.

Merging Clinical Trial Drugs with GlobalData and Drugs@FDA

Given the extracted drug names from the trial, we then merge these drug names with the drug-level dataset we create from GlobalData and Drugs@FDA. Importantly, GlobalData provides an extensive list of aliases for each drug, include names used in early-stage clinical trials. We check all possible aliases when merging. Table B.1.2 shows summary statistics for merging drugs reported in clinical trials with drugs reported in GlobalData.

Table B.1.2: Clinical Trial GlobalData Drug Merge Rates by # Drugs in Trial

# Drugs	# Trials	Avg Frac Drugs Merged	Frac Merged All Drugs
1 Drug	11959	0.8	0.8
2 Drugs	8537	0.89	0.82
3 Drugs	4175	0.88	0.77
4 Drugs	1800	0.85	0.68
5 Drugs	938	0.86	0.64
All	28555	0.85	-

Notes: Table shows summary statistics of the merge rates between drugs in the clinical trial data and drugs in GlobalData. Each row (except for the last) presents summary statistics subsetting to clinical trials that have a certain number of drugs. # Trials gives how many trials satisfy that condition (e.g., have a certain number of drugs), Avg Frac Drugs Merged is the average (over trials) fraction of drugs in a trial that merge to GlobalData, and Frac Merged All Drugs is the fraction of trials for which *all* drugs in a trial merge.

Trials per Drug

Figure B.1.2: Number of Trials Involving a Drug

(a) Approved Drugs

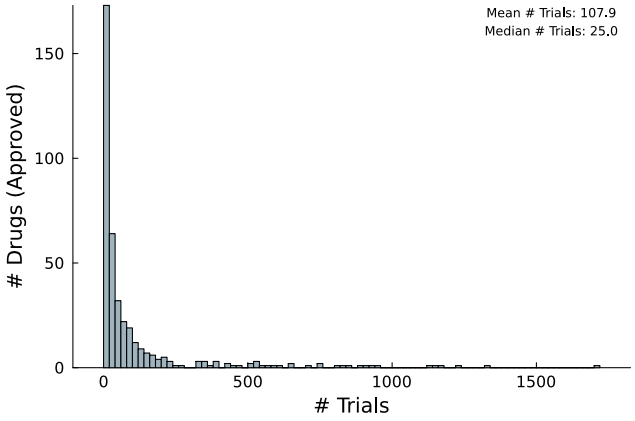

(b) Not Approved Drugs

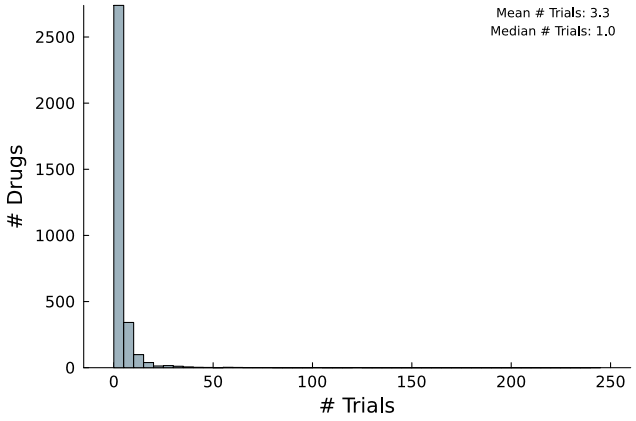

Notes: Figure shows the distribution of the number of trials that is a drug is involved in (in the treatment arm), separately for approved and not approved drugs.

Figure B.1.3: Fraction of Trials for Combinations

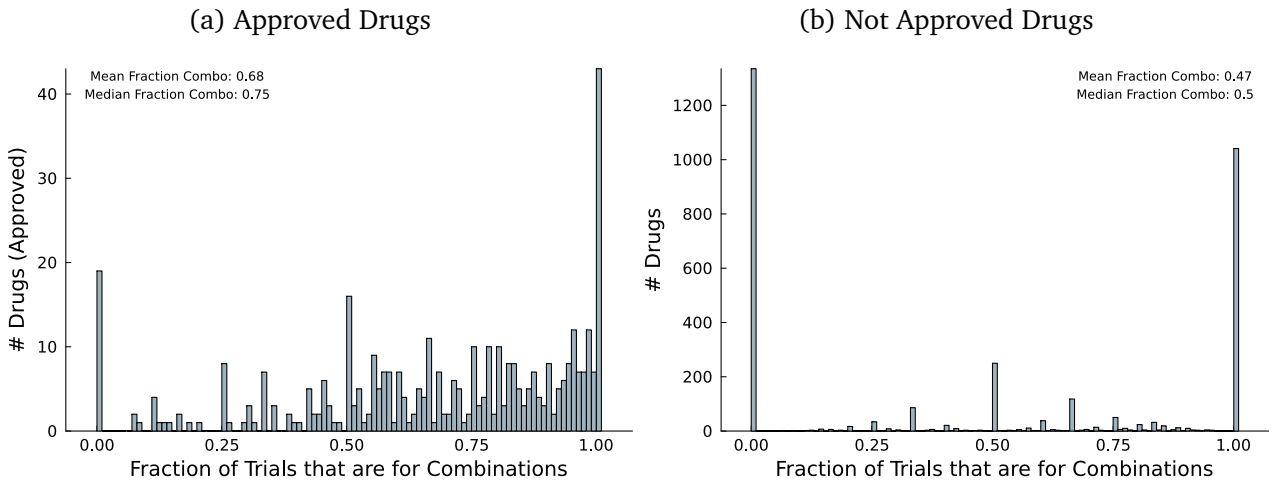

Notes: Figure shows the distribution of the fraction of trials that is a drug is involved in that are for combinations (in the treatment arm), separately for approved and not approved drugs.

Trials for New Uses

Figure B.1.4: Clinical Trials by Funding Type

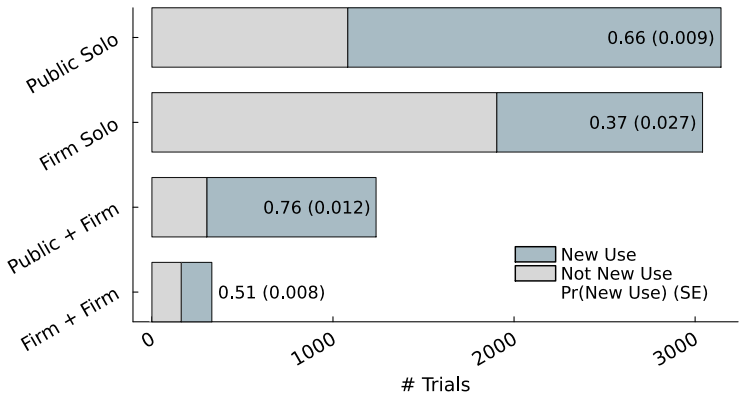

Notes: Figure shows the counts of fractions of single drug trials for the initial condition versus new conditions, separately by different trial sponsors.

B.1.3 Cancer Regimens

Figure B.1.5 shows an example of how regimens are recorded in Chu and DeVita (2019). Table B.1.3 summarizes the number of regimens available for each cancer, the number of drugs used in each regimen, and the publication year of related scientific papers.

Figure B.1.5: Chu and DeVita (2019) example

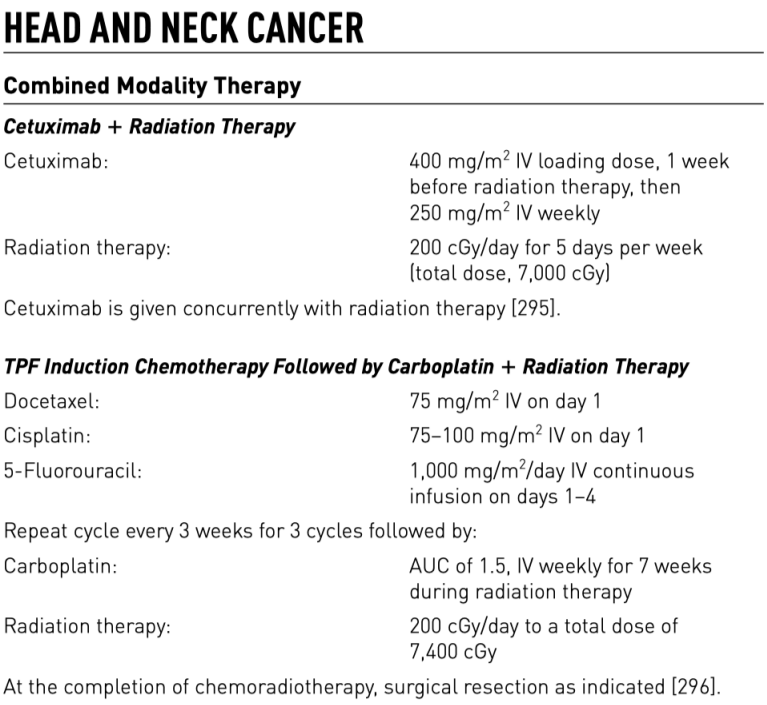

HEAD AND NECK CANCER

Combined Modality Therapy

Cetuximab + Radiation Therapy

Cetuximab:	400 mg/m^2 IV loading dose, 1 week before radiation therapy, then 250 mg/m^2 IV weekly
Radiation therapy:	200 cGy/day for 5 days per week (total dose, 7,000 cGy)

Cetuximab is given concurrently with radiation therapy [295].

TPF Induction Chemotherapy Followed by Carboplatin + Radiation Therapy

Docetaxel:	75 mg/m^2 IV on day 1
Cisplatin:	75–100 mg/m^2 IV on day 1
5-Fluorouracil:	1,000 mg/m^2/day IV continuous infusion on days 1–4

Repeat cycle every 3 weeks for 3 cycles followed by:

Carboplatin:	AUC of 1.5, IV weekly for 7 weeks during radiation therapy
Radiation therapy:	200 cGy/day to a total dose of 7,400 cGy

At the completion of chemoradiotherapy, surgical resection as indicated [296].

Notes: Figure shows example example treatment regimens for head and neck cancer, taken from Chu and DeVita (2019).

Additional Summary Statistics

Table B.1.3: Cancer Regimen Summary

Cancer	# Regimens	Mean # Drugs	Min # Drugs	Max # Drugs	Unique # Drugs	Average Paper Year
adrenocortical	2	3.0	2	4	5	2012.0
anal	5	1.8	1	2	6	2007.6
basal cell	2	1.0	1	1	2	2013.5
biliary tract	9	1.67	1	2	6	2004.56
bladder	20	1.9	1	4	18	2004.3
brain	10	1.7	1	3	8	2003.2
breast	70	2.04	1	5	35	2003.99
unknown primary	8	2.5	2	3	11	2001.25
carcinoid and neuroendocrine	10	1.5	1	2	12	2004.0
cervical	21	1.9	1	4	15	2003.19
colorectal	39	2.51	1	5	20	2006.33
endometrial	15	1.93	1	4	14	2001.33
esophageal	12	2.25	1	3	9	2004.75
gastric	22	2.32	1	3	13	2004.82
gastrointestinal stromal tumor	5	1.0	1	1	5	2008.4
head and neck	20	2.05	1	4	13	1999.85
hepatocellular	10	1.4	1	3	12	2006.4
kaposi sarcoma	9	1.33	1	3	11	1995.44
all	17	2.76	1	6	22	1999.12
aml	21	1.67	1	3	19	2002.67
cll	23	1.61	1	3	16	2001.65
cml	10	1.1	1	2	10	2002.9
hairy cell leukemia	3	1.0	1	1	3	1989.33
non-small cell lung	51	1.86	1	3	28	2006.76
small cell lung	12	1.92	1	3	10	1997.0
hodgkin lymphoma	18	3.06	1	8	24	2001.17
non-hodgkin lymphoma	44	3.39	1	11	38	2000.82
primary cns lymphoma	5	2.6	1	6	8	2008.0
malignant melanoma	18	1.44	1	3	17	2005.56
malignant mesothelioma	11	1.73	1	3	8	2004.82
merkel cell	4	2.0	1	3	7	1999.0
multiple myeloma	27	2.26	1	4	16	2005.63
myelodysplastic syndrom	5	1.2	1	2	6	2007.0
osteogenic sarcoma	4	3.25	2	4	7	1997.5
ovarian epithelial	26	1.5	1	3	18	2003.0
ovarian germ cell	1	3.0	3	3	3	2004.0
pancreatic	13	2.08	1	4	10	2003.0
prostate	27	1.48	1	3	21	1999.56
renal cell	16	1.31	1	2	14	2007.38
soft tissue sarcoma	15	2.07	1	5	15	2004.2
testicular	12	2.92	1	5	12	1993.67
thymoma	6	2.83	2	4	11	2000.33
thyroid	8	1.12	1	2	8	2004.25
waldenstrom's macroglobulinemia	5	2.6	1	3	6	2007.0
Total	691	1.99	1.13636	3.52273	193	2003.19

Notes: Table shows summary statistics of recommended regimens in Chu and DeVita (2019). We drop regimens that are missing paper publication years. # regimens is the number of recommend regimens for that cancer. Mean # Drugs is the average number of drugs in a regimen for that cancer. Min # Drugs and Max # Drugs are the minimum and maximum number of drugs in a regimen for that cancer, respectively. Unique # Drugs is the total number of unique drugs used in regimens for that cancer. Average Paper Year is the average paper publication year for the pivotal study for regimens for that cancer.

Figure B.1.6: Number of Firms, Drugs Per Regimen in Chu and DeVita (2019)

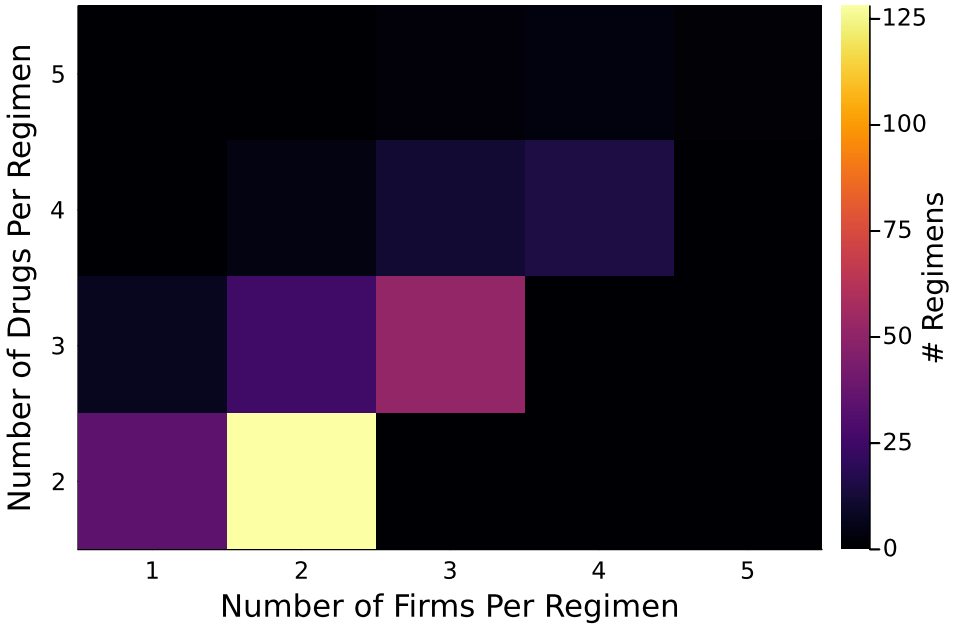

Notes: Figure shows the number of firms and drugs per regimen in Chu and DeVita (2019).

Figure B.1.7: Fraction of Drug Market Share from Combinations

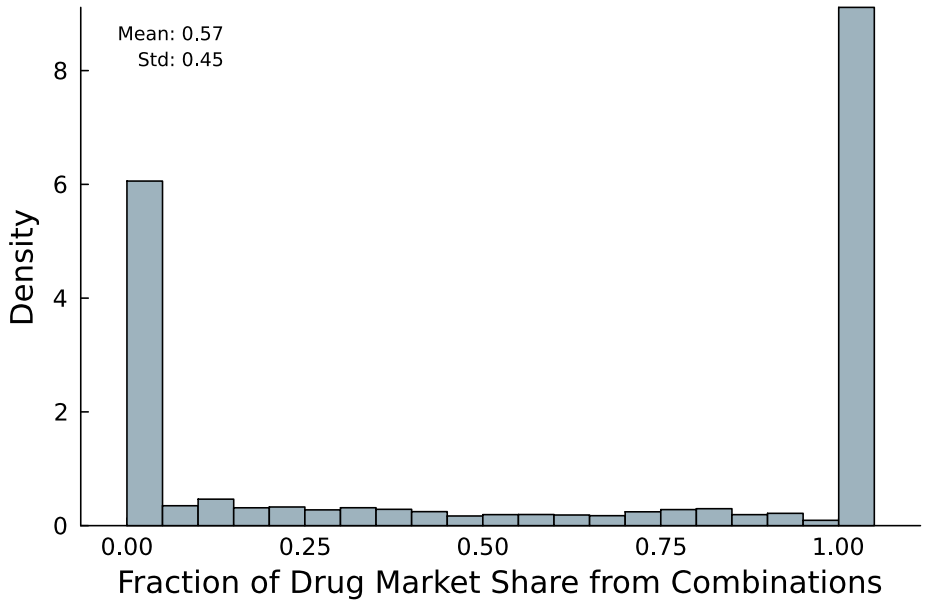

Notes: Figure shows the density of the fraction of a drug's market share (within time, cancer tuples) that comes from combination regimens. A drug's market share is calculated as the sum of market shares of regimens the drug is a component of.

Figure B.1.8: Observed versus Constructed Regimen Price

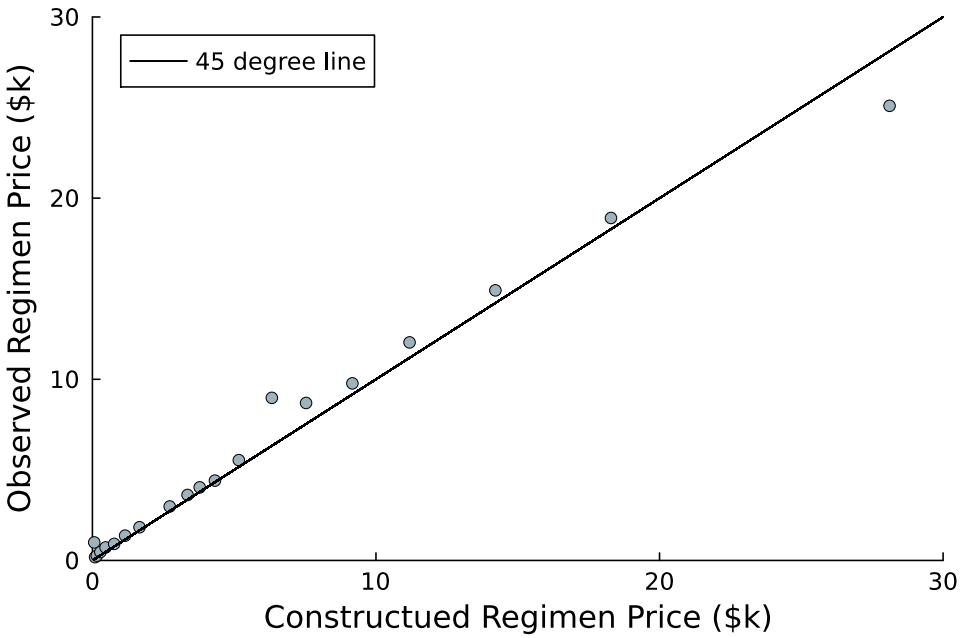

Notes: Figure shows a binscatter of the relationship between observed regimen prices and constructed regimen prices. Constructed regimen prices are calculated from average drug prices then summed to the regimen level, while observed prices sum over drugs and then average. This figure demonstrates the lack of price discrimination based on use in a particular regimen or not.

Table B.1.4: Observed versus Constructed Regimen Price

	Constructed Regimen Price
(Intercept)	-0.050
	(0.058)
Observed Regimen Price	0.953
	(0.006)
N	3,550
R^2	0.875

Notes: Table shows the relationship between observed regimen prices and constructed regimen prices. Constructed regimen prices are calculated from average drug prices then summed to the regimen level, while observed prices sum over drugs and then average. This table demonstrates the lack of price discrimination based on use in a particular regimen or not.

B.1.4 Claims Data

This section describes how we use microdata from CMS and Marketscan to construct cancer patient treatment regimen usage and prices.

Populations

Figure B.1.9 displays counts of patients in the Medicare and Marketscan data. The solid line and dashed line show the total number of beneficiaries enrolled in Traditional Medicare and Medicare Advantage, respectively, calculated using the Master Beneficiary Summary File. We use a 20% random sample of the beneficiaries, where we observe claims for patients enrolled in Traditional Medicare. The dash-dotted line shows the number of patients we observe in the Marketscan data, calculated using the Annual Summary Enrollment file.

In order to make our demand system representative of the US population, we scale the counts of Traditional Medicare patients (and their drug usage etc.) to meet match the total number of Traditional Medicare and Medicare Advantage patients. We scale up the counts of Marketscan patients to meet the residual US population.

Figure B.1.9: Number of Enrollees by Insurance Type

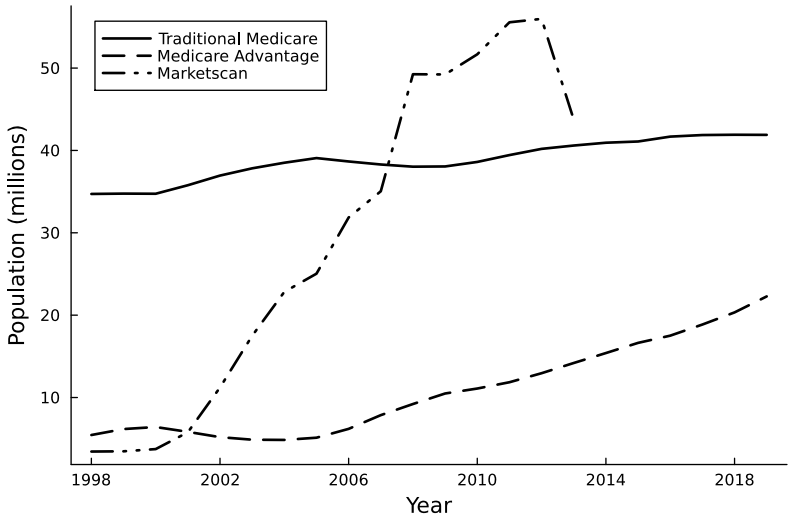

Notes: Figure shows the number of individuals enrolled in Traditional Medicare and Medicare Advantage, and the number of individuals we observe in the Marketscan data.

Identifying Cancer Patients

To identify cancer patients, we use the 42 types of cancer identified by Chu and DeVita (2019). In the Medicare and Marketscan data, patients are assigned ICD 10 codes after 2015, and ICD 9 before 2015, to indicate diagnoses. We consider codes C00-D48 (neoplasms), and in particular codes C00-D09 (malignant and in situ neoplasms), leaving out codes D10-D36 (benign neoplasms) and D37-D48 (neoplasms of uncertain or unknown behavior). We manually assign each code to a set of relevant Chu and DeVita (2019) cancer types. The mapping is one to

many (e.g., one ICD 10 code can map to many Chu and DeVita (2019) cancer types). In these cases, the diagnosis codes do not allow us to distinguish between different types of cancers. An important example of this is lung cancer, where the diagnosis codes do not distinguish between non-small and small-cell lung cancer.

To convert between ICD 9 and ICD10 codes, we use the NBER ICD9 to ICD10 crosswalk.[5]

We use the 20% sample of Medicare beneficiaries for years 1998-2019. A beneficiary is defined to have a type of cancer in a year if a claim is submitted with a primary diagnosis of that cancer in either the inpatient, outpatient, or carrier files. We use the 100% sample of the Marketscan dataset. A beneficiary is defined to have a type of cancer in a year if a claim is submitted with a primary diagnosis code of that cancer in either the outpatient, inpatient services, or inpatient admissions files.

Table B.1.5 shows counts of patients by year for the Medicare and Marketscan datasets.

[5]Available here: https://www.nber.org/research/data/icd-9-cm-and-icd-10-cm-and-icd-10-pcs-crosswalk-or-general-equivalence-mappings. These files were downloaded on January 21, 2023.

Table B.1.5: Patient Counts by Cancer

Cancer	Median Patients Per Year	
	Medicare	Marketscan
Prostate	205,334	47,097
Head and Neck	110,665	26,738
Colorectal	71,295	27,008
Bladder	54,078	13,270
Basal Cell Carcinoma	53,614	118,782
Non-Hodgkin Lymphoma (NHL)	46,927	26,732
Breast	43,554	41,182
Small Cell Lung (SCLC)	29,755	11,404
Non-Small Cell Lung (NSCLC)	29,755	11,404
Osteogenic Sarcoma	26,327	12,208
Malignant Melanoma	20,618	17,930
Renal Cell Carcinoma	19,976	10,119
Chronic Lymphocytic Leukemia (CLL)	19,397	4,942
Hepatocellular Carcinoma	18,718	7,697
Malignant Mesothelioma	18,625	11,072
Anal	18,467	8,244
Endometrial	17,299	9,243
Multiple Myeloma	16,274	4,920
Biliary Tract	16,128	7,077
Brain	14,552	10,754
Hodgkin Lymphoma	14,369	13,088
Ovarian Germ Cell Tumor	12,410	10,616
Ovarian Epithelial	12,410	10,616
Pancreatic	11,338	3,597
Gastric	7,821	2,993
Central Nervous System (CN)	5,548	3,526
Chronic Myeloid Leukemia (CML)	5,496	2,789
Acute Myeloid Leukemia (AML)	5,092	3,414
Acute Lymphoblastic Leukemia (ALL)	4,965	3,990
Esophageal	3,976	1,577
Cervical	3,059	13,833
Hairy Cell Leukemia	2,995	1,210
Waldenstrom Macroglobulinemia (WM)	2,133	576
Testicular	1,158	5,308
Adrenocortical Carcinoma	1,094	1,492
Kaposi Sarcoma	713	610
Thyroid	661	822
Thymoma	456	321
Merkel Cell Carcinoma	451	293

Notes: Table shows counts of patients by cancer, taking a median over years.

Calculating Regimen Market Shares

Table B.1.6 shows the average number of drugs cancer patients take that are not included in the regimen they are assigned to.

Table B.1.6: Regimen Procedure – Extra Drugs

Cancer	Extra Drugs
Prostate	0.18
Head and Neck	0.36
Colorectal	0.72
Bladder	0.11
Non-Hodgkin Lymphoma (NHL)	1.14
Breast	0.23
Non-Small Cell Lung (NSCLC)	0.12
Chronic Lymphocytic Leukemia (CLL)	0.13
Hepatocellular Carcinoma	0.36
Malignant Mesothelioma	0.15
Endometrial	0.39
Multiple Myeloma	0.33
Biliary Tract	0.1
Brain	0.15
Hodgkin Lymphoma	0.5
Ovarian Epithelial	0.18
Pancreatic	0.16
Gastric	0.28
Esophageal	0.13
Mean	0.3

Notes: Table shows the average number of "extra" drugs patients take in addition to drugs in their assigned regimens, by cancer.

Outside Option

Figure B.1.10 shows the fraction of patients identified to have a particular type of cancer that take cancer drugs in a given year. We take the mean over the cancers, and plot this trend separately for Medicare and Marketscan patients.

Figure B.1.10: Fraction of Patients Taking Cancer Drugs

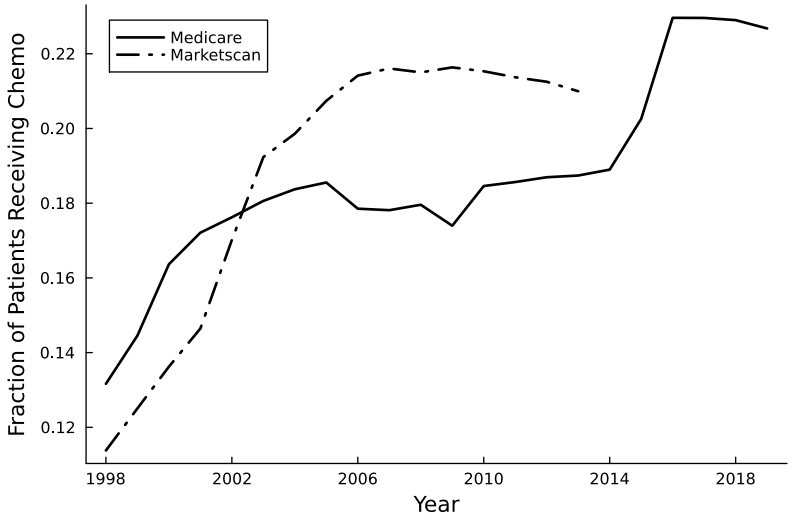

Notes: Figure shows the fraction of cancer patients taking drugs. We take a mean over cancers.

Patient Regimen Summary Statistics

Table B.1.7 presents summary statistics at the patient level about patient characteristics and regimen usage, separately for each cancer in our estimation sample.

Table B.1.7: Regimen - Patient Summary Statistics

Cancer	Mean Age	Total Regimen Days (Mean)	Total Regimen Days (Std)	Regimen Days (Mean)	Regimen Days (Std)	Num Regimens (Mean)	Num Regimens (Std)
Biliary Tract	92.6	133	91	120	81	1.1	0.3
Bladder	91.6	170	138	126	95	1.3	0.7
Brain	81.2	184	181	167	165	1.1	0.4
Breast	89.5	355	365	213	261	1.7	1.1
Chronic Lymphocytic Leukemia	88.8	282	325	182	218	1.5	0.9
Colorectal	77.6	266	242	152	128	1.8	1.2
Endometrial	91.0	198	189	160	154	1.2	0.5
Esophageal	88.3	131	90	113	68	1.2	0.4
Gastric	85.4	168	134	128	94	1.3	0.6
Head and Neck	78.3	150	139	117	100	1.3	0.6
Hepatocellular	89.5	145	155	128	136	1.1	0.4
Hodgkin Lymphoma	85.2	153	123	143	103	1.1	0.3
Malignant Mesothelioma	85.6	146	113	117	88	1.2	0.6
Multiple Myeloma	77.5	469	541	226	289	2.1	1.5
Non-Hodgkin Lymphoma	77.8	261	265	201	215	1.3	0.6
Non-Small Cell Lung	76.3	199	195	128	119	1.6	0.9
Ovarian Epithelial	81.4	346	315	154	129	2.2	1.5
Pancreatic	77.7	201	172	141	119	1.4	0.8
Prostate	78.0	430	530	287	376	1.5	0.9

Notes: Table includes summary statistics for patients about demographic characteristics and regimen usage by cancer.

B.2 Facts

B.2.1 Who Funds Combination Innovation

Figure B.2.1: Clinical Trials by Funding Type

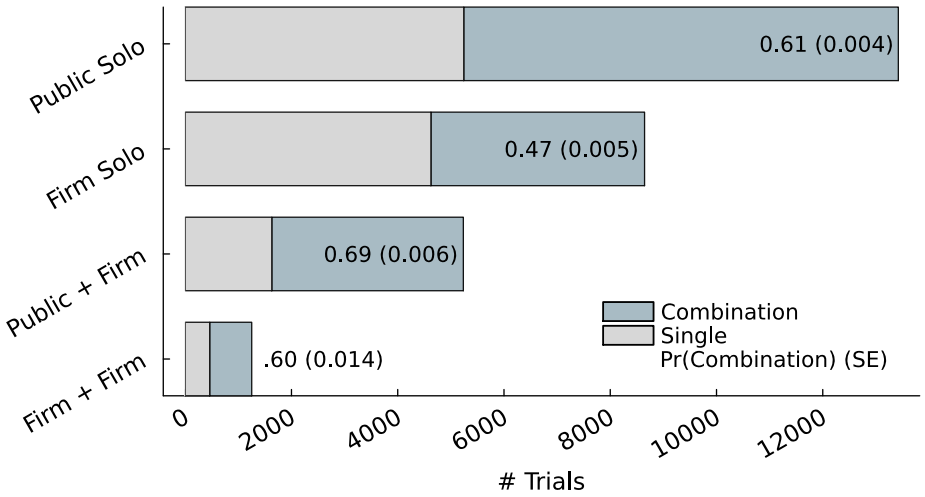

Notes: Figure shows the number of trials run by different sponsors broken in combination (blue) and single drug trials (gray) and the fraction of that sponsor's trials that are for combinations (black number). Robust standard errors are in parentheses on each bar, which come from a trial-level regression of an indicator for being a combination trial on sponsor type.

Table B.2.1: Probability of Combination Trialing

	Combination Indicator			
	(1)	(2)	(3)	(4)
Sponsor Type: Firm + Firm	0.632	0.651	0.013	0.046
	(0.014)	(0.014)	(0.014)	(0.040)
Sponsor Type: Firm Solo	0.465	0.479	-0.139	-0.074
	(0.005)	(0.006)	(0.007)	(0.032)
Sponsor Type: Public + Firm	0.689	0.685	0.074	0.089
	(0.006)	(0.007)	(0.008)	(0.050)
Sponsor Type: Public Solo	0.609	0.600		
	(0.004)	(0.005)		
Cancer Type Fixed Effects			Yes	Yes
Trial Submission Year Fixed Effects			Yes	Yes
N	28,555	21,807	28,555	26,485
R^2	0.593	0.590	0.060	0.333
Within-R^2			0.025	0.013
Year Restriction	None	Post 2007	None	None
Weights	None	None	None	Enrollment

Notes: Table shows estimates of a trial-level regression of an indicator for being a combination trial on indicators for different multi-sponsor types. The first column includes the full sample of clinical trials, while the second checks robustness with respect to only including trials run after 2007. The third column additionally conditions on cancer type (the first type of cancer reported in the trial) and trial submission year. We choose the "Public Solo" group as the base. The fourth column additional weights be number of enrolled patients (e.g., trial size). Robust standard errors are in parentheses.

B.2.2 What is the Direction of Combination Innovation

Levels

Figure B.2.2: Two-Drug Trials by Type

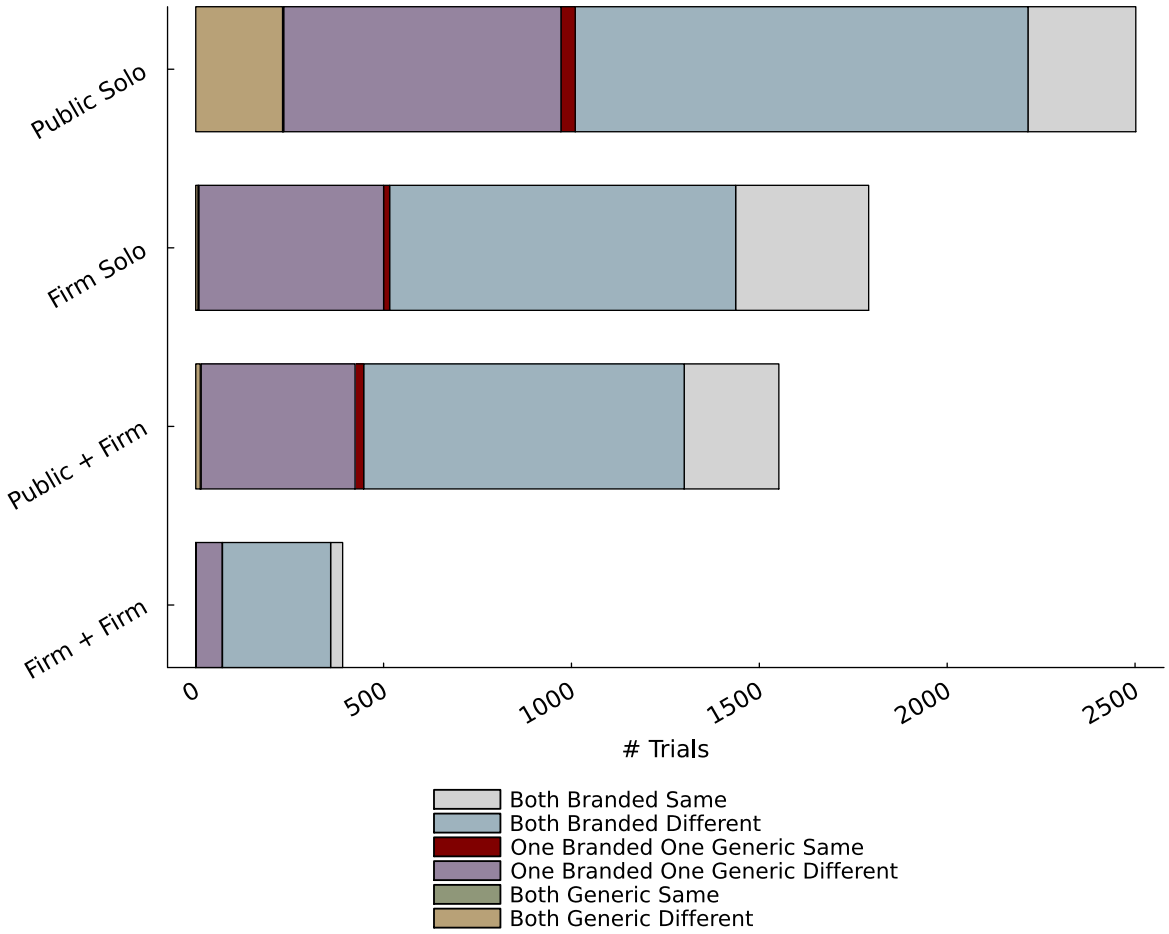

Baseline Regression Tables

Table B.2.2: Relative Probability of Trialing 2 Drugs Together by Drug Ownership Status – Private

	Trial Run (Normalized), Private
Both Branded, Same	10.86546
	(0.938387)
Both Branded, Diff	0.46238
	(0.019609)
Brand and Generic	1.02547
	(0.051401)
N	33,356,609
R^2	0.000094

Notes: Table shows estimates of Equation (2.3). Robust standard errors are in parentheses.

Table B.2.3: Relative Probability of Trialing 2 Drugs Together by Drug Ownership Status – Public

	Trial Run (Normalized), Public
Both Branded	0.61313
	(0.015553)
One Branded, One Generic	1.34790
	(0.042284)
Both Generic	5.05438
	(0.300375)
N	33,356,609
R^2	0.000132

Notes: Table shows estimates of Equation (2.3) for public innovators. Robust standard errors are in parentheses.

Figure B.2.3: Relative Probability of Trialing 2 Drugs Together by Drug Ownership Status

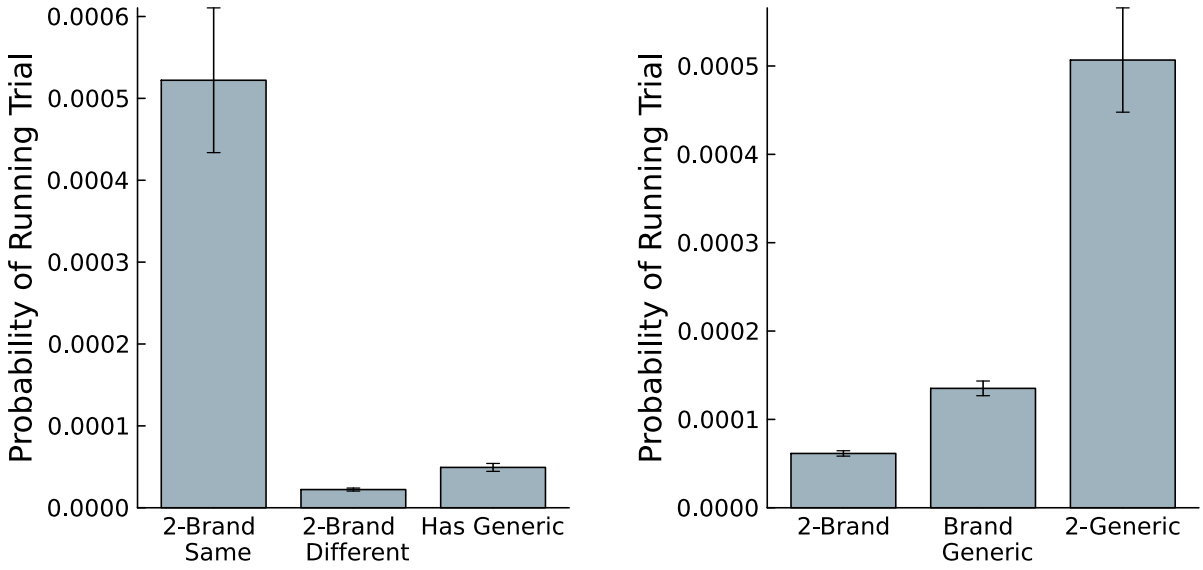

(a) Firm

(b) Public

Notes: Figure shows estimated γ_k coefficients of Equation (2.3) where we do not normalize the trialing indicator. Private innovators (firms) are on the left panel, and public innovators on the right. 95% confidence intervals for the regression coefficients calculated from robust standard errors are displayed on each bar. Regression tables and additional robustness checks are given in Appendix B.2.

Baseline Regression Tables, Not Normalized

Extended Regression Results

$$\begin{aligned}\frac{\text{Trial Now}_{rft}}{\beta^{\text{unif}}\left(L_{ft}\right)} &= \gamma_1 \,\text{Both Branded Same}_{rft} + \gamma_2 \,\text{Both Branded Different}_{rft} \\ &\quad + \gamma_3 \,\text{One Branded One Generic Same}_{rft} + \gamma_4 \,\text{One Branded One Generic Different}_{rft} \\ &\quad + \gamma_5 \,\text{Both Generic Same}_{rft} + \gamma_6 \,\text{Both Generic Different}_{rft} + \varepsilon_{rft}. \qquad \text{(B.1)}\end{aligned}$$

Table B.2.4: Relative Probability of Trialing 2 Drugs Together by Drug Ownership Status – Private

	Trial Run (Normalized), Private
Both Branded, Same	10.86546
	(0.938387)
Both Branded, Diff	0.46238
	(0.019609)
Branded and Generic, Same	1.76066
	(0.556746)
Branded and Generic, Diff	1.00675
	(0.053207)
Both Generic, Same	0.80991
	(0.809899)
Both Generic, Diff	1.13274
	(0.210338)
N	33,356,609
R^2	0.000094

Notes: Table shows estimates of Equation (B.1) for private innovators. Robust standard errors are in parentheses.

Table B.2.5: Relative Probability of Trialing 2 Drugs Together by Drug Ownership Status – Public

	Trial Run (Normalized), Public
Both Branded, Same	7.69509
	(0.546655)
Both Branded, Diff	0.54049
	(0.014677)
Branded and Generic, Same	4.81012
	(0.636962)
Branded and Generic, Diff	1.29260
	(0.041737)
Both Generic, Same	4.65828
	(1.344416)
Both Generic, Diff	5.07348
	(0.308113)
N	33,356,609
R^2	0.000183

Notes: Table shows estimates of Equation (B.1) for public innovators. Robust standard errors are in parentheses.

Intra-Firm Complementarity

Using data from the NCI-ALAMANC, A Large Matrix of Anti-Neoplastic Agent Combinations, from Holbeck et al. (2017) we can compare whether two-drug combinations that are owned by the same firm have the same efficacy as measured by laboratory tests on human tumor cell lines (NCI-60) of all pairwise combinations of a large set of marketed cancer drugs ($\approx$ 100 drugs, leading to $\approx$ 5000 combinations) for various dosages.

We first look at the effect of being owned by the same firm on tumor growth rates. The results are in Table B.2.6. While the coefficient on the indicator of being owned by the same firm is negative (indicating the tumor decreasing in size), it is relatively small in magnitude compared to the mean growth rate, and reduces in magnitude even further when including drug fixed effects. Figure B.2.4 shows the distribution of combination score (a measure of how effective the combination is overall and relative to single-agent therapies, and find similar results), separately for combinations consisting of drugs owned by the firm or not.

Table B.2.6: NCI Almanac Tumor Growth Rates

	Percent Growth		
	(1)	(2)	(3)
(Intercept)	72.771		
	(0.028)		
Same Firm	-1.300	-1.301	-0.667
	(0.166)	(0.163)	(0.152)
Cell and Panel Fixed Effects Fixed Effects		Yes	Yes
Drug 1 Fixed Effects			Yes
Drug 2 Fixed Effects			Yes
N	2,578,833	2,578,833	2,578,833
F	61.189	63.811	19.383
Within-R^2		0.000	0.000
Adjusted R^2	0.000	0.057	0.235

Notes: Table shows estimates of a regression of tumor growth rate (percentage) of a specific combination on an indicator of whether the combination consists of two drugs owned by the same firm.

Table B.2.7: NCI Almanac ComboScore

	ComboScore		
	(1)	(2)	(3)
(Intercept)	-2.438		
	(0.008)		
Same Firm	0.175	0.174	-0.166
	(0.042)	(0.042)	(0.043)
Cell and Panel Fixed Effects Fixed Effects		Yes	Yes
Drug 1 Fixed Effects			Yes
Drug 2 Fixed Effects			Yes
N	2,578,833	2,578,833	2,578,833
F	17.387	17.385	14.801
Within-R^2		0.000	0.000
Adjusted R^2	0.000	0.005	0.041

Notes: Table shows estimates of a regression of combination score (higher scores are better) of a specific combination on an indicator of whether the combination consists of two drugs owned by the same firm.

Figure B.2.4: Distribution of ComboScore

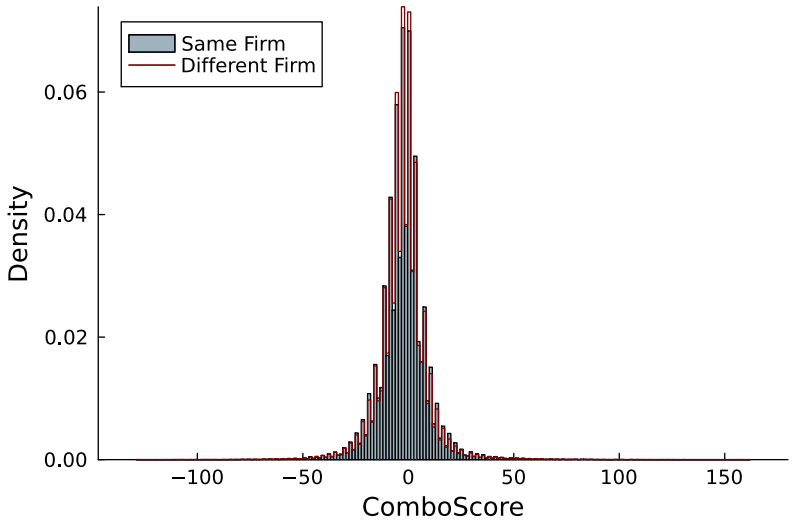

Notes: Figure shows the distribution of combination scores, separately for combinations consisting of drugs owned by the firm or not. Each distribution is normalized to show a density.

Multi-Brand Combination Trials

For two-drug combination trials consisting of branded drugs owned by different firms, the firm that has the drug furthest from patent expiry is significantly more like to run the clinical trial

(64% of trials) than the firm with the "newer" drug (36% of trials).

B.2.3 When Does Combination Innovation Occur

We present results from an alternative specification of Equation 2.4 where we omit drug fixed effects.

Figure B.2.5: Combination Trials and Generic Entry

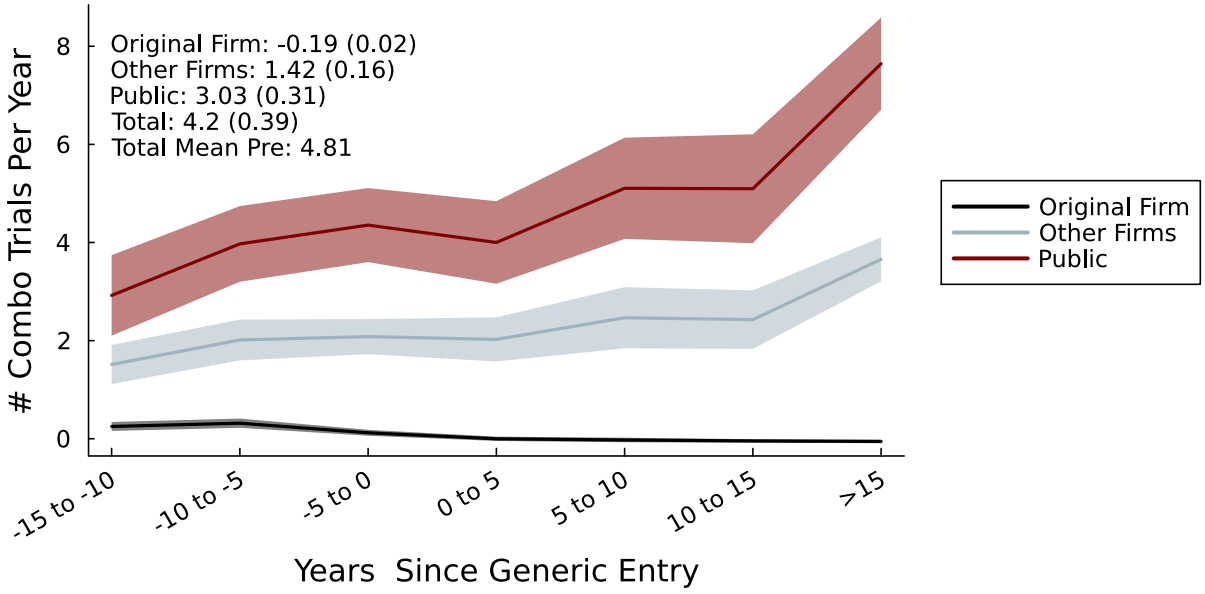

Notes: Figure includes dynamic estimates of Equation 2.4 without drug fixed effects.

Table B.2.8: Combination Trials and Generic Entry Regression

	Original Firm	Public	Other Firms	Total Trials
	(1)	(2)	(3)	(4)
Generic Indicator	-0.172	3.207	1.505	4.291
	(0.020)	(0.277)	(0.142)	(0.350)
Year Fixed Effects	Yes	Yes	Yes	Yes
N	3,882	3,882	3,882	3,882
Mean Pre	0.229	3.894	1.462	4.815

Notes: Table shows coefficient on Generic Indicator in Equation (2.4), excluding drug fixed effects. # Trials (Total) is the total number of combination trials a drug is used in within a year. # Trials Firm is the number of combination trials run by the original owner of the drug in a year. # Trials Other Firm is the number of combination trials run by other firms (i.e., not the original owner of the drug) in a year. Generic Indicator is an indicator of whether drug d is generic at time t (i.e., has had generic entry). Robust standard errors are in parentheses.

B.3 Demand

B.3.1 Instrument Details

Table B.3.1: Instrument First Stage – Medicare

	Regimen Price ($k)
# Entrants	-0.27509
	(0.006316)
Generic Entry Time	-0.15212
	(0.010108)
Regimen Fixed Effects	Yes
Year Fixed Effects	Yes
N	34,751
F	1488.304852
First-stage F statistic	
Within-R^2	0.079509
Adjusted R^2	0.893844

Notes: Table shows estimates of the instrument first-stage (the relationship between regimen price and the instrument) for Medicare patients.

Table B.3.2: Instrument First Stage – Marketscan

	Regimen Price ($k)
# Entrants	-0.27509
	(0.006316)
Generic Entry Time	-0.15212
	(0.010108)
Regimen Fixed Effects	Yes
Year Fixed Effects	Yes
N	34,751
F	1488.304852
First-stage F statistic	
Within-R^2	0.079509
Adjusted R^2	0.893844

Notes: Table shows estimates of the instrument first-stage (the relationship between regimen price and the instrument) for Marketscan patients.

B.3.2 Micro Moments

We calculate micro moments in the data as follows. Let $\hat{Z}$ denote the set of demographic bins. For each $\hat{z} \in \hat{Z}$ calculate the average characteristic of the taken inside regimens as

$$\bar{p}_{\hat{z}t} = \frac{\sum_{r \in \mathcal{R}_t} s_{\hat{z}rt} p_{rt}}{\sum_{r' \in \mathcal{R}_t} s_{\hat{z}r't}} \tag{B.2}$$

$$\overline{\text{fraction biologic}}_{\hat{z}t} = \frac{\sum_{r \in \mathcal{R}_t} s_{\hat{z}rt} \mathbf{1}_{r \text{ biologic}}}{\sum_{r' \in \mathcal{R}_t} s_{\hat{z}r't}} \tag{B.3}$$

$$\overline{\text{fraction combo}}_{\hat{z}t} = \frac{\sum_{r \in \mathcal{R}_t} s_{\hat{z}rt} \mathbf{1}_{r \text{ combo}}}{\sum_{r' \in \mathcal{R}_t} s_{\hat{z}r't}} \tag{B.4}$$

We then calculate the covariance between $\hat{z}$ and and each of these average characteristics (weighted by the number of patients in that bin).

In the model, these moments are calculated similarly, but instead of observing $s_{\hat{z}rt}$ directly, we impute $s_{\hat{z}rt}$ from $s_{r\hat{t}t}$ and the integration weights of the demographic bins. These quantities are simple to compute in the model.

Micro moment targets are show in Table B.3.3. Additional patterns are displayed below.

Table B.3.3: Micro Moment Targets

Micro Moment	**Medicare**	**Medicare Dual**	**Marketscan**
$\text{Cov}(a_{it}, p_{rt})$	-0.0157	-0.0157	0.0217
$\text{Cov}(a_{it}, \mathbf{1}_{r \text{ combo}})$	-0.0076	-0.0076	0.0008
$\text{Cov}(a_{it}, \mathbf{1}_{r \text{ biologic}})$	-0.004	-0.004	0.0004

Notes: Table shows micro moment targets by insurance type. Prices are in thousands of dollars and age is divided by 100. These patterns are shown in detail in Appendix B.3.

Figure B.3.1: Regimen Price and Patient Age – Medicare

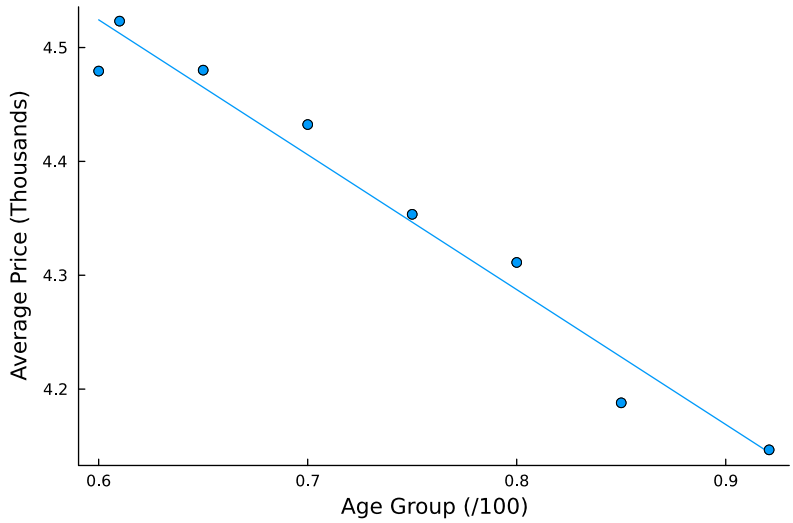

Notes: Figure shows a binscatter of regimen price on age for Medicare patients. Dataset is (share, age-group) monthly average characteristics (conditional on inside good).

Figure B.3.2: Regimen Price and Patient Age – Marketscan

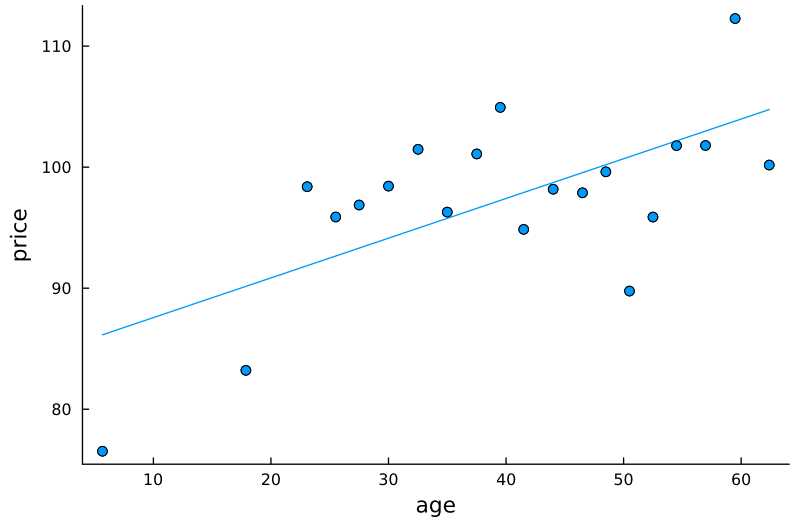

Notes: Figure shows a binscatter of regimen price on age for Marketscan patients. Dataset is patient-level monthly regimen usage (conditional on inside good).

B.3.3 Additional Results

Figure B.3.3: Distribution of Own Price Elasticities – Medicare

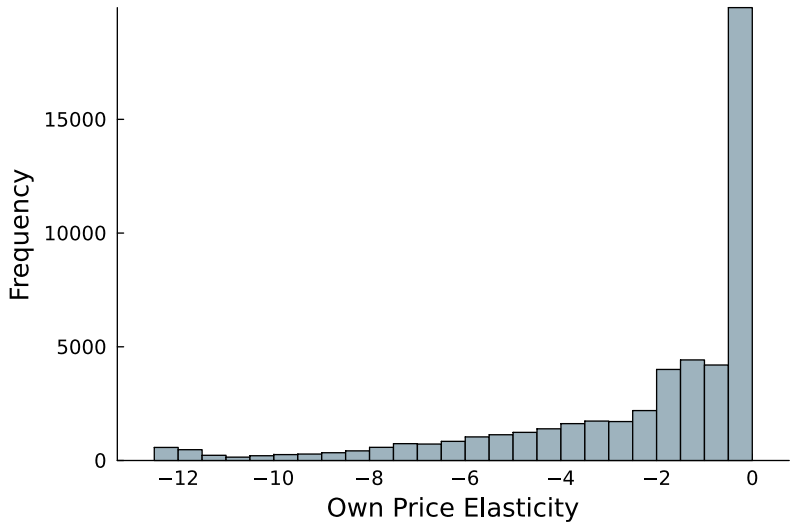

Notes:. Figure shows the distribution of own-price elasticities of demand for Medicare patients.

Figure B.3.4: Distribution of Own Price Elasticities – Medicare Dual

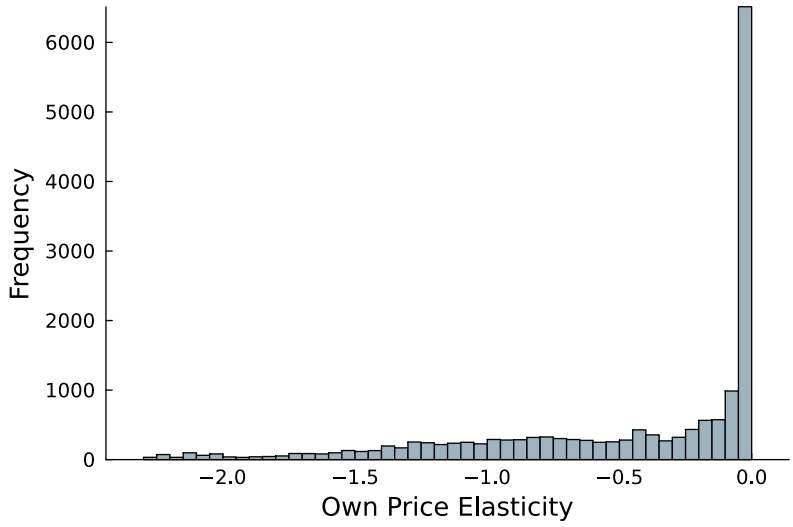

Notes:. Figure shows the distribution of own-price elasticities of demand for Medicare dual-enrolled patients.

Figure B.3.5: Distribution of Own Price Elasticities – Marketscan

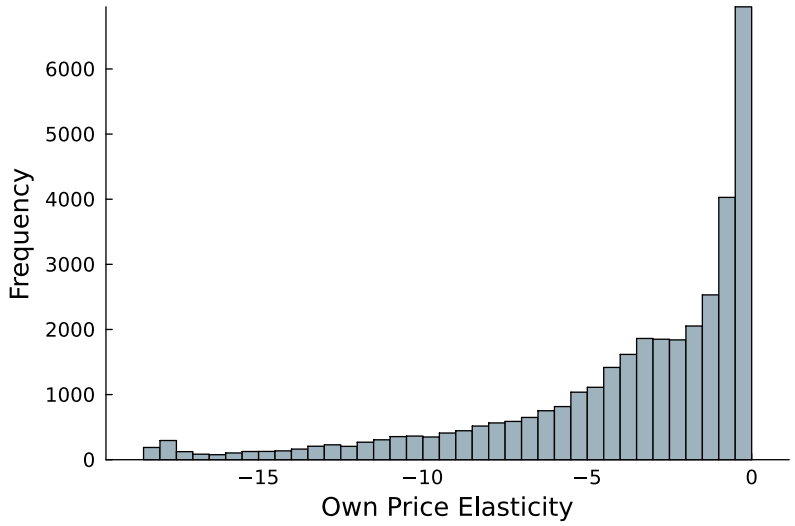

Notes: Figure shows the distribution of own-price elasticities of demand for Marketscan patients.

Figure B.3.6: Distribution of Regimen Fixed Effects – Medicare

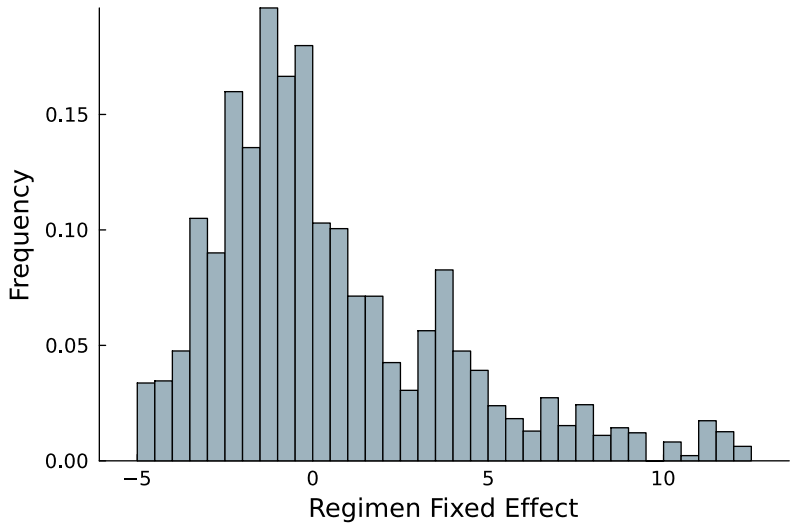

Notes: Figure shows the distribution of regimen fixed effects for Medicare patients.

Figure B.3.7: Regimen Fixed Effect Comparison Across Insurance Types

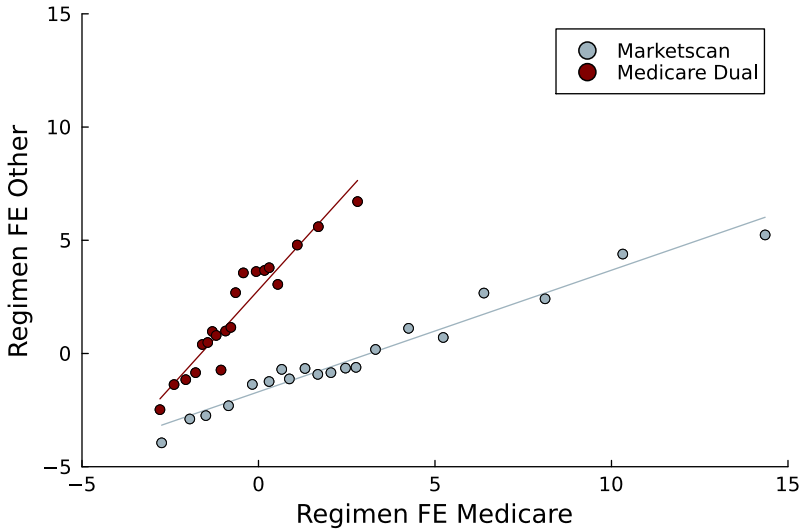

Notes: Figure shows binscatters of the relationship between the regimen FEs estimated in the Medicare sample versus the Marketscan and Medicare Dual samples.

Table B.3.4: Regimen Fixed Effect Comparison Across Insurance Types

	Regimen FE Medicare	
	(1)	(2)
(Intercept)	-1.688	2.807
	(0.042)	(0.068)
Regimen FE Marketscan	0.536	
	(0.008)	
Regimen FE Medicare Dual		1.741
		(0.045)
N	2,444	1,604
R^2	0.634	0.482

Notes: Tables shows regressions of the relationship between the regimen FEs estimated in the Medicare sample versus the Marketscan and Medicare Dual samples.

Table B.3.5: Regimen Fixed Effect Comparison – Complementarities – Colorectal Cancer

Drug 1	**Drug 2**	**Drug 3**	**Drug 4**	**Regimen FE**
regorafenib				7.59
nivolumab				6.72
pembrolizumab				5.35
oxaliplatin	fluorouracil	leucovorin	bevacizumab	4.88
irinotecan	cetuximab			4.77
cetuximab	bevacizumab	irinotecan		4.06
panitumumab				3.84
cetuximab				3.78
irinotecan	fluorouracil	leucovorin	bevacizumab	3.69
irinotecan	fluorouracil	leucovorin	cetuximab	3.59
cetuximab	bevacizumab			3.47
fluorouracil	leucovorin	bevacizumab		2.59
oxaliplatin	fluorouracil	leucovorin		2.08
irinotecan	fluorouracil	leucovorin	panitumumab	1.76
irinotecan				0.86
oxaliplatin	fluorouracil	levoleucovorin		0.71
fluorouracil	leucovorin			0.54
irinotecan	fluorouracil	leucovorin		0.44
fluorouracil	oxaliplatin			-0.22
fluorouracil				-0.34
irinotecan	oxaliplatin	fluorouracil	leucovorin	-0.51
fluorouracil	levoleucovorin			-0.52
capecitabine				-3.1

Notes: Table shows estimated regimen fixed effects ξ_r for colorectal cancer. Supersets of drugs often have higher fixed effects that the subsets of drugs, indicating complementarity.

Table B.3.6: Regimen Fixed Effect Comparison – Complementarities – Breast Cancer

Drug 1	Drug 2	Drug 3	Drug 4	Regimen FE
palbociclib	fulvestrant			10.6
everolimus	exemestane			10.15
palbociclib	letrozole			9.48
docetaxel	carboplatin	pertuzumab	trastuzumab	7.6
ado-trastuzumab emtansine				5.81
trastuzumab				2.61
anastrozole				1.86
eribulin				1.77
docetaxel	carboplatin	trastuzumab		1.73
paclitaxel	bevacizumab			1.53
trastuzumab	paclitaxel			1.51
fulvestrant				1.2
letrozole				1.15
tamoxifen				1.05
trastuzumab	docetaxel			0.72
trastuzumab	vinorelbine			0.17
exemestane				0.16
trastuzumab	gemcitabine			-0.36
docetaxel				-0.47
paclitaxel				-0.51
doxorubicin	cyclophosphamide			-0.73
docetaxel	cyclophosphamide			-0.96
gemcitabine				-0.99
carboplatin	paclitaxel			-1.33
cyclophosphamide	methotrexate	fluorouracil		-1.65
megestrol				-1.87
doxorubicin	cyclophosphamide	docetaxel		-1.87
vinorelbine				-1.88
doxorubicin	cyclophosphamide	paclitaxel		-2.02
carboplatin	docetaxel			-2.08
fluorouracil	epirubicin	cyclophosphamide		-2.49
docetaxel	doxorubicin			-2.92
gemcitabine	paclitaxel			-3.16
doxorubicin				-3.17
cyclophosphamide	doxorubicin	fluorouracil		-3.21
doxorubicin	paclitaxel			-3.83

Notes: Table shows estimated regimen fixed effects ξ_r for breast cancer. Supersets of drugs often have higher fixed effects that the subsets of drugs, indicating complementarity.

B.4 Bargaining

B.4.1 Consumer Surplus Weight

We describe two sets of reasons why the insurer might weight consumer surplus with some weight $\rho > 1$ in its objective.

Patient Price Misperception: The first can be microfounded by patient misperceptions of drug costs. Suppose the insurance market is competitive, so that insurers must offer some minimum level of consumer surplus CS^*. For simplicity, assume there is one drug with demand $x(p)$. The insurer chooses premium Φ and drug price p. The consumer pays $p^c(p)$, but perceives price to be $\mu p^c(p)$ (e.g., as suggested by Abaluck Gruber). The insurer's objective is

$$\max_{\Phi,p} \quad \Phi - (p - p^c(p))x(p) \quad \text{s.t.} \quad -\Phi + \int_0^{x(p^c(p))} x^{-1}(s)ds - \mu \times p^c(p)x(p^c(p)) = CS^* \tag{B.5}$$

We can simplify

$$\int_0^{x(p^c(p))} x^{-1}(s)ds = \int_{p^c(p)}^{\infty} x(s)ds + p^c(p)x(p^c(p)) \tag{B.6}$$

Substituting for Φ, the objective function becomes

$$\max_p \quad \int_{p^c(p)}^{\infty} x(s)ds + p^c(p)x(p) - \mu p^c(p)x(p) - (p - p^c(p))x(p) - CS^* \tag{B.7}$$

Suppose $p^c(p) = \zeta p$ (i.e., constant coinsurance rate). Then the objective can be further simplified as

$$\max_p \quad \int_{p^c(p)}^{\infty} x(s)ds + (\zeta - \mu\zeta - (1-\zeta))px(p) - CS^* \tag{B.8}$$

$$= \quad \int_{p^c(p)}^{\infty} x(s)ds + (\zeta(2-\mu) - 1)px(p) - CS^* \tag{B.9}$$

$$= \quad \int_{p^c(p)}^{\infty} x(s)ds + \frac{(\zeta(2-\mu)-1)}{1-\zeta}(1-\zeta)px(p) - CS^* \tag{B.10}$$

Suppose $\zeta = .2$ and $\mu = .8$. This gives a weight of .95 on insurer costs. This weight is decreasing in μ and ζ. Other reasons that might decrease μ (i.e., lower consumer expectation of cost):

- underestimating sickness / duration of treatment

- overestimating probability of getting high cost drugs (i.e., overestimating probability of prior authorization) – think of as making ex-post cost smaller than expected
- physicians steering to high-cost drugs

Patient Value Misperception: A second reason for $\rho > 1$ could be that consumers themselves "overweight" the (expected) consumer surplus from these drugs. If a patient overweights the probability she will need to take drugs for a cancer, or overweights the value of those drugs (i.e., because of advertising), then in order to get the consumer to choose a particular insurance plan, the insurer might as well.

B.4.2 Estimation Procedure Details

Define $\Delta V_{ft} = \mathrm{V}_{\iota t}(\mathcal{R}_t, p_t) - \mathrm{V}_{\iota t}\left(\mathcal{R}_t \setminus \mathcal{R}_{ft}, p_t\right)$ as the disagreement payoff for firm f at time t. The first-order condition of the Nash Bargaining problem of firm f in Equation 2.5 for p_{dt}, where $d \in \mathcal{D}_{ft}(\mathcal{R}_t)$ is

$$0 = \gamma_f \frac{1}{\pi_{ft}} \frac{\partial \pi_{ft}}{\partial p_{dt}} + \left(1 - \gamma_f\right) \frac{1}{\Delta V_{ft}} \frac{\partial V_t}{\partial p_{dt}}. \tag{B.11}$$

The derivatives are given by

$$\frac{\partial \pi_{ft}}{\partial p_{dt}} = \sum_{\iota} \sum_{c \in \mathcal{C}} M_{\iota ct} \sum_{d' \in \mathcal{D}_{ft}} \sum_{r \in \mathcal{R}_{ct}(d')} \left\{ \left(p_{d't} - mc_{d't}\right) \frac{\partial s_{\iota rt}}{\partial p_{dt}} + s_{\iota rt} \mathbb{1}\left[d = d'\right] \right\} \tag{B.12}$$

Given a bargaining weight, we can invert the FOCs for the vector of marginal costs $(mc_{dt})_{d \in \mathcal{D}}$. When firms are single-product, marginal costs are given by

$$mc_{dt} = \frac{\gamma_f \Delta V_t \sum_{\iota} \sum_{c} M_{\iota ct} \sum_{r \in \mathcal{R}(d)} \left(s_{\iota rt} + p_{dt} \frac{\partial s_{\iota rt}}{\partial p_{dt}} \right) + (1 - \gamma_f) \frac{\partial V_t}{\partial p_{dt}} \sum_{\iota} \sum_{c} M_{\iota ct} \sum_{r \in \mathcal{R}(d)} p_{dt} s_{\iota rt}}{\gamma_f \Delta V_t \sum_{\iota} \sum_{c} M_{\iota ct} \sum_{r \in \mathcal{R}(d)} \frac{\partial s_{\iota rt}}{\partial p_{dt}} + (1 - \gamma_f) \frac{\partial V_t}{\partial p_{dt}} \sum_{\iota} \sum_{c} M_{\iota ct} \sum_{r \in \mathcal{R}(d)} s_{\iota rt}} \tag{B.13}$$

When firms own multiple drugs, we can solve for marginal costs by inverting the stacked first-order conditions.

B.4.3 Solving for Equilibrium Prices

The first-order condition in Equation B.11 may be numerically ill-conditioned when $\pi_{ft} \downarrow 0$ or $\Delta V_{ft} \downarrow 0$, so we instead solve using the transformed FOC

$$0 = \gamma_f \frac{\partial \pi_{ft}}{\partial p_{dt}} \Delta V_{ft} + \left(1 - \gamma_f\right) \frac{\partial V_t}{\partial p_{dt}} \pi_{ft}. \tag{B.14}$$

B.4.4 Additional Results

Figure B.4.1: Biologic Drug Markups

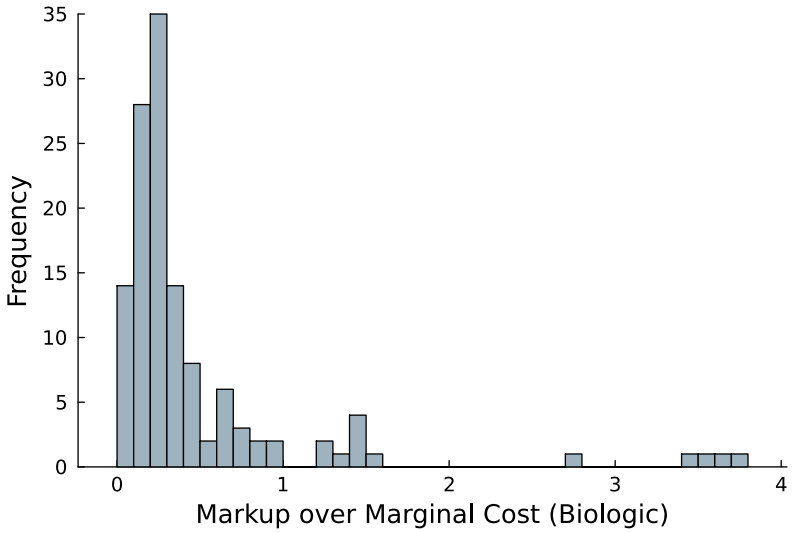

Notes: Figure shows the distribution of markups (calculated as margin over marginal cost) of biologic drugs.

Table B.4.1: Bargaining Weight Robustness

ρ	γ	γ **SE**	**Markup**	**Markup SD**
8.04	0.69	0.29	0.25	0.44
2.5	1.0	0.0	0.11	0.37
5.0	0.89	0.51	0.22	0.41
10.0	0.52	0.14	0.26	0.44

Notes: Table shows how the manufacturer bargaining varies with the consumer surplus weigh ρ. The first row contains our baseline estimates. The remaining rows fix ρ and different levels and estimate the bargaining weight. Model implied markups are computed for each specification.

B.5 Combination Innovation Externalities

B.5.1 Model Fit

Figure B.5.1: Externality Model Fit: Predicted versus Observed Market Shares after Innovation

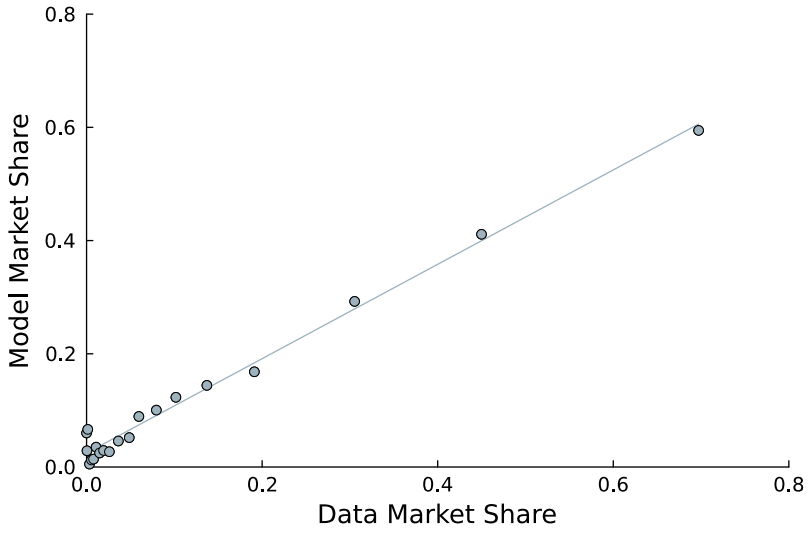

(a) Single Drug Innovation

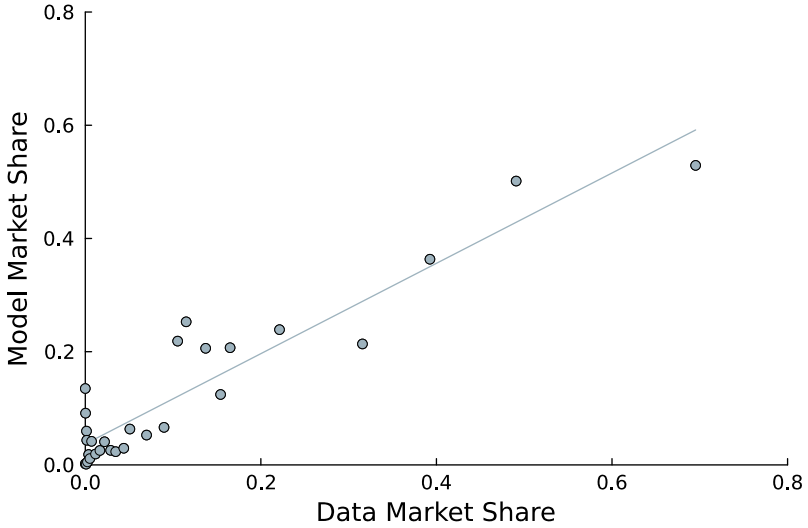

(b) Combination Therapy Innovation

Notes: Figures show binscatters of the relationship between drug market shares predicted by the model and drug market shares observed in the data after innovation events. Panel (a) contains market shares after the introduction of a single drug. Panel (b) contains market shares after the introduction of a combination. Drug market shares are calculated within cancer, summing over market shares of regimens (either single or combination) that drug is used in. Figure only includes branded drugs.

B.6 Dynamic Model

B.6.1 Model Solution Details

Approximation Error Monte Carlo

We compare the exact and approximate expected value functions (over states) for a model with 3 periods, 9 potential regimens, and 4 quality states for each regimen (not trialed, trialed and failed, trailed and success with low quality, and trialed and success with high quality). Solving the model exactly requires evaluating over 250,000 states each period. We choose small number of states ($\approx$ 5000) for the approximation. Figure B.6.1 shows the relationship between the exact expected value function versus the approximation value function. We observe minimal approximation error. Importantly, the error does significantly increase in earlier periods of the game.

Figure B.6.1: Exact EV versus Approximate EV

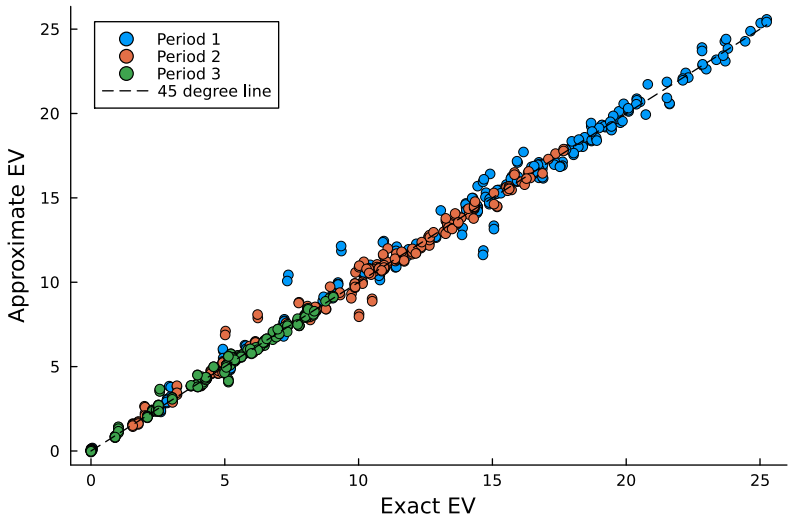

Notes: Figure shows the relationship between the exact expected value function and the approximate expected value function, by period.

B.6.2 Computing Counterfactuals

For focal regimen r, set s_{r0} to be the initial state in data. For each simulation, we proceed as follows. For each t until the model predicts trialing of the focal regimen:

1. Compute flow surplus for each innovator at s_{rt}.

2. Compute the ex-ante choice-specific value functions for each innovator using the sieve

computed at the counterfactual parameters.

3. Draw random variables to determine state transition:

   (a) Draw uniform random variables to determine trialing decisions of each innovator using the conditional choice probabilities that are derived from the ex-ante choice-specific value functions.

   (b) Draw uniform random variables to determine the trialing decisions of tracked (non-focal) regimens.

   (c) Draw uniform random variables to determine trial success and random variables from quality distribution G to determine trial outcomes of non-focal and focal regimens (if applicable).

4. Update the state.

We repeat this procedure a large number of times (250) for each focal regimen and compute average trial times (and fraction of trials by each innovator) over these repetitions.

B.6.3 Estimation Sample

Table B.6.1: Dynamic Model Regimens: Summary Statistics

	Public Trial Mean	**Public Trial** Std Dev	**Private Trial** Mean	**Private Trial** Std Dev
Time to First Trial	5.65	3.91	4.44	3.49
Start Year	2003.96	5.15	2007.74	5.68
Number of Drugs	2.63	0.73	2.49	0.67
Success Indicator	0.01382		0.00947	
Number of Regimens	5138		4014	

Notes: Table shows summary statistics of the regimens included in the dynamic model.

B.6.4 Model Fit

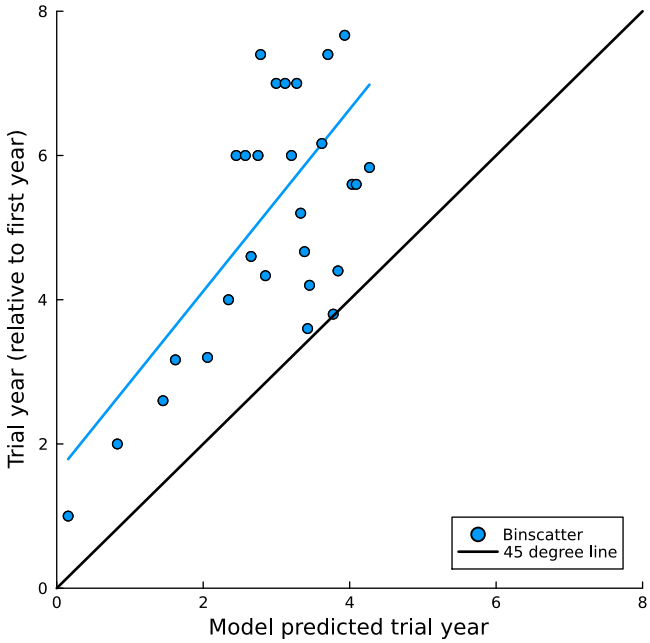

Figure B.6.2: Model Fit - Trial Year

Notes: Figure shows the relationship between the trial year in the data and the model predicted trial year through a binscatter.

B.6.5 Nonparametric CCP Estimation

We produce first-stage estimates of conditional choice probabilities by regressing trialing decisions on functions of state variables, separately for public innovators (Table B.6.2) and private firms (Table B.6.3). We focus on the timing of the trial, summary statistics of other regimens (e.g., number of untested regimens in the state), and cancer fixed effects.

Table B.6.2: Public CCPs

	action				
	(1)	(2)	(3)	(4)	(5)
(Intercept)	0.040	0.357	0.366	0.366	0.306
	(0.003)	(0.020)	(0.020)	(0.020)	(0.023)
Time since first year	0.023				
	(0.001)				
Time since first year (0-5)		-0.286	-0.260	-0.260	-0.262
		(0.020)	(0.020)	(0.020)	(0.020)
Time since first year (5-10)		-0.137	-0.125	-0.125	-0.131
		(0.020)	(0.020)	(0.020)	(0.020)
Time since first year (10-15)		-0.077	-0.071	-0.071	-0.076
		(0.022)	(0.021)	(0.021)	(0.021)
Number of untested regimens in state			-0.011	-0.011	-0.015
			(0.001)	(0.001)	(0.001)
Time since generic (owned drug in regimen)				-0.002	-0.002
				(0.002)	(0.002)
cancer: bladder					0.045
					(0.017)
cancer: breast					0.098
					(0.013)
cancer: cll					0.070
					(0.013)
cancer: colorectal					0.087
					(0.015)
cancer: head and neck					0.053
					(0.014)
cancer: hepatocellular					0.028
					(0.016)
cancer: hodgkin lymphoma					0.052
					(0.013)
cancer: multiple myeloma					0.095
					(0.014)
cancer: non-hodgkin lymphoma					0.074
					(0.013)
cancer: non-small cell lung					0.108
					(0.014)
cancer: ovarian epithelial					0.077
					(0.015)
cancer: pancreatic					0.068
					(0.014)
cancer: prostate					0.058
					(0.014)
N	30,759	30,759	30,759	30,759	30,759
R^2	0.053	0.055	0.063	0.063	0.067

Notes: Table shows CCP estimates for public innovators.

Table B.6.3: Private CCPs

	action				
	(1)	(2)	(3)	(4)	(5)
(Intercept)	0.021	0.005	0.009	0.009	-0.002
	(0.001)	(0.005)	(0.005)	(0.005)	(0.006)
Time since first year	-0.001				
	(0.000)				
Time since first year (0-5)		0.015	0.027	0.027	0.028
		(0.005)	(0.005)	(0.005)	(0.005)
Time since first year (5-10)		0.011	0.016	0.016	0.015
		(0.005)	(0.005)	(0.005)	(0.005)
Time since first year (10-15)		0.001	0.004	0.004	0.003
		(0.006)	(0.006)	(0.006)	(0.006)
Number of untested regimens in state			-0.005	-0.005	-0.008
			(0.000)	(0.000)	(0.000)
Time since generic (owned drug in regimen)				-0.001	-0.001
				(0.000)	(0.000)
cancer: bladder					0.004
					(0.004)
cancer: breast					0.039
					(0.003)
cancer: cll					0.011
					(0.003)
cancer: colorectal					0.023
					(0.004)
cancer: head and neck					0.007
					(0.004)
cancer: hepatocellular					0.013
					(0.004)
cancer: hodgkin lymphoma					0.001
					(0.003)
cancer: multiple myeloma					0.031
					(0.004)
cancer: non-hodgkin lymphoma					0.013
					(0.003)
cancer: non-small cell lung					0.047
					(0.003)
cancer: ovarian epithelial					0.014
					(0.004)
cancer: pancreatic					0.015
					(0.004)
cancer: prostate					0.001
					(0.004)
N	74,430	74,430	74,430	74,430	74,430
R^2	0.001	0.001	0.010	0.010	0.020

Notes: Table shows CCP estimates for private innovators.

B.6.6 Alternative Estimation Procedure: Euler Perturbations

This section derives estimating equations for exploiting finite dependence to estimate the model via Euler perturbations.

Estimation Equations

Let $p_{rft}(s_{rt})$ be the CCP for focal player given state s_{rt} (focal regimen r, focal player f, time t), and let $p_{rftc}(s_{rt})$ be the (aggregate) CCP of competitors.

Compare two sequences of *planned* trialing actions for an agent

1. $(0,1,1,\ldots)$
2. $(1,1,1,\ldots)$

Then compare choice-specific *conditional* value functions for these choices, and observe that the game ends up in same state by period $t+2$:

$$v_{rft}(s_{rt};0,1) = \pi_{rft}(s_{rt}) + \beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ doesn't trial at } t) + (1-p_{rftc}(s_t))\kappa_{rft+1}\right] + \beta^2\mathbb{E}\left[v_{rft+2}(s_{t+2};1)\right],$$
$$v_{rft}(s_{rt};1) = \pi_{rft}(s_{rt}) - \kappa_{rft} + \beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ trials at } t)\right] + \beta^2\mathbb{E}\left[v_{rft+2}(s_{t+2};1)\right].$$

Subtracting these terms gives

$$\begin{aligned} v_{rft}(s_{rt};1) - v_{rft}(s_{rt};0,1) = &-\kappa_{rft} + \beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ trials at } t)\right] \\ &- \beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ doesn't trial at } t) + (1-p_{rftc}(s_t))\kappa_{rft+1}\right]. \end{aligned}$$

This leads to a two-step estimation approach. The first step is to nonparametrically estimate CCPs. In the second step, we use the recovered CCPs and the Hotz-Miller inversion to estimate dynamic parameters:

$$\ln\left(\frac{p_{rft}(s_{rt})}{1-p_{rft}(s_{rt})}\right) = \frac{1}{\theta^{\epsilon}}\left(v_{rft}(s_{rt};1) - v_{rft}(s_{rt};0)\right).$$

Notice however the choice-specific conditional value function $v_{rft}(s_{rt};0,1)$ is not the same as the choice-specific value function $v_{rft}(s_{rt};0)$. However, we can estimate the difference between

these terms. We have that

$$\begin{aligned}
V_{rft}(s_t) &= \theta^{\epsilon}\mathbb{E}\left[\ln\left(\exp\left(\frac{1}{\theta^{\epsilon}}v_{rft}(s_t;0)\right)+\exp\left(\frac{1}{\theta^{\epsilon}}v_{rft}(s_t;1)\right)\right)\right] \\
&= v_{rft}(s_t;1)-\theta^{\epsilon}\ln p_{rft}(s_{rt}) \\
&= v_{rft}(s_t;0)-\theta^{\epsilon}\ln(1-p_{rft}(s_{rt}))
\end{aligned}$$

and

$$\begin{aligned}
v_{rft}(s_{rt};0) &= \pi_{rft}(s_{rt})+\beta\mathbb{E}\left[V_{rft+1}(s_{t+1})\right] \\
&= \pi_{rft}(s_{rt})+\beta\mathbb{E}\left[v_{rft+1}(s_{t+1};1)-\theta^{\epsilon}\ln p_{rft+1}(s_{t+1})\right]
\end{aligned}$$

and

$$\begin{aligned}
v_{rft}(s_{rt};0,1) &= \pi_{rft}(s_{rt})+\beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ doesn't trial at } t)+(1-p_{rftc}(s_t))\kappa_{rft+1}\right]+\beta^2\mathbb{E}\left[v_{rft+2}(s_{t+2};1)\right] \\
&= \pi_{rft}(s_{rt})+\beta\mathbb{E}\left[(1-p_{rftc})v_{rft+1}(s_{t+1}|f \text{ trials});1)+p_{rftc}(v_{rft+1}(s_{t+1}|f \text{ doesn't trial};1))\right] \\
&= \pi_{rft}(s_{rt})+\beta\mathbb{E}\left[v_{rft+1}(s_{t+1};1)+p_{rftc}\kappa_{rft+1}\right]
\end{aligned}$$

This gives

$$\begin{aligned}
v(s_{rt};1)-v(s_{rt};0) &= v(s_{rt};1)-(v(s_{rt};0,1)+\text{correction}) \\
&= -\kappa_{rft}+\beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ trials at } t)\right]-\beta\mathbb{E}\left[\pi_{rft+1}(s_{t+1}|f \text{ doesn't trial at } t)+(1-p_{rftc}(s_t))\kappa_{rft+1}\right] \\
&\quad +\beta\mathbb{E}\left[\theta^{\epsilon}\ln p_{rft+1}(s_{t+1})+p_{rftc}(s_{rt})\kappa_{rft+1}\right]
\end{aligned}$$

This is a similar result to Arcidiacono and Ellickson (2011) (and applied, in, e.g., Scott, 2014), but with a correction for free-rider problem: there is some probability of ending up in trialed state but not actually having to run trial.

We can then apply the standard rational expectations assumption. Under the assumption of rational expectations, the conditional expectation of CCPs and profits at $t+1$ is equal to the realized variables minus an expectational error. We assume this error is orthogonal to state variables at period t. That expectational error will be included in the regression error term.

The second step estimation equation is then

$$\ln\left(\frac{p_{rft}(s_{rt})}{1-p_{rft}(s_{rt})}\right)-\beta\ln p_{rft+1}(s_{rt+1})=$$
$$\frac{1}{\theta^{\epsilon}}(-\kappa_{rft}+\beta\pi_{rft+1}(s_{rt+1}|f \text{ trials at } t)-\beta\pi_{rft+1}(s_{rt+1}|f \text{ doesn't trial at } t)+\beta p_{rftc}(s_{rt})\kappa_{rft+1}$$
$$-\beta(1-p_{rftc}(s_{rt}))\kappa_{rft+1}+\text{expectation error})$$
$$=\frac{1}{\theta^{\epsilon}}(\beta\pi_{rft+1}(1-p_{rftc}(s_{rt}))+\kappa_{rft}(-1+\beta(2p_{rftc}(s_{rt})-1)))+\text{expectation error}$$

B.7 Institutional Details

B.7.1 Treatment Decisions

Cancer treatment decisions are influenced by a variety of factors, including the type and stage of cancer, patient characteristics, regimen characteristics, prior therapies, and the individual response to treatment. Treatments regimens, which may include single agents or combination therapies, are typically classified as first-line, second-line, or beyond. However, the same regimen might be recommended at different stages depending on these factors, as the choice of therapy is tailored to the patient's needs and the progression of their disease.

Figures B.7.1, B.7.2, and B.7.3 illustrate potential treatment pathways for patients with advanced or metastatic colon cancer who are candidates for intensive therapy (National Comprehensive Cancer Network, 2024). Notably, certain regimens may appear as options across multiple lines of treatment—whether as first-line, second-line, or later therapies—highlighting their substitutability in different clinical contexts. These regimens may include single agents or drug combinations, tailored to the patient's disease progression and treatment history.

Figure B.7.1: Example Treatment Progression (1)

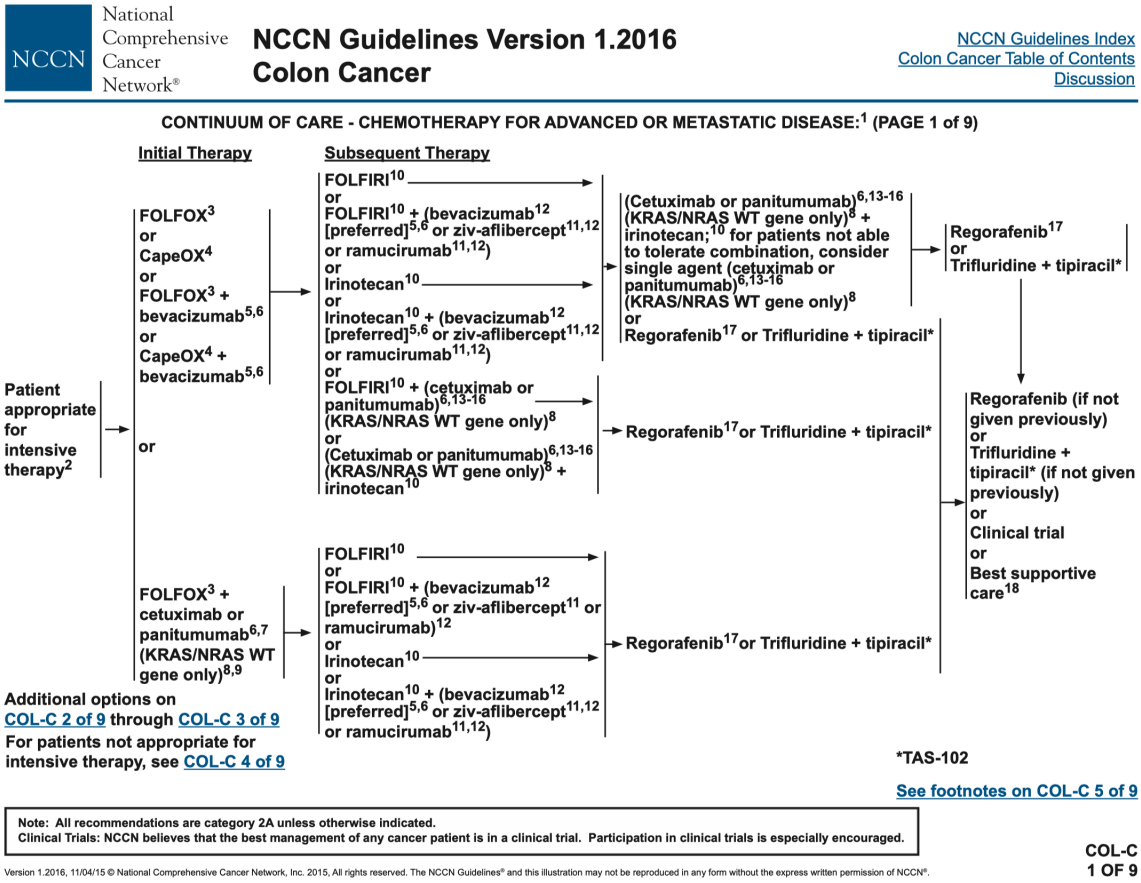

Notes: Figure shows treatment pathways for patients with advanced or metastatic colon cancer (National Comprehensive Cancer Network, 2024).

Figure B.7.2: Example Treatment Progression (2)

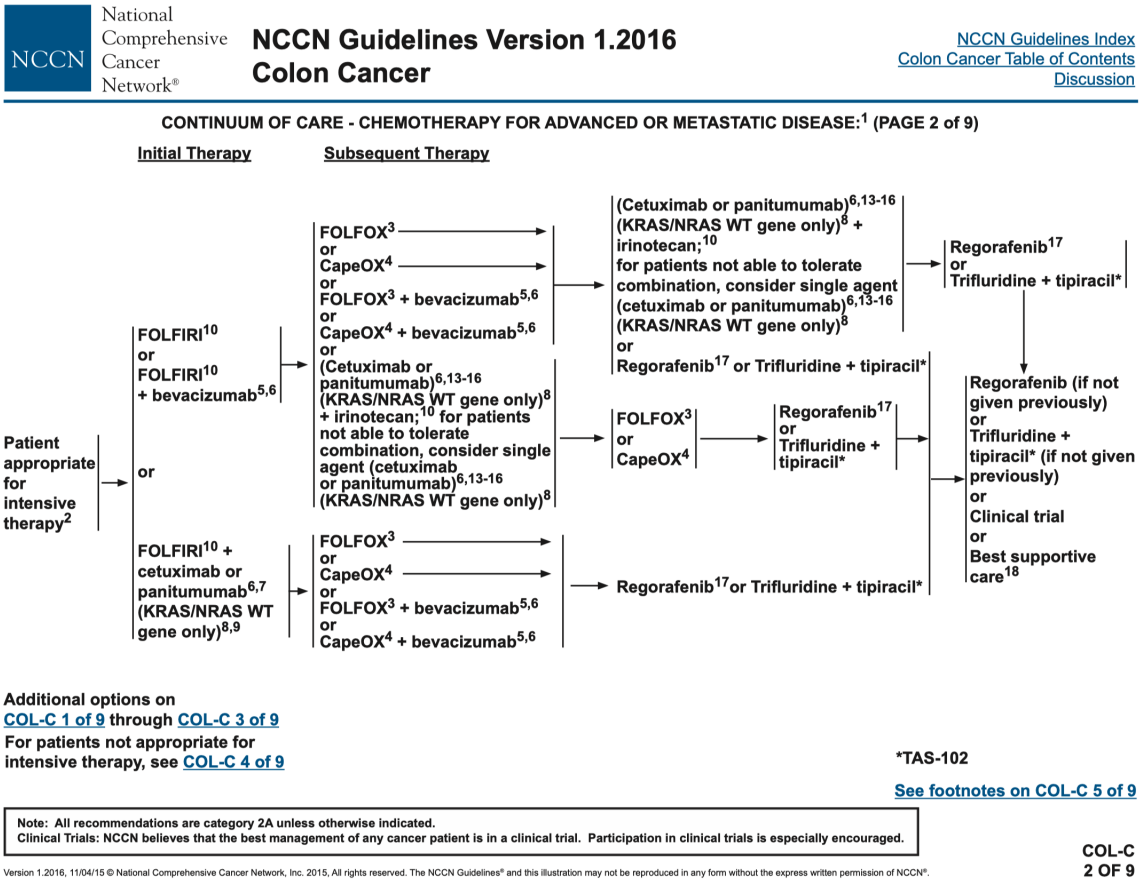

Notes: Continuation of Figure B.7.1.

Figure B.7.3: Example Treatment Progression (3)

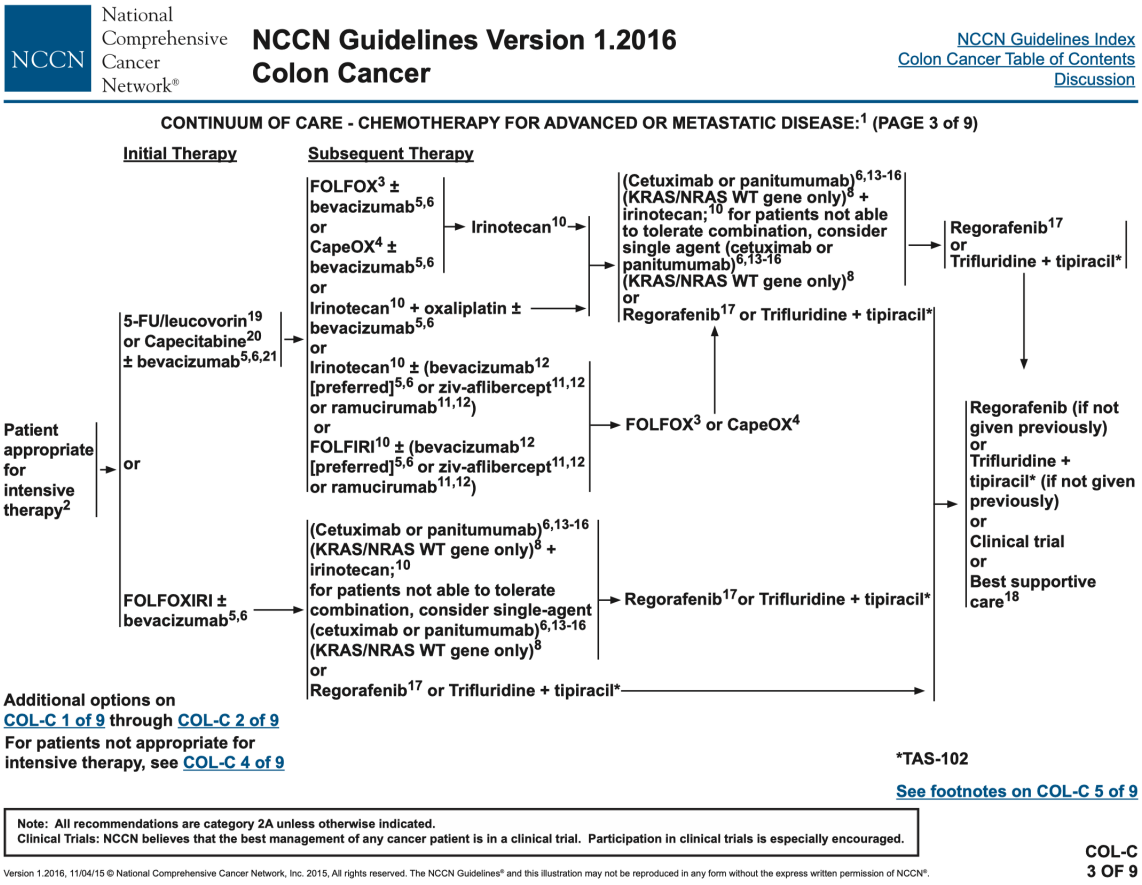

Notes: Continuation of Figure B.7.1.

B.7.2 Clinical Trial Cost Components

Table B.7.1: Clinical Trial Cost Components

Cost Component	Cost Group	Mean Percent
Data Management Costs	Administrative	1.0
Cost Per IRB Approvals	Administrative	0.9
Cost of IRB Amendments	Administrative	0.0
Administrative Staff Costs	Administrative	19.6
Patient Recruitment Costs	Patient Recruitment	2.3
Patient Retention Costs	Patient Recruitment	0.2
Clinical Procedure Total	Clinical Procedure	19.2
SDV Costs	Site	6.8
RN/CRA Costs	Site	7.4
Physician Costs	Site	5.8
Central Lab Costs	Site	8.4
Site Recruitment Costs	Site	2.6
Site Retention Costs	Site	12.9
Site Monitoring Costs	Site	12.9

Notes: Table shows estimates of clinical trial cost components (as a percentage of total cost) from Sertkaya et al. (2014). The cost of drugs used in the trial would be recorded under the category "Clinical Procedure Total," which also likely includes costs associated with administering the drugs.

B.7.3 Contracting Frictions

The development of cancer drug combination therapies faces several significant contracting frictions that hinder collaboration between firms and the ability of firms to internalize innovation externalities through Coasian bargaining.

One major challenge is the uncertainty surrounding the commercial potential of combination regimens. This uncertainty complicates negotiations between firms, often leading to a "hold up" problem where each party tries to extract more value than their contribution warrants. As Humphrey et al. (2011) notes, this can result in negotiations being viewed as a "zero-sum exercise" with perceived economic winners and losers.

Intellectual property (IP) concerns also pose substantial barriers to collaboration. Companies are often reluctant to combine their drugs with those from other firms due to worries about potential IP violations and liabilities. This extends to concerns about "secondary IP" that might arise from unexpected therapeutic benefits of drug combinations. As a result, firms tend to prefer developing combinations using only their own drugs, as it simplifies the IP landscape:

- "One company might have a candidate therapy that would make sense to test with a drug from a different firm. But because the two firms hold the patents to each separately, both parties might worry about future liabilities, intellectual property (IP) rights, and secondary IP (that is, IP issues that might arise from unexpected new therapeutic benefits from combining the drugs)" (Deng, 2015).

- "They would rather do it with two of their own drugs, because it makes life easy" (Institute of Medicine et al., 2012).

- "A major impediment to companies sharing their cell lines and drug candidates preclinically is intellectual property issues, while others stressed that intellectual property rights impede clinical trials of combination cancer therapies" (Institute of Medicine et al., 2012).

Another significant friction arises from the potential for adverse effects in combination therapies. Companies fear that unexpected negative outcomes in combination trials could harm the development prospects of their individual drugs. This risk is particularly acute for drugs still in the investigational stage:

- "There was a day, years ago, when the only drugs worth testing were coming from CTEP [NCI's Cancer Therapy Evaluation Program]," said Johnson. But these days, most drugs come from companies, whose corporate cultures—and in particular, their legal departments—render them reluctant to collaborations with potential competitors" (Goldman, 2003).

- "Such combination trials are relatively easy when the compounds are all owned or licensed by one company, but what about combining an investigational drug with another investigational drug from a different company? Companies developing one or both compounds have real concerns—not the least of which is the fear of fortuitous bad reactions. "Suppose we did a combination trial," Johnson mused, "and had some catastrophic result—like the first three patients just up and died within 2 hours of being treated. That would put a cold chill on both drugs. That's the conundrum that companies face" (Goldman, 2003).

Antitrust concerns and pricing issues further complicate the landscape. Some companies express reluctance to engage in collaborative R&D due to perceived antitrust risks (Institute of Medicine et al., 2012):

- "Some drug companies have expressed reluctance to conduct collaborative R&D on investigational drugs with other companies because of concerns about violating antitrust laws" (Institute of Medicine et al., 2012).

Additionally, competition laws in various jurisdictions can prevent manufacturers from agreeing on pricing strategies for combination treatments without impacting standalone drug prices, limiting their ability to optimize pricing for different use cases. Podkonjak et al. (2021) highlights this friction in the UK, and similar institutional details also apply to the US context:

- "UK competition law, enforced by the Competition Markets Authority (CMA), prevents individual manufacturers agreeing prices for their treatments as part of an agreement for splitting revenues from combination treatments, where this has the effect of also impacting prices for the treatments when sold on a standalone basis. It also prohibits the exchange of pricing or other sensitive commercial information that could have the effect of limiting competition between the manufacturers when supplying their treatments on a standalone basis" (Podkonjak et al., 2021) (statement for UK, but something similar true in US).

B.7.4 Exclusivity

A drug is protected from competition through two mechanisms: patents and FDA "data exclusivity." A firm can directly patent a drug (more precisely an "active moiety," the part of a drug responsible for the physiological effect) after discovery, in which case it has the sole right to manufacture and market the drug for 20 years, with the exception of the Bolar exemption which protects the rights of others to use the drug in research. It can further apply for patents related to other formulations and indications at any time that would extend the firm's exclusive right to market the drug for these new purposes, but would not interfere with generic competition under the original formulation or indications.

The FDA's data exclusivity provisions function similarly: The firm(s) applying for New Chemical Entity (NCE) exclusivity must submit a New Drug Application (NDA) describing the drug and the clinical trials used to verify the drug's safety and efficacy for a given formulation and indication. A successful NDA gives the firm(s) an exclusive right to market the drug in the US for 5 years, unless another firm undertakes duplicate clinical trials that verify safety and efficacy. (In practice, the latter does not occur either because a drug with an NDA is also under patent protection or because the time and expense needed to conduct clinical trials would exceed the profits that could be earned before the expiry of the initial NDA.) Note that the FDA requires the NDA applicant to have "conducted or sponsored the study by providing 50 percent of the funding or by purchasing exclusive rights to the study." While a given drug is still under NCE exclusivity, the "exclusive firm(s)" can conduct additional trials to show safety and efficacy for new formulations and indications, including the use of the drug in combination with others.

These additional trials allow the firm(s) to apply for New Clinical Investigation (NCI) exclusivity, which protects the right to market the drug according to its new uses for an additional 3 years regardless of the patent or NCE exclusivity term. Other firms can also conduct clinical trials and attempt to receive NCI exclusivity while NCE exclusivity is still in force, but this would again be precluded by patent protection and may require additional clinical trials if data produced for the original NDA is needed to establish safety/efficacy. For these reasons, we expect other firms to apply for NCI exclusivity only after patent protection and the original NCE exclusivity period have lapsed.

It is not clear how valuable additional "method-of-use" patents or NCI exclusivity provisions are for firms after the original patent and NCE exclusivity have lapsed. In principle, off-label use allows doctors to prescribe generic versions of a drug for use in combinations or for indications that have only been tested using the branded version of a drug, though insurance reimbursement restrictions may affect this. There is some concern that method-of-use patents or NCI exclusivity provisions may have a chilling effect on generic entry after the original patent and NCE exclusivity have lapsed if the entrants take any action to market the drug for new uses, but it is not clear how important this is (Strohbehn et al., 2021; Feldman, 2022; Tu and Sarpatwari, 2023). In general, it appears that firms only have incentives to conduct new clinical trials for their own drugs that are still under original patent or NCE exclusivity protection.

We make a first attempt at understanding the prevalence of patenting for combinations by using data from Durvasula et al. (2023) to extract patents and exclusivity extensions associated with combinations. We first subset to patents and exclusivity instances that pertain to oncology drugs. We count the mean number of patent use codes and exclusivity codes that a particular drug has per year, and we calculate the same means for patents and exclusivity that pertain to combinations. We say a particular patent use code or exclusivity code is associated with combinations if it includes more than one drug in the description or says the word combination (or close variants). There is an increasing number of patents for each drug associated with combinations (Figure B.7.4), but these patents do not seem to translate into explicit exclusivity extensions by the FDA (Figure B.7.5).

Figure B.7.4: Patent Use Codes and Combinations

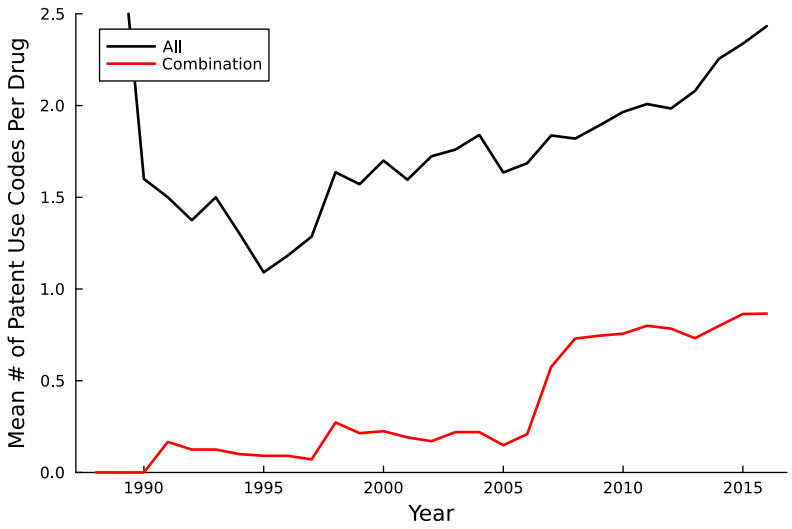

Notes: Figure shows the mean number of patent use codes per drug by year in total and that are associated with combinations.

Figure B.7.5: Exclusivity Codes and Combinations

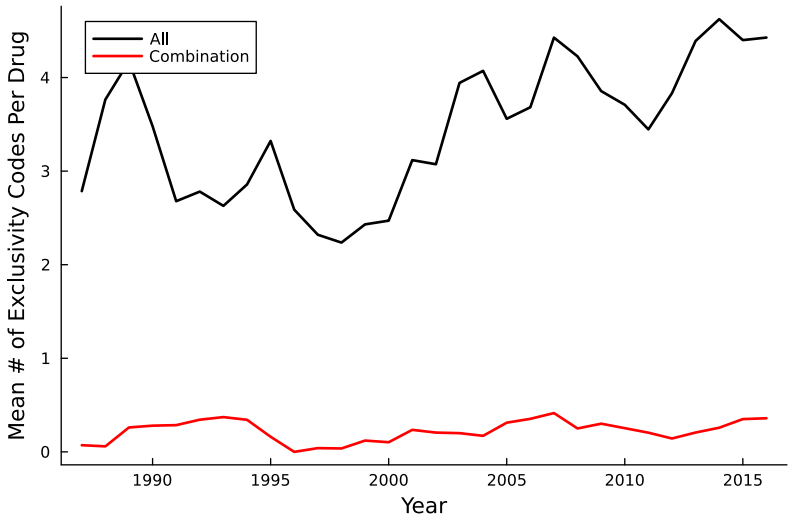

Notes: Figure shows the mean number of exclusivity codes per drug by year in total and that are associated with combinations.

B.7.5 Clinical Trials

ClinicalTrials.gov defines a "sponsor" as the entity that initiates the study, while a "collaborator" is any organization other than the sponsor that provides support (funding, design, implementation, data analysis, or reporting). The legal definition of "sponsor" is provided in 21 CFR 50.3:

"A person who initiates a clinical investigation, but who does not actually conduct the investigation, i.e., the test article is administered or dispensed to or used involving, a subject under the immediate direction of another individual. A person other than an individual (e.g., corporation or agency) that uses one or more of its own employees to conduct a clinical investigation it has initiated is considered to be a sponsor (not a sponsor-investigator), and the employees are considered to be investigators."

The Food and Drug Administration Amendments Act of 2007 requires that clinical trial information submissions include the sponsor, but not collaborators (42 CFR 11). It is not clear how extensive reporting of collaborators is.

Appendix C

Appendix to Regulating Transformative Technologies

Appendix C.1 contains proofs for the results in the main text. In Appendices C.2, C.3, and C.4, we analyze extensions of our benchmark model and discuss the robustness of our main results.

C.1 Proofs for the Main Text

In this part of the Appendix, we provide proofs for results in the main text.

C.1.1 Proofs for Section 3.3

Proof of Proposition 8. Let $L(t) \equiv L(\mu(t), q(t))$ denote the damage threshold at time $t \in (0, \infty)$. Making use of Proposition 7 and the equation $q(t) = q(0) + (g_N - g_O)t$, the damage threshold equals

$$L(t) = \alpha + (\rho - g_N)\left[\frac{\alpha - \exp(-[q(0) + (g_N - g_O)t])}{\mu(t)\lambda\eta} - \frac{\exp(-[q(0) + (g_N - g_O)t])}{\rho - g_O}\right]. \tag{C.1}$$

It is immediate that $L(t)$ is strictly decreasing in g_O, so adoption at time t is nonincreasing in g_O. Note that adoption is also strictly decreasing whenever $L(t) \in (\underline{\delta}, \bar{\delta})$.

Considering instead the comparative static with respect to g_N, we can differentiate to find

$$\frac{\partial L(t)}{\partial g_N} = (1+(\rho-g_N)t)\left[\frac{1}{\mu(t)\lambda\eta}+\frac{1}{\rho-g_O}\right]\exp(-[q(0)+(g_N-g_O)t])-\frac{\alpha}{\mu(t)\lambda\eta}.$$

This derivative is positive iff

$$(1+(\rho-g_N)t)\left[1+\frac{\mu(t)\lambda\eta}{\rho-g_O}\right]\exp(-[z(0)+(g_N-g_O)t])\geq\alpha.$$

The left side limits to zero as $t\to\infty$, so there exists an earliest time $\bar{t}<\infty$ such that $\partial L(t)/\partial g_N<0$ for $t>\bar{t}$. If $\bar{t}>0$, then the left side of the inequality above must be decreasing in t at $t=\bar{t}$. Since the left side is also decreasing in g_N, we must have that $\bar{t}$ is decreasing in g_N.

Finally, the bracketed term in (C.1) limits to a finite value as g_N increases to ρ, which implies that $L(t)$ limits to α. Since the lower support of F satisfies $\alpha\leq\underline{\delta}$, we conclude that $X(\mu(t),q(t))$ limits to zero. ■

Proof of Proposition 9. Using the expression for the damage threshold (3.8), we can calculate

$$\frac{\dot{L}(\mu,q)}{\rho-g_N}=\frac{1-\mu}{\mu}\frac{\alpha-\exp(-q)}{\eta}+\left(\frac{1}{\mu\lambda\eta}+\frac{1}{\rho-g_O}\right)(g_N-g_O)\exp(-q).$$

This equation implies that $\dot{L}(\mu,q)$ is strictly decreasing in g_O. Differentiating implies that $\dot{L}(\mu,q)$ is strictly decreasing in g_N iff

$$\alpha\exp(q)-1>\frac{(\rho-g_N)-(g_N-g_O)}{1-\mu}\left(\frac{1}{\lambda}+\frac{\mu\eta}{\rho-g_O}\right).$$

Differentiating $\dot{L}(\mu,q)$ again yields

$$\frac{\ddot{L}(\mu,q)}{\rho-g_N}=\lambda\frac{1-\mu}{\mu}\frac{\alpha-\exp(-q)}{\eta}+\left[2\frac{1-\mu}{\mu}-\left(\frac{1}{\mu\lambda}+\frac{\eta}{\rho-g_O}\right)(g_N-g_O)\right](g_N-g_O)\frac{\exp(-q)}{\eta}.$$

Provided that $\alpha>\exp(-q)$, this expression immediately implies that $\ddot{L}(\mu,q)<0$ iff $g_N-g_O>G(\mu,q)$, where $G(\mu,q)$ is the largest solution to the quadratic equation

$$\lambda\frac{1-\mu}{\mu}\frac{\alpha-\exp(-q)}{\eta}+\left[2\frac{1-\mu}{\mu}-\left(\frac{1}{\mu\lambda}+\frac{\eta}{\rho-g_O}\right)G(\mu,q)\right]G(\mu,q)\frac{\exp(-q)}{\eta}=0.$$

Equivalently,

$$G(\mu,q)=\lambda\frac{1+\sqrt{1+\left(1+\frac{\mu\lambda\eta}{\rho-g_O}\right)(1-\mu)^{-1}(\alpha\exp(q)-1)}}{\left(1+\frac{\mu\lambda\eta}{\rho-g_O}\right)(1-\mu)^{-1}}.$$

We observe that G is decreasing in μ and increasing in q, so that $\dot{G}(\mu,q)>0$ with

$$\lim_{t\to\infty}G(\mu(t),q(t))=\infty.$$

■

Proof of Proposition 10. See the proof of Proposition C.3.2 in Appendix C.3. ■

C.1.2 Proofs for Section 3.5

Proof of Proposition 14. With a sector-independent use tax $\tau(\mu,Q)$, it is privately optimal to use technology N before the disaster iff

$$\alpha Q_N-Q_O-\tau(\mu,Q)>\mu\lambda\eta\left[\frac{1}{\rho-g_O}Q_O-\frac{\alpha-\gamma_i}{\rho-g_N}Q_N\right]. \tag{C.2}$$

The right side of this inequality is strictly increasing in γ_i. Given an initial state $(\mu(0),Q(0))$, let $\tilde{t}_i$ denote the time at which firm i begins using technology N. For any sector j with private damages γ_j, we immediately observe that $\gamma_i\leq\gamma_j$ iff $\tilde{t}_i\leq\tilde{t}_j$. The latter inequality is strict if $\gamma_i<\gamma_j$ and $\tilde{t}_j>0$.

If γ_i and δ_i are positively affiliated, the tax (3.13) suffices to implement socially optimal technology choices in equilibrium. To see this, note that the private optimality condition (C.2) implies that firm i uses technology N in state (μ,Q) iff $\gamma_i<\hat{L}(\mu,q)$, where

$$\frac{\hat{L}(\mu,q)-\alpha+L(\mu,q)-\kappa(L(\mu,q))}{\rho-g_N}=\frac{\alpha-\exp(-q)}{\mu\lambda\eta}-\frac{\exp(-q)}{\rho-g_O}.$$

Using the definition of the damage threshold $L(\mu,q)$ from Proposition 7, this equation reduces to $\hat{L}(\mu,q)=\kappa(L(\mu,q))$. Since κ is strictly increasing, we conclude that equilibrium technology choices are efficient: Firm i uses technology N iff

$$\delta_i=\kappa^{-1}(\gamma_i)<\kappa^{-1}\left(\hat{L}(\mu,q)\right)=L(\mu,q).$$

Finally, fix an initial state $(\mu(0), Q(0))$ such that $L(\mu(0), Q(0)) < \underline{\delta}$, so that it is inefficient for any sector to use technology N at $t = 0$. Suppose that a given sector-independent tax $\tau(\mu, Q)$ implements socially optimal technology choices in equilibrium. We can define the affiliation function κ as follows: For any value of social damages $\delta \in \left[\underline{\delta}, \bar{\delta}\right]$, let $t(\delta) > 0$ be the time at which sectors with social damages δ (socially) optimally begin using technology N. Since $\tau(\mu, Q)$ implements socially optimal technology choices in equilibrium, these same sectors must find it privately optimal to begin using technology N at time $t(\delta)$. These sectors must have a common value of private damages $\gamma(t(\delta))$: If one sector had a larger value of private damages $\gamma' > \gamma(t(\delta))$, it would find it privately optimal to delay using technology N, contradicting the assumption that τ implements socially optimal technology choices. As a result, the affiliation function $\kappa(\delta) = \gamma(t(\delta))$ is well-defined, and we conclude that social and private damages must be positively affiliated. ■

Proof of Proposition 15. Given a threshold $\hat{\delta}$ and wait time $\hat{T}$, the planner's objective discounted to $t=0$ can be written

$$\begin{aligned}V\left(\hat{\delta},\hat{T}\right)=&\int_0^{\hat{T}}\exp\left(-\rho t\right)\int_{\delta_i<\hat{\delta}}\Big\{\left(1-x\left(\mu(t),q(t),\gamma_i\right)\right)\left[1+\mu(t)\lambda\eta\frac{1}{\rho-g_O}\right]Q_O(t)\\&\qquad+x\left(\mu(t),q(t),\gamma_i\right)\left[\alpha+\mu(t)\lambda\eta\frac{\alpha-\delta_i}{\rho-g_N}\right]Q_N(t)\Big\}didt\\&+\int_0^{\hat{T}}\exp\left(-\rho t\right)\int_{\delta_i\geq\hat{\delta}}\left[1+\mu(t)\lambda\eta\frac{1}{\rho-g_O}\right]Q_O(t)didt\\&+\int_{\hat{T}}^{\infty}\exp\left(-\rho t\right)\int_0^1\Big\{\left(1-x\left(\mu(t),q(t),\gamma_i\right)\right)\left[1+\mu(t)\lambda\eta\frac{1}{\rho-g_O}\right]Q_O(t)\\&\qquad+x\left(\mu(t),q(t),\gamma_i\right)\left[\alpha+\mu(t)\lambda\eta\frac{\alpha-\delta_i}{\rho-g_N}\right]Q_N(t)\Big\}didt.\end{aligned}$$

Here $x\left(\mu,q,\gamma_i\right)$ denotes the unrestricted equilibrium technology choice given state (μ,q) and private damages γ_i:

$$x\left(\mu,q,\gamma_i\right)=\begin{cases}1 & \text{if } \alpha_i-\exp(-q)>\mu\lambda\eta\left[\frac{1}{\rho-g_O}\exp(-q)-\frac{\alpha-\gamma_i}{\rho-g_N}\right],\\ 0 & \text{else.}\end{cases}$$

With $\hat{\delta}$ fixed, we can differentiate V with respect to $\hat{T}$ to find

$$\exp\left(\rho\hat{T}\right)\frac{\partial V\left(\hat{\delta},\hat{T}\right)}{\partial\hat{T}}=-\int_{\delta_i\geq\hat{\delta}}x\left(\mu,q,\gamma_i\right)\left\{\alpha Q_N-Q_O-\mu\lambda\eta\left[\frac{1}{\rho-g_O}Q_O-\frac{\alpha-\delta_i}{\rho-g_N}Q_N\right]\right\}di.$$

To simplify notation, we have left the dependence of the state (μ,Q) on the wait time $\hat{T}$ implicit. First observe that the optimal wait time is bounded:

$$\lim_{\hat{T}\to\infty}\frac{1}{Q_N\left(\hat{T}\right)}\exp\left(\rho\hat{T}\right)\frac{\partial V\left(\hat{\delta},\hat{T}\right)}{\partial\hat{T}}=-\alpha\int_{\delta_i\geq\hat{\delta}}di<0.$$

Let $\tilde{t}_i$ denote the equilibrium time of adoption for sector i when unrestricted, and let $\underline{t}(\hat{\delta})\geq 0$ denote the greatest lower bound for these times across all sectors above the threshold ($\delta_i\geq\hat{\delta}$). Note that we can write

$$\exp\left(\rho\hat{T}\right)\frac{\partial V\left(\hat{\delta},\hat{T}\right)}{\partial\hat{T}}=-\int_{\delta_i\geq\hat{\delta},\tilde{t}_i\leq\hat{T}}\alpha Q_N-Q_O-\mu\lambda\eta\left[\frac{1}{\rho-g_O}Q_O-\frac{\alpha-\delta_i}{\rho-g_N}Q_N\right]di.$$

Clearly $\partial V(\hat{\delta},\hat{T})/\partial \hat{T} = 0$ for $\hat{T} \leq \underline{t}\left(\hat{\delta}\right)$. But $\partial V(\hat{\delta},\hat{T})/\partial \hat{T} > 0$ for $\hat{T}$ just above $\underline{t}\left(\hat{\delta}\right)$, because

$$\frac{\partial}{\partial \hat{T}} \exp\left(\rho \hat{T}\right) \frac{\partial V\left(\hat{\delta},\hat{T}\right)}{\partial \hat{T}}\bigg|_{\hat{T}=\underline{t}(\hat{\delta})} = -\int_{\delta_i \geq \hat{\delta}, \tilde{t}_i = \underline{t}(\hat{\delta})} \alpha Q_N - Q_O - \mu\lambda\eta\left[\frac{1}{\rho - g_O} Q_O - \frac{\alpha - \delta_i}{\rho - g_N} Q_N\right] di.$$

On the right-hand side, the state (μ, Q) is evaluated at $\underline{t}(\hat{\delta})$. Since $\gamma_i < \delta_i$ for all sectors above the threshold, the right-hand side must be strictly positive. This implies that V is strictly increasing in $\hat{T}$ just above $\underline{t}(\hat{\delta})$, so the optimal wait time $\hat{T}$ must be interior. It satisfies the first-order condition

$$0 = -\int_{\delta_i \geq \hat{\delta}} x\left(\mu, q, \gamma_i\right)\left\{\alpha Q_N - Q_O - \mu\lambda\eta\left[\frac{1}{\rho - g_O} Q_O - \frac{\alpha - \delta_i}{\rho - g_N} Q_N\right]\right\} di.$$

Setting $\hat{T} = \underline{t}(\hat{\delta})$ replicates the laissez-faire equilibrium, so this argument establishes that a sandbox policy with $\hat{\delta} > \underline{\delta}$ can strictly improve upon the laissez-faire equilibrium. ■

C.2 Extensions

In this part of the Appendix, we discuss one generalization and two extensions.

C.2.1 Unrestricted Heterogeneity across Sectors

In this section, we briefly describe how our analysis generalizes without the simplifying assumption that α_i and η_i are constant across sectors (but maintaining the "large damages" assumption (3.6)). The basic optimality condition (3.7) describes the solution to the planner's problem, and we can rearrange this condition to observe that there exists a *sector-specific* damage threshold $L_i(\mu, q)$ such that it is optimal to adopt the new technology in sector i iff $\delta_i < L_i(\mu, q)$. This damage threshold satisfies the analogue to (3.8) with α_i and η_i in place of α and η:

$$\frac{L_i(\mu,q)-\alpha_i}{\rho-g_N} = \frac{\alpha_i-\exp(-q)}{\mu\lambda\eta_i} - \frac{\exp(-q)}{\rho-g_O}.$$

The comparative statics of Proposition 7 naturally hold for this sector-specific damage threshold, and they should now be interpreted as comparative statics for the adoption of the new technology *within sector* i as opposed to economy-wide.

Adoption of the new technology is defined exactly as in the benchmark model:

$$X(\mu,q) = \int_0^1 x_i(\mu,q)di. \tag{C.3}$$

It is straightforward to observe that the comparative statics of Proposition 8 apply immediately to sector-level adoption $x_i(\mu,q) = \mathbb{1}[\delta_i < L_i(\mu,q)]$ and, provided that the joint distribution of the sector characteristics $(\delta_i, \alpha_i, \eta_i)$ is well-behaved (e.g., compact support), to overall adoption X. The slope and curvature results of Proposition 9 extend immediately to the sector-specific damage threshold $L_i(\mu,q)$. In particular, the second part of this proposition describes a sense in which adoption is convex *within* each sector, and with sufficient dispersion in damages δ_i this provides a force for convex adoption *across* sectors. If we additionally relax the assumption that damages are large (3.6), the key implication of Proposition 10 holds: With $q(0)$ sufficiently low, it remains optimal to delay adoption in sector i even if g_N increases toward the discount rate ρ. The optimality of gradual adoption is robust to all of these generalizations.

Finally, we note that the same adjustments can be made to our analysis of equilibrium adoption, and similarly for our regulation results: Proposition 14 generalizes in the sense that sector-independent use taxes can restore socially optimal adoption if and only if the laissez-faire and socially optimal orders of adoption coincide. Positive affilitation of damages again suffices for

this, but with heterogeneity in α_i and η_i, it is typically stricter than necessary. By a similar argument as in the proof of Proposition 15, it can again be verified that a regulatory sandbox can generally improve upon the laissez-faire equilibrium.

C.2.2 Heterogeneous α_i

Suppose that η_i and δ_i are constant across sectors, and let F_α denote the smooth distribution function for α_i with support $\left[\underline{\alpha}, \bar{\alpha}\right]$. We maintain the assumption that $\alpha_i \leq \delta$ for each sector i, which requires $\bar{\alpha} \leq \delta$. Making use of the planner's optimality condition (3.7), we observe that there exists a *productivity threshold* $A(\mu, q)$ such that it is optimal to use the new technology in sector i iff $\alpha_i > A(\mu, q)$. Total adoption of the new technology is then the fraction of sectors above the productivity threshold:

$$X(\mu, q) = 1 - F_\alpha(A(\mu, q)).$$

The following proposition characterizes the productivity threshold and is analogous to Proposition 7 in Section 3.3.2.

Proposition C.2.1. *It is socially optimal to use technology N in sector i iff $\alpha_i > A(\mu, q)$, where*

$$A(\mu, q) + \mu\lambda\eta \frac{A(\mu, q) - \delta}{\rho - g_N} = \left(1 + \frac{\mu\lambda\eta}{\rho - g_O}\right) \exp(-q). \tag{C.4}$$

$A(\mu, q)$ (and thus $1 - X(\mu, q)$) is strictly decreasing in q; strictly increasing in g_O and δ; and strictly increasing in λ, η, μ, and g_N provided that $A(\mu, q) < \delta$.

Proof. The characterizing equation (C.4) follows from the planner's optimality condition (3.7). The comparative statics are immediate from (C.4). ■

The analogue of Proposition 8 also holds:

Proposition C.2.2. *For all $t > 0$:*

1. *$X(\mu(t), q(t))$ is decreasing in g_O.*
2. *There exists an earliest time $\bar{t} < \infty$ such that $X(\mu(t), q(t))$ is decreasing in g_N if $t > \bar{t}$. The time $\bar{t}$ is decreasing in g_N.*
3. *Adoption falls to zero as g_N approaches ρ, i.e., $\lim_{g_N \uparrow \rho} X(\mu(t), q(t)) = 0$.*

Comparative statics for the evolution of the productivity threshold $A(\mu, q)$ over time are less tractable than for the damage threshold $L(\mu, q)$ in the benchmark model. The following

proposition provides some guidance about $\dot{A}(\mu,q)$ and $\ddot{A}(\mu,q)$ for the limiting case in which the new and old technologies grow at the same rate.

Proposition C.2.3. *When $g = g_O = g_N$:*

1. *$\dot{A}(\mu,q)$ is negative and increasing in g.*
2. *There exists a posterior $\hat{\mu} \in (0,1/2)$ such that if $\mu \leq \hat{\mu}$, $\ddot{A}(\mu,q)$ is positive.*

Proof. When $g = g_O = g_N$, the characterizing equation (C.4) becomes

$$A(\mu,q) = \frac{1}{1+\frac{\rho-g}{\mu\lambda\eta}}\delta + \exp(-q).$$

The quality gap q is constant since $g = g_O = g_N$. Differentiating in t then yields

$$\dot{A}(\mu,q) = \dot{\mu}\frac{\frac{\rho-g}{\lambda\eta}}{\left(\mu+\frac{\rho-g}{\lambda\eta}\right)^2}\delta,$$
$$\ddot{A}(\mu,q) = \left[\ddot{\mu} - 2\dot{\mu}^2\frac{1}{\mu+\frac{\rho-g}{\lambda\eta}}\right]\frac{\frac{\rho-g}{\lambda\eta}}{\left(\mu+\frac{\rho-g}{\lambda\eta}\right)^2}\delta.$$

Clearly $\dot{A}(\mu,q) < 0$ because $\dot{\mu} < 0$. Using the equations $\dot{\mu} = -\lambda\mu(1-\mu)$ and $\ddot{\mu} = -\lambda\dot{\mu}(1-2\mu)$, we observe that $\ddot{A}(\mu,q) > 0$ iff

$$1 - 2\mu > 2\frac{\mu(1-\mu)}{\mu+\frac{\rho-g}{\lambda\eta}}.$$

This inequality is violated at $\mu = 1/2$, but it is satisfied at $\mu = 0$. Hence there exists a cutoff $\hat{\mu} \in (0,1/2)$ such that it is satisfied for $\mu \leq \hat{\mu}$. ■

Corollary C.2.1. *If $g = g_O = g_N$ and $\mu \in (0,\hat{\mu}]$, adoption is concave over time: $\ddot{X}(\mu,q) < 0$.*

These results imply that learning dynamics favor concave adoption over time when sectors are heterogeneous according to comparative advantage, in contrast to the case with heterogeneous damages considered in the main text.

C.2.3 Constant Damages

In this section, we assess the role of the assumption that post-disaster damages scale with quality Q_N by revisiting the analysis of Section 3.3 under an alternative assumption: Post-disaster damages in sector i are a fixed constant $\Delta_i \geq 0$. In this case, the planner's HJB equations (3.4,

3.5) are still valid, but total damages $D(x)$ are now independent of Q and satisfy

$$D(x) = \int_0^1 x_i \Delta_i di.$$

The planner uses technology N in sector i after the disaster iff $\bar{x}_i = 1$ or $\alpha_i Q_N - \Delta_i > Q_O$. If the disaster strikes when the quality vector is Q and the technology choice in sector i is unconstrained, the planner uses technology O for a time period of length $\bar{T}(Q, g, \Delta_i)$, after which she switches to technology N. The time period $\bar{T}(Q, g, \Delta_i)$ is equal to zero if $\alpha_i Q_N - \Delta_i \geq Q_O$, and otherwise it is the unique solution to the equation

$$\alpha_i Q_N \exp\left(g_N \bar{T}(Q, g, \delta_i)\right) - \Delta_i = Q_O \exp\left(g_O \bar{T}(Q, g, \delta_i)\right).$$

The solution always exists and is unique since $g_N > g_O$.

By the same argument as in Section 3.3.1, technology N is used in sector i before the disaster if the increase in flow output $\alpha_i Q_N - Q_O$ dominates the expected loss due to the disaster. The latter is the product of the expected arrival rate of the disaster $\mu\lambda$, the probability of irreversibility η_i, and the difference between the discounted value of net output when technology choice is unconstrained and when it is constrained to technology N. If the technology choice in sector i is unconstrained after the disaster, the sector produces discounted net output

$$\int_0^{\bar{T}(Q,g,\delta_i)} \exp(-\rho t) \exp(g_O t) Q_O dt + \int_{\bar{T}(Q,g,\delta_i)}^{\infty} \exp(-\rho t) [\alpha_i \exp(g_N t) Q_N - \Delta_i] dt.$$

When constrained to technology N, the sector's discounted net output is

$$\int_0^{\infty} \exp(-\rho t) [\alpha_i \exp(g_N t) Q_N - \Delta_i] dt.$$

We then that it is optimal to use technology N in sector i before the disaster iff

$$\alpha_i Q_N - Q_O > \mu\lambda\eta_i \int_0^{\bar{T}(Q,g,\delta_i)} \exp(-\rho t) \{\exp(g_O t) Q_O - [\alpha_i \exp(g_N t) Q_N - \Delta_i]\} dt. \quad \text{(C.5)}$$

This optimality condition is analogous to (3.7) in the benchmark model, but with three differences. First, we have not explicitly integrated the integral in (C.5) as we have in (3.7). Second, in (C.5) the fixed damages Δ_i replace the quality-dependent damages $Q_N \delta_i$ in (3.7). Finally, with quality-independent damages Δ_i it is always optimal to use technology N at some point

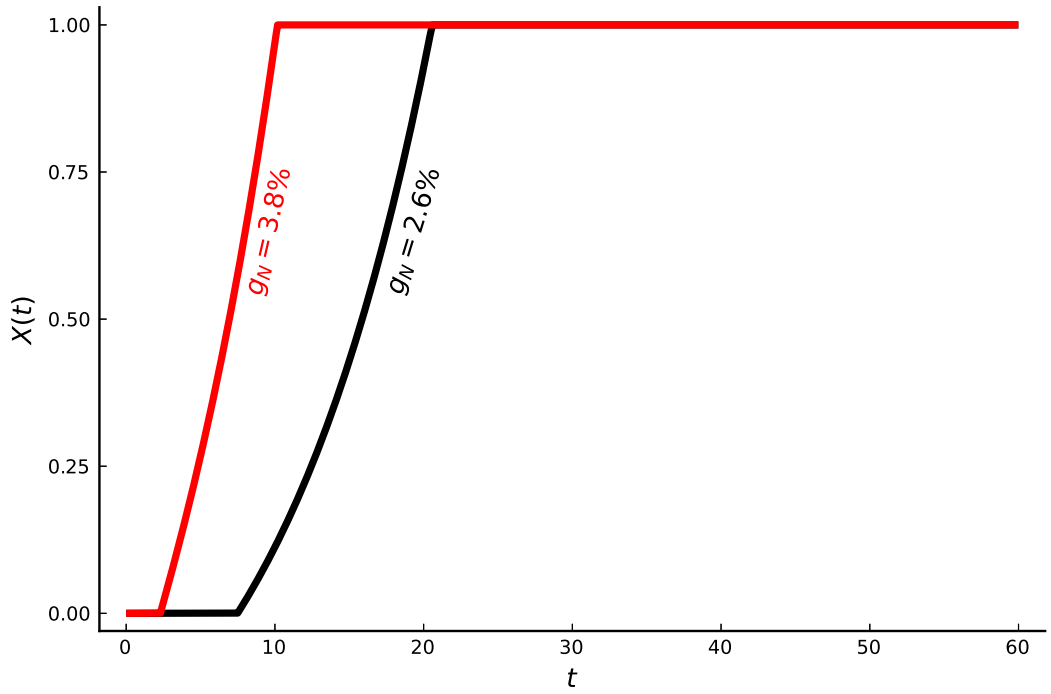

Figure C.2.1: Adoption curves $X(t) \equiv X(\mu(t), q(t))$ for different values of g_N. The parameterization is the same as in Figure 3.1, but with $\underline{\Delta} = 1$ and $\bar{\Delta} = 5$.

after the disaster in sector i: $\bar{T}(Q, g, \delta_i) < \infty$. This contrasts with the benchmark model, in which the assumption $\alpha_i \leq \delta_i$ implies that the planner will always use technology O after the disaster when possible.

Suppose as in Section 3.3.2 that α_i and η_i are constant across sectors. The following proposition is analogous to Proposition 7 in Section 3.3.2. It demonstrates that optimal technology choices can be described using a damage threshold $L(\mu, Q)$ and provides comparative statics.

Proposition C.2.4. *It is socially optimal to use technology N in sector i before the disaster iff $\Delta_i < L(\mu, Q)$, where $L(\mu, Q)$ is the unique solution to the equation*

$$\begin{aligned} &\alpha Q_N - Q_O \\ &\quad = \mu\lambda\eta \int_0^{\bar{T}(Q,g,L(\mu,Q))} \exp(-\rho t)\{\exp(g_O t) Q_O - [\alpha \exp(g_N t) Q_N - L(\mu, Q)]\}\, dt. \end{aligned} \tag{C.6}$$

$L(\mu, Q)$ (and thus $X(\mu, Q)$) is strictly increasing in α, Q_N, and g_N and strictly decreasing in g_O, λ, μ, and Q_O.

We omit the proof details, because the argument is almost identical to the proof of Proposition C.3.1 in Appendix C.3.2. This proposition demonstrates that, when damages from technology N do not scale with its quality Q_N, optimal adoption is increasing in the growth rate g_N. This

contrasts with the corresponding result in Proposition 7, demonstrating that the assumption of proportional damages has significant implications for optimal adoption. We argue that many of the conjectured dangers of (generative) AI more naturally correspond to the case in which damages scale with the capabilities (quality) of rapidly improving models.

We illustrate these results in Figure C.2.1. We modify the calibration of Figure 3.1 only by assuming constant damages Δ_i uniformly distributed over $\left[\underline{\Delta}, \bar{\Delta}\right]$, where $\underline{\Delta} = 1$ and $\bar{\Delta} = 5$, and by initializing $Q(0) = (1, 1)$. As a result, the initial value of the damages in each sector is the same as in the quantitive example in the main text, as is the quality gap $q(0) = 0$. Consistent with Proposition C.2.4, we observe that adoption is increasing in the growth rate g_N. Moreover, adoption is much faster than in Figure 3.1 because (potential) damages do not increase over time as technology N improves.

C.3 Analysis with Unrestricted Damages

In this part of the Appendix, we analyze the benchmark model without assuming large damages (3.6).

C.3.1 Socially Optimal Technology Choice

As described in the main text, the planner uses technology N after the disaster iff $\bar{x}_i = 1$ or $(\alpha_i - \gamma_i) Q_N > Q_O$. Letting $q = \log(Q_N/Q_O)$ denote the log quality gap between the technologies, we can equivalently define a threshold gap q_i such that the planner uses technology N after the disaster iff $\bar{x}_i = 1$ or $q \geq q_i$:

$$q_i = \begin{cases} -\log(\alpha_i - \delta_i) & \text{if } \alpha_i > \delta_i, \\ \infty & \text{else.} \end{cases} \tag{C.7}$$

At the onset of the disaster, if $q < q_i$ the planner optimally reverts to using technology O in sector i if possible. If $q_i < \infty$, the planner eventually uses technology N again when it attains a sufficiently large lead over technology O.

With this characterization, we can directly integrate the post-disaster HJB equation (3.5) and take expectations with respect to $\bar{x}$:

$$\mathbb{E}[W(\bar{x}, Q) | x] = \int_0^1 (1 - x_i \eta_i) \left\{ \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O \right. \\ \left. + \exp\left(-\frac{\rho - g_N}{g_N - g_O} (q_i - q)_+ \right) \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N \right\} + x_i \eta_i \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N di.$$

Here we use the notation $(q_i - q)_+ = \max\{q_i - q, 0\}$. Considering the planner's problem before the disaster (3.4), we observe that it is optimal to use technology N in sector i iff

$$\alpha_i Q_N - Q_O > \mu \lambda \eta_i \left\{ \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O \right. \\ \left. - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N \right\}. \tag{C.8}$$

This optimality condition differs from (3.7) because the discounted future net output from using technology O at the time of the disaster now accounts for the possibility that technology N is used after the quality gap q exceeds q_i.

C.3.2 Comparative Statics for Socially Optimal Adoption

Suppose as in Section 3.3.2 that α_i and η_i are constant across sectors, but make no assumption about the ranking between δ_i and α. Let $\bar{q}(\delta_i) = q_i$ denote the quality gap above which it is optimal to use technology N in sector i after the disaster (C.7), making explicit the dependence on δ_i. The following proposition shows that optimal technology choices can be described using a damage threshold $L(\mu, q)$ and provides comparative statics, generalizing Proposition 7 from Section 3.3.2.

Proposition C.3.1. *It is socially optimal to use technology N in sector i before the disaster iff $\delta_i < L(\mu, q)$, where $L(\mu, q)$ is the unique solution to the equation*

$$\begin{aligned}\alpha - \exp(-q) = \mu\lambda\eta \Bigg\{ &\left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O}\left(\bar{q}\left(L(\mu,q)\right) - q\right)_+\right)\right]\frac{1}{\rho - g_O}\exp(-q) \\ &- \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O}\left(\bar{q}\left(L(\mu,q)\right) - q\right)_+\right)\right]\frac{\alpha - L(\mu,q)}{\rho - g_N}\Bigg\}.\end{aligned} \tag{C.9}$$

$L(\mu, q)$ (and thus $X(\mu, q)$) is strictly increasing in α and q and strictly decreasing in g_O, λ, and μ. It is strictly decreasing in g_N if $L(\mu, q) > \alpha$ and strictly increasing in g_N if $L(\mu, q) < \alpha$.

Proof. Throughout the proof, we suppress the arguments of the damage threshold $L(\mu, q)$ to simplify notation. The results described in the proposition are easier to prove if we re-write the discounted values on the right-hand side of (C.9) as integrals over time. To do this, given a quality gap q, let $\bar{T}(q, g, \delta)$ denote the length of time after the disaster during which it is optimal to use technology O instead of technology N in a sector with damages δ:

$$\bar{T}(q, g, \delta) = \begin{cases} \max\left\{\frac{-\log(\alpha - \delta) - q}{g_N - g_O}, 0\right\} & \text{if } \alpha > \delta, \\ \infty & \text{else.} \end{cases}$$

If the sector is not constrained to technology N, its discounted net output after the disaster is

$$\int_0^{\bar{T}(q,g,\delta)} \exp(-\rho t)\exp(g_O t) Q_O dt + \int_{\bar{T}(q,g,\delta)}^{\infty} \exp(-\rho t)\exp(g_N t)(\alpha - \delta) Q_N dt. \tag{C.10}$$

Similarly, its discounted net output when constrained to technology N is

$$\int_0^{\infty} \exp(-\rho t)\exp(g_N t)(\alpha - \delta) Q_N dt. \tag{C.11}$$

The bracketed term in (C.9) is the difference between the previous two terms above for the

marginal sector (with $\delta = L$), divided by Q_N. The right-hand side of (C.9) can then be written

$$\text{RHS} = \mu\lambda\eta \int_0^{\bar{T}(q,g,L)} \exp(-\rho t)[\exp(g_O t)\exp(-q) - \exp(g_N t)(\alpha - L)]\,dt.$$

We first demonstrate that, when $\alpha > \exp(-q)$ so that technology N is more productive than technology O, there always exists a unique solution L to (C.9). We observe that RHS is continuous in L, equals zero when $L \leq \alpha - \exp(-q)$, and limits to infinity as $L \to \infty$. Moreover, RHS is strictly increasing in L when $L > \alpha - \exp(-q)$: This condition implies $\bar{T}(q,g,L) > 0$, and we can differentiate RHS to find

$$\begin{aligned}
\frac{\partial \text{RHS}}{\partial L} &= \mu\lambda\eta \exp\left(-\rho\bar{T}\right)\left[\exp\left(g_O\bar{T}\right)\exp(-q) - \exp\left(g_N\bar{T}\right)(\alpha - L)\right]\frac{\partial \bar{T}}{\partial L} \\
&\quad + \mu\lambda\eta \int_0^{\bar{T}(q,g,L)} \exp(-\rho t)\exp(g_N t)\,dt \\
&= \mu\lambda\eta \int_0^{\bar{T}(q,g,L)} \exp(-\rho t)\exp(g_N t)\,dt \\
&> 0.
\end{aligned}$$

Note that the second equality holds by the Envelope Theorem: $\bar{T}$ maximizes the discounted net output from the marginal sector after the disaster, assuming its technology choice is unconstrained. As a result RHS does not vary locally with respect to $\bar{T}$ ($\partial \text{RHS}/\partial \bar{T} = 0$). Given these properties of RHS, the Intermediate Value Theorem guarantees a unique solution L to (C.9) when $\alpha > \exp(-q)$. Moreover, it follows from the optimality condition (C.8) that it is socially optimal to use technology N in sector i before the disaster iff $\delta_i < L(\mu, q)$.

The comparative statics for the damage threshold L follow from the Implicit Function Theorem. Holding L fixed, we immediately observe that RHS is decreasing in α and increasing in μ, λ, and η. Differentiating with respect to q, g_O, and g_N yields

$$\begin{aligned}
\frac{\partial \text{RHS}}{\partial q} &= -\mu\lambda\eta \int_0^{\bar{T}} \exp(-\rho t)\exp(g_O t)\exp(-q)\,dt, \\
\frac{\partial \text{RHS}}{\partial g_O} &= \mu\lambda\eta \int_0^{\bar{T}} \exp(-\rho t)\exp(g_O t)\exp(-q)\,t\,dt, \\
\frac{\partial \text{RHS}}{\partial g_N} &= -\mu\lambda\eta \int_0^{\bar{T}} \exp(-\rho t)\exp(g_N t)(\alpha - L(\mu, q))\,dt.
\end{aligned}$$

These expressions imply that RHS is decreasing in q, increasing in g_O, and decreasing (increasing) in g_N iff $\alpha > (<) L(\mu, q)$. Collecting these results, the Implicit Function Theorem delivers the comparative statics stated in the proposition. ■

The proposition demonstrates that almost all comparative statics from Proposition 7 hold without the assumption that social damages always exceed output from technology N after the disaster ($\alpha_i \leq \delta_i$). However, the comparative static with respect to g_N is sensitive to this assumption. When damages in the marginal sector exceed output ($L(\mu, q) > \alpha$), the damage threshold is decreasing in g_N as in Proposition 7. When damages in the marginal sector are below output ($L(\mu, q) < \alpha$), the damage threshold is instead increasing in g_N.

The following proposition generalizes Proposition 8 to provide full comparative statics for adoption with respect to the growth rates g_O and g_N, including both the direct effects described in Proposition C.3.1 and the indirect effects through the state $(\mu(t), q(t))$.

Proposition C.3.2. *Suppose α_i and η_i are constant across sectors. For all t with $L(\mu(t), q(t)) < \alpha$:*

1. *$X(\mu(t), q(t))$ is decreasing in g_O.*
2. *$X(\mu(t), q(t))$ is increasing in g_N.*
3. *If $q(0)$ is sufficiently low and $X(\mu(t), q(t)) < F(\alpha)$, $X(\mu(t), q(t))$ is bounded strictly below $F(\alpha)$ as g_N approaches ρ, i.e., $\lim_{g_N \uparrow \rho} X(\mu(t), q(t)) < F(\alpha)$.*

Proof. The first two results follow from Proposition C.3.1 after noting that the damage threshold L is increasing in the quality gap q, and in turn the quality gap $q(t)$ at time t is decreasing in g_O and increasing in g_N. The final result of the proposition follows by taking the limit $g_N \uparrow \rho$ in (C.9). More precisely, recall from the optimality condition (C.8) that it is socially optimal to use technology N in sector i at time t before the disaster only if

$$\begin{aligned} \alpha - \exp(-q(t)) \geq \mu(t)\lambda\eta \Bigg\{ \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (q_i - q(t))_+ \right) \right] \frac{1}{\rho - g_O} \exp(-q(t)) \\ - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (q_i - q(t))_+ \right) \right] \frac{\alpha - \delta_i}{\rho - g_N} \Bigg\}. \end{aligned} \tag{C.12}$$

Just as in the proof of Proposition C.3.1, the right-hand side of this inequality can be written

$$\text{RHS} = \mu(t)\lambda\eta \int_0^{\bar{T}(q(t), g, \delta_i)} \exp(-\rho s) \left[\exp(g_O s) \exp(-q(t)) - \exp(g_N s)(\alpha - \delta_i) \right] ds,$$

For $g_N < \rho$, this function is continuous and nondecreasing in δ_i. At $\delta_i = \alpha$, it takes the value

$$\text{RHS} = \frac{\mu(t)\lambda\eta}{\rho - g_O}\exp(-q(t)).$$

In the limit as $g_N \uparrow \rho$, (C.12) implies that it is socially optimal to use technology N in sector i with damages $\delta_i = \alpha$ only if

$$\alpha\exp(q(0) + (\rho - g_O)t) - 1 \geq \frac{\mu(t)\lambda\eta}{\rho - g_O}.$$

For t fixed, this inequality is violated for $q(0)$ sufficiently small (but potentially still positive). Since RHS is continuous in $\delta_i \leq \alpha$, this implies that $L(\mu(t), q(t))$ must remain bounded strictly below α. If $X(\mu(t), q(t)) < F(\alpha)$, this immediately implies that $X(\mu(t), q(t))$ is bounded strictly below $F(\alpha)$. ■

Notably, in the limit $g_N \uparrow \rho$ adoption does not tend to either of the extreme values 0 or $F(\alpha)$, in contrast to the corresponding result in Proposition 8. This holds because, for any sector i with $\delta_i < \alpha$, the discounted net output after the disaster tends to infinity as $g_N \uparrow \rho$ regardless of whether the sector is constrained to use technology N after the disaster. However, the *difference* between the discounted net output when unconstrained and the discounted net output when constrained tends to a finite limit. Socially optimal technology choices before the disaster depend on this difference (see C.8), so provided that δ_i is sufficiently close to α and the initial quality gap $q(0)$ sufficiently low, it can remain optimal to delay using technology N in sector i before the disaster even when $g_N \uparrow \rho$.

We illustrate these results in Figure C.3.1 by depicting adoption curves for a stylized parameterization of the model. We modify the calibration of Figure 3.1 only by assuming that the distribution of damages δ_i is uniform over $[0, 5]$ instead of $[1, 5]$. Technology choices for sectors with $\delta_i \in [1, 5]$ are exactly as in Section 3.3, and since these sectors comprise $5/6$ of all sectors in this calibration, the adoption curves in Figure C.3.1 when $X(t) \geq 5/6$ are identical to the adoption curves in Figure 3.1.

When instead $X(t) \in (0, 1/6)$, the sectors adopting technology N produce positive net output after the disaster, so the analysis in this appendix becomes relevant. In this region, we observe that adoption is increasing in g_N, consistent with Proposition C.3.2.

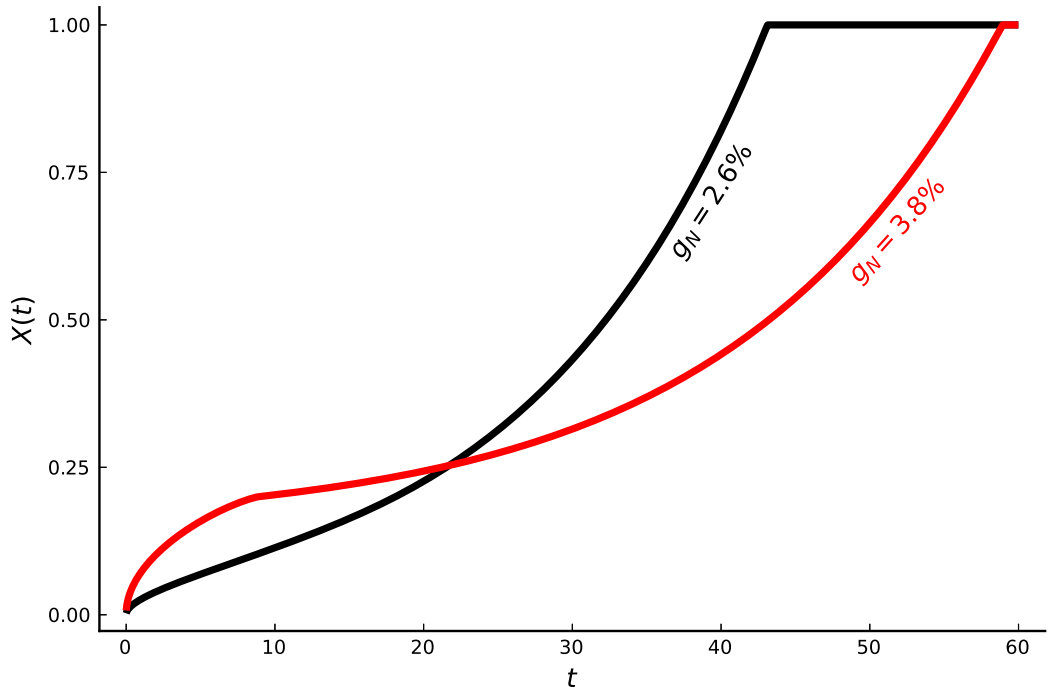

Figure C.3.1: Adoption curves $X(t)$ for different values of g_N. The parameterization is the same as in Figure 3.1, but with $\underline{\delta} = 0$.

C.3.3 Equilibrium Technology Choice

Using the same derivations as for the optimal technology choice, firm i uses technology N after the disaster iff $\bar{x}_i = 1$ or $q \geq \tilde{q}_i$, where

$$\tilde{q}_i = \begin{cases} -\log(\alpha_i - \gamma_i) & \text{if } \alpha_i > \gamma_i, \\ \infty & \text{else.} \end{cases}$$

Note that $\tilde{q}_i \leq q_i$ since $\gamma_i \leq \delta_i$. This implies that the private firm returns to using technology N more quickly after the disaster than the planner. Integrating the firm's post-disaster HJB equation (3.10) and taking expectations with respect to $\bar{x}_i$ yields

$$\mathbb{E}[\Phi_i(\bar{x}_i, Q)|x_i] = (1 - x_i\eta_i)\left\{\left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O}(\tilde{q}_i - q)_+\right)\right]\frac{1}{\rho - g_O}Q_O \right.$$
$$\left. + \exp\left(-\frac{\rho - g_N}{g_N - g_O}(\tilde{q}_i - q)_+\right)\frac{\alpha_i - \gamma_i}{\rho - g_N}Q_N\right\} + x_i\eta_i\frac{\alpha_i - \gamma_i}{\rho - g_N}Q_N.$$

It is then privately optimal to use technology N in sector i before the disaster iff

$$\alpha_i Q_N - Q_O > \mu\lambda\eta_i \left\{ \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O \right. \\ \left. - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{\alpha_i - \gamma_i}{\rho - g_N} Q_N \right\}. \tag{C.13}$$

We observe two differences between this condition and the planner's optimality condition (C.8), First, as in the main text, private damages γ_i appear in (C.13) instead of the social damages that appear in (C.8). Second, the firm begins using technology N more quickly after the disaster than the planner ($\tilde{q}_i \leq q_i$). Both effects tend to reduce the net private cost of irreversibility and incentivize the firm to use technology N more often than the planner before the disaster.

Lemma C.3.1. *If the social planner uses technology N in sector i in state (μ, Q) before the disaster, then so does firm i.*

Proof. The statement holds provided that firm i's opportunity cost to using technology N instead of technology O at the time of the disaster is smaller than the planner's opportunity cost:

$$\left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{\alpha_i - \gamma_i}{\rho - g_N} Q_N \\ \leq \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N.$$

Replacing γ_i with δ_i yields the intermediate inequality

$$\left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{\alpha_i - \gamma_i}{\rho - g_N} Q_N \\ \leq \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N.$$

Optimality of q_i in the planner's problem after the disaster yields the remaining inequality

$$\left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (\tilde{q}_i - q)_+ \right) \right] \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N \\ \leq \left[1 - \exp\left(-\frac{\rho - g_O}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{1}{\rho - g_O} Q_O - \left[1 - \exp\left(-\frac{\rho - g_N}{g_N - g_O} (q_i - q)_+ \right) \right] \frac{\alpha_i - \delta_i}{\rho - g_N} Q_N.$$

■

C.4 Other Issues in Regulation

In this part of the Appendix, we explore second-best tax regulation schemes when private and social damages are not positively affiliated, and we provide additional details about optimal regulatory sandboxes and discuss their advantages relative to sector-independent taxes.

C.4.1 Second-Best Tax Regulation

Aside from the special case in which social and private damages are positively affiliated, a sector-independent tax cannot implement the optimal technology choices in equilibrium. More generally, use taxes can allow the planner to improve upon laissez-faire technology choices even when optimal ones cannot be implemented. Suppose as in Proposition 14 that α_i and η_i are constant across sectors, but make no assumptions on the joint distribution of δ_i and γ_i. In each state (μ, Q) before the disaster, the planner chooses the use tax $\tau(\mu, Q)$ to maximize output less the expected discounted social cost from the disaster:

$$\max_{\tau} \int_0^1 \left\{ (1 - x(\mu, Q, \gamma_i, \tau)) \left[Q_O + \mu\lambda\eta \frac{1}{\rho - g_O} Q_O \right] + x(\mu, Q, \gamma_i, \tau) \left[\alpha Q_N + \mu\lambda\eta \frac{\alpha - \delta_i}{\rho - g_N} Q_N \right] \right\} di.$$

Here $x(\mu, Q, \gamma_i, \tau)$ describes the equilibrium technology choice for firm i when subject to the tax:

$$x(\mu, Q, \gamma_i, \tau) = \begin{cases} 1 & \text{if } \alpha Q_N - Q_O - \tau > \mu\lambda\eta \left[\frac{1}{\rho - g_O} Q_O - \frac{\alpha - \gamma_i}{\rho - g_N} Q_N \right], \\ 0 & \text{else.} \end{cases}$$

Firms adopt technology N in order of increasing γ_i, so we can equivalently assume that the planner selects a private damage threshold $\hat{L}(\mu, q)$ such that firm i uses technology N iff $\gamma_i < \hat{L}(\mu, q)$. The optimal threshold trades off flow consumption against the expected social cost of the disaster. When interior, it satisfies

$$\alpha Q_N - Q_O = \mu\lambda\eta \left[\frac{1}{\rho - g_O} Q_O - \frac{\alpha - \bar{\delta}\left(\hat{L}(\mu, q)\right)}{\rho - g_N} Q_N \right]. \tag{C.14}$$

Here $\bar{\delta}(\gamma) = \mathbb{E}[\delta_i | \gamma_i = \gamma]$ is the average social damages across all firms with private damages γ. The optimality condition (C.14) is analogous to the original optimality condition (3.7), but it replaces a single sector's social damages δ_i with the expectation $\bar{\delta}(\gamma)$. The planner's problem is concave iff $\bar{\delta}(\gamma)$ is increasing, in which case an interior solution can be optimal. If, for example,

$\bar{\delta}(\gamma)$ is decreasing, then the planner cannot incentivize sectors with low social damages to use technology N while sectors with high social damages use technology O. As a result, the planner chooses $\hat{L}(\mu, q) = 0$ (no use of N) or $\hat{L}(\mu, q) = \infty$ (full use of N). The latter is optimal when

$$\alpha Q_N - Q_O > \mu\lambda\eta\left[\frac{1}{\rho - g_O}Q_O - \frac{\alpha - \mathbb{E}[\delta_i]}{\rho - g_N}Q_N\right].$$

C.4.2 Analysis of Sandbox Regulation

Proposition 15 in the main text demonstrates that it is generally optimal for the planner to implement a regulatory sandbox with a strictly positive wait time $\hat{T}$. The optimal wait time $\hat{T}$ must satisfy the following interior first-order condition, which is derived in the proof of the proposition in Appendix C.1:

$$0 = -\int_{\delta_i \geq \hat{\delta}} x(\mu, q, \gamma_i)\left\{\alpha Q_N - Q_O - \mu\lambda\eta\left[\frac{1}{\rho - g_O}Q_O - \frac{\alpha - \delta_i}{\rho - g_N}Q_N\right]\right\} di. \tag{C.15}$$

Here the state (μ, Q) is evaluated at the optimal time $\hat{T}$, and $x(\mu, q, \gamma_i) = 1$ iff sector i would use technology N in the laissez-faire equilibrium. Two forces determine the optimal wait time $\hat{T}$: If sector i is above the threshold ($\delta_i \geq \hat{\delta}$) and would inefficiently use technology i at time $\hat{T}$, its laissez-faire technology choice would decrease social welfare, favoring a longer wait time:

$$x(\mu, q, \gamma_i)\left\{\alpha Q_N - Q_O - \mu\lambda\eta\left[\frac{1}{\rho - g_O}Q_O - \frac{\alpha - \delta_i}{\rho - g_N}Q_N\right]\right\} < 0.$$

If sector i would instead efficiently use technology i at time $\hat{T}$, its laissez-faire technology choice would increase social welfare, favoring a shorter wait time.[1]

We can similarly derive the following interior first-order condition for the optimal threshold $\hat{\delta}$, keeping $\hat{T}$ fixed:

$$0 = \int_0^{\hat{T}} \exp(-\rho t)\int_{\delta_i = \hat{\delta}} x(\mu, q, \gamma_i)\left\{\alpha Q_N - Q_O - \mu\lambda\eta\left[\frac{1}{\rho - g_O}Q_O - \frac{\alpha - \hat{\delta}}{\rho - g_N}Q_N\right]\right\} dt.$$

If the threshold $\hat{\delta}$ is too high, a large fraction of sectors i face no restrictions on their technology choices, and they subtract too much from social welfare between $t = 0$ and $t = \hat{T}$ as they begin using the new technology too quickly. If $\hat{\delta}$ is too low, then too many sectors i are forced to use technology O between $t = 0$ and $t = \hat{T}$, foregoing the benefits of using technology N in these

[1] As this intuition suggests, it is straightforward to verify that, under the assumptions of Proposition 15, the optimal wait time $\hat{T}$ is nondecreasing in $\hat{\delta}$.

sectors when it is efficient to do so. This analysis demonstrates that the optimal parameters $\left(\hat{\delta}, \hat{T}\right)$ are chosen to resolve a trade-off between restricting early use of the new technology in sectors where expected damages are large, while allowing broad use later as the probability of a disaster falls and the quality gap grows.

We conclude this section by observing that regulatory sandboxes are likely to dominate (or complement) sector-independent taxes when the order of adoption differs substantially between the equilibrium and social optimum. For example, suppose that private and social damages are negatively affiliated: $\gamma_i = \kappa(\delta_i)$, where κ is strictly decreasing. Then Proposition 14 implies that, for any sector-independent tax $\tau(\mu, Q)$, the order in which sectors adopt the new technology in equilibrium is exactly the opposite of the optimal order. Moreover, the analysis in Appendix C.4.1 implies that the optimal sector-independent tax is such that there exists a time $\hat{T}$ before which no sector uses technology N and after which every sector uses technology N. This time is characterized by the equation

$$\alpha Q_N\left(\hat{T}\right) - Q_O\left(\hat{T}\right) = \mu\left(\hat{T}\right)\lambda\eta\left[\frac{1}{\rho - g_O}Q_O\left(\hat{T}\right) - \frac{\alpha - \mathbb{E}[\delta_i]}{\rho - g_N}Q_N\left(\hat{T}\right)\right].$$

These technology choices can also be implemented using the sandbox policy with threshold $\hat{\delta} = \underline{\delta}$ and wait time $\hat{T}$. Hence a regulatory sandbox can achieve weakly greater social welfare than any sector-independent tax when the misalignment in the order of adoption between the equilibrium and the social optimum is severe.

References

Abaluck, J., & Gruber, J. (2011). Heterogeneity in choice inconsistencies among the elderly: Evidence from prescription drug plan choice. *American Economic Review*, *101*(3), 377–381.

Acemoglu, D. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *The quarterly journal of economics*, *113*(4), 1055–1089.

Acemoglu, D. (2002). Directed technical change. *The review of economic studies*, *69*(4), 781–809.

Acemoglu, D. (2011). *Diversity and technological progress* (tech. rep.). National Bureau of Economic Research.

Acemoglu, D. (2021). Ai's future doesn't have to be dystopian.' *Boston Review*, 23–29.

Acemoglu, D., Aghion, P., Bursztyn, L., & Hemous, D. (2012). The environment and directed technical change. *American economic review*, *102*(1), 131–166.

Acemoglu, D., Aghion, P., & Zilibotti, F. (2006). Distance to frontier, selection, and economic growth. *Journal of the European Economic association*, *4*(1), 37–74.

Acemoglu, D., Akcigit, U., Hanley, D., & Kerr, W. (2016). Transition to clean technology. *Journal of political economy*, *124*(1), 52–104.

Acemoglu, D., Antràs, P., & Helpman, E. (2007). Contracts and technology adoption. *American Economic Review*, *97*(3), 916–943.

Acemoglu, D., & Cao, D. (2015). Innovation by entrants and incumbents. *Journal of Economic Theory*, *157*, 255–294.

Acquatella, A., Ericson, K. M., & Starc, A. (2023). Lagged-Price Reimbursement Contracts: The Impact of Medicare Part B on Pharmaceutical Price Growth.

Agha, L., Kim, S., & Li, D. (2022). Insurance design and pharmaceutical innovation. *American Economic Review: Insights*, *4*(2), 191–208.

Aghion, P., Bergeaud, A., De Ridder, M., & Van Reenen, J. (2024). Lost in transition: Financial barriers to green growth.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., & Howitt, P. (2005). Competition and innovation: An inverted-u relationship. *The quarterly journal of economics*, *120*(2), 701–728.

Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., & Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, *124*(1), 1–51.

Aghion, P., Harris, C., Howitt, P., & Vickers, J. (2001). Competition, imitation and growth with step-by-step innovation. *The Review of Economic Studies*, *68*(3), 467–492.

Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, *60*(2), 323–351.

Agrawal, A., Gans, J., & Goldfarb, A. (2019). Economic policy for artificial intelligence. *Innovation policy and the economy*, *19*(1), 139–159.

Aguirregabiria, V., & Mira, P. (2010). Stability of Equilibria and Properties of Sequential Estimators: Problems and Solutions. *Manuscript. University of Toronto*.

Ahmed, N., Wahed, M., & Thompson, N. C. (2023). The growing influence of industry in ai research. *Science*, *379*(6635), 884–886.

Akcigit, U., & Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, *13*(1), 257–298.

Akcigit, U., & Ates, S. T. (2023). What happened to us business dynamism? *Journal of Political Economy*, *131*(8), 2059–2124.

Akcigit, U., & Goldschlag, N. (2023). *Where have all the" creative talents" gone? employment dynamics of us inventors* (tech. rep.). National Bureau of Economic Research.

Akcigit, U., & Kerr, W. R. (2018). Growth through heterogeneous innovations. *Journal of Political Economy*, *126*(4), 1374–1443.

Alberer, M., Gnad-Vogt, U., Hong, H. S., Mehr, K. T., Backert, L., Finak, G., Gottardo, R., Bica, M. A., Garofano, A., Koch, S. D., et al. (2017). Safety and immunogenicity of a mrna rabies vaccine in healthy adults: An open-label, non-randomised, prospective, first-in-human phase 1 clinical trial. *The Lancet*, *390*(10101), 1511–1520.

Alchian, A. A., & Demsetz, H. (1972). Production, information costs, and economic organization. *The American economic review*, *62*(5), 777–795.

Alferov, Kroemer, & Kilby. (2000). The Nobel Prize in Physics 2000 [Accessed: 2024-9-15].

American Cancer Society Cancer Action Network. (2024, May). Survivor Views: Majority of Cancer Patients & Survivors Have or Expect to Have Medical Debt [Accessed: 2024-7-15].

Arcidiacono, P., Bayer, P., Bugni, F. A., & James, J. (2013, December). Approximating high-dimensional dynamic models: Sieve value function iteration. In *Structural Econometric Models* (pp. 45–95). Emerald Group Publishing Limited.

Arcidiacono, P., & Ellickson, P. (2011). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics*, *3*(1), 363–394.

Arcidiacono, P., & Miller, R. (2019). Nonstationary dynamic models with finite dependence. *Quantitative Economics*, *10*(3), 853–890.

Argente, D., Baslandze, S., Hanley, D., & Moreira, S. (2020). Patents to products: Product innovation and firm dynamics.

Arora, A., Belenzon, S., Cioaca, L., Sheer, L., Shin, H. M. (, & Shvadron, D. (2024). DISCERN 2: Duke Innovation & SCientific Enterprises Research Network. https://doi.org/10.5281/zenodo.13619821

Arrow, K. (1962). Economic Welfare and the Allocation of Resources for Invention. *The Rate and Direction of Inventive Activity: Economic and Social Factors*.

Arthur, B. (2009). *The nature of technology: What it is and how it evolves*. Simon; Schuster.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, *99*(394), 116–131.

Aryal, G., Ciliberto, F., Farmer, L. E., & Khmelnitskaya, E. (2022). Valuing pharmaceutical drug innovations.

Aschenbrenner, L. (2020). *Existential risk and growth* (tech. rep.). GPI Working Paper.

Bajari, P., Benkard, L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, *75*(5), 1331–1370.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, *121*(3), 823–866.

Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review*, *99*(1), 243–264.

Benhabib, J., Perla, J., & Tonetti, C. (2021). Reconciling models of diffusion and innovation: A theory of the productivity distribution and technology frontier. *Econometrica*, *89*(5), 2261–2301.

Benkard, C. L., Jeziorski, P., & Weintraub, G. Y. (2015). Oblivious equilibrium for concentrated industries. *The Rand journal of economics*, *46*(4), 671–708.

Berry, S., & Jia, P. (2010). Tracing the Woes: An Empirical Analysis of the Airline Industry. *American Economic Journal: Microeconomics*, *2*(3), 1–43.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, *63*(4), 841–890.

Bloom, N., Schankerman, M., & Van Reenen, J. (2013a). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica*, *81*(4), 1347–1393.

Bloom, N., Schankerman, M., & Van Reenen, J. (2013b). Identifying technology spillovers and product market rivalry. *Econometrica*, *81*(4), 1347–1393.

Bodéré, P. (2023). Dynamic spatial competition in early education: An equilibrium analysis of the preschool market in pennsylvania. *Job Market Paper*.

Boshuizen, J., & Peeper, D. (2020). Rational cancer treatment combinations: An urgent clinical need. *Molecular Cell*, *78*(6), 1002–1018.

Bovenberg, A. L., & Smulders, S. (1995). Environmental quality and pollution-augmenting technological change in a two-sector endogenous growth model. *Journal of public Economics*, *57*(3), 369–391.

Brekke, K. R., Dalen, D. M., & Straume, O. R. (2023). Taking the competitor's pill: when combination therapies enter pharmaceutical markets.

Brock, W. A., & Durlauf, S. N. (1999). A formal model of theory choice in science. *Economic theory*, *14*, 113–130.

Brynjolfsson, E. (2023). The turing trap: The promise & peril of human-like artificial intelligence. In *Augmented education in the global age* (pp. 103–116). Routledge.

Budish, E., Roin, B. N., & Williams, H. (2015). Do firms underinvest in long-term research? evidence from cancer clinical trials. *American Economic Review*, *105*(7), 2044–2085.

Bukowski, K., Kciuk, M., & Kontek, R. (2020). Mechanisms of multidrug resistance in cancer chemotherapy. *International Journal of Molecular Sciences*, *21*(9), 3233.

Califf, R. M., Zarin, D. A., Kramer, J. M., Sherman, R. E., Aberle, L. H., & Tasneem, A. (2012). Characteristics of clinical trials registered in ClinicalTrials.gov, 2007-2010. *JAMA: The Journal of the American Medical Association*, *307*(17), 1838–1847.

Cao, S., & Chatterjee, C. (2023). Equilibrium effects of pharmaceutical bundling: Evidence from India.

Capps, C., Dranove, D., & Satterthwaite, M. (2003). Competition and market power in option demand markets. *The RAND Journal of Economics*, *34*(4), 737–763.

Carey, C., Daly, M., & Li, J. (2024). Nothing for Something: Marketing Cancer Drugs to Physicians Increases Prescribing Without Improving Mortality.

CDC. (2024, May). Leading Causes of Death [Accessed: 2024-7-10].

Christensen, C. M. (1997). *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business Review Press.

Chu, E., & DeVita, V. (2019). *Physicians' cancer chemotherapy drug manual 2019*. Jones & Bartlett Learning.

Chu, E., & DeVita, V. T., Jr. (2023, December). *Physicians' cancer chemotherapy drug manual 2024* (24th ed.). Jones & Bartlett.

Chui, M., Roberts, R., Yee, L., Hazan, E., Singla, A., Smaje, K., Sukharevsky, A., & Zemmel, R. (2023). *The economic potential of generative ai* (tech. rep.). McKinsey & Company.

Clancy, M. S. (2018). Combinations of technology in us patents, 1926–2009: A weakening base for future innovation? *Economics of Innovation and new Technology*, *27*(8), 770–785.

Coase, R. H. (1960). The problem of social cost. *The Journal of Law and Economics*, *3*, 1–44. https://doi.org/10.1086/466560

Coase, R. (1937). The nature of the firm. *Economica*, *4*(16), 386–405.

Comin, D., & Mestieri, M. (2014). Technology diffusion: Measurement, causes, and consequences. In *Handbook of economic growth* (pp. 565–622, Vol. 2). Elsevier.

Conlon, C. T., & Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*.

Conlon, C. T., & Gortmaker, J. (2023). Incorporating micro data into differentiated products demand estimation with PyBLP. *SSRN Electronic Journal*, *2023*.

Connor, R. (2024). *UpToDate* (W. Kluwer, Ed.).

Conrad, R., & Lutter, R. (2019). *Generic Competition and Drug Prices: New Evidence Linking Greater Generic Competition and Lower Generic Drug Prices* (tech. rep. No. FDA).

Conti, R., Athey, S., Frank, R., & Gruber, J. (2020, December). Generic drug repurposing for public health and national security: COVID-19 and beyond [Title of the publication associated with this dataset: Forefront Group].

Cuddy, E. (2021). Competition and collusion in the generic drug market.

Cunningham, C., Ederer, F., & Ma, S. (2021). Killer acquisitions. *Journal of political economy*, *129*(3), 649–702.

Cutler, D. (2008). Are we finally winning the war on cancer? *The Journal of Economic Perspectives*, *22*(4), 3–26.

Dafny, L., Ho, K., & Kong, E. (2023). How do copayment coupons affect branded drug prices and quantities purchased?

de Gramont, A., Figer, A., Seymour, M., Homerin, M., Hmissi, A., Cassidy, J., Boni, C., Cortes-Funes, H., Cervantes, A., Freyer, G., Papamichael, D., Le Bail, N., Louvet, C., Hendler, D., de Braud, F., Wilson, C., Morvan, F., & Bonetti, A. (2000). Leucovorin and fluorouracil with or without oxaliplatin as first-line treatment in advanced colorectal cancer. *Journal of Clinical Oncology*, *18*(16), 2938–2947.

de Soyres, F., Herrero, J. G.-C., Goernemann, N., Jeon, S., Lofstrom, G., & Moore, D. (2024). Why is the us gdp recovering faster than other advanced economies?

Dechezleprêtre, A., Märtin, R.-P., & Mohnen, M. (2014). *Knowledge spillovers from clean and dirty technologies*. Centre for Economic Performance.

Deng, B. (2015). Approval may embolden industry to combine cancer therapies. *Nature Medicine*, *21*(2), 105.

Dhanyamraju, P. K. (2024). Drug resistance mechanisms in cancers: Execution of pro-survival strategies. *Journal of Biomedical Research*, *38*(2), 95–121.

Dixit, A. K. (1993). *The art of smooth pasting* (Vol. 55). Taylor & Francis.

Dolgin, E. (2021). The tangled history of mrna vaccines. *Nature*, *597*(7876), 318–324.

Doraszelski, U., & Satterthwaite, M. (2010). Computable Markov-perfect industry dynamics. *The RAND Journal of Economics*.

Dosi, G. (1982). Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research policy*, *11*(3), 147–162.

Dubois, P., & Lasio, L. (2018). Identifying industry margins with price constraints: Structural estimation on pharmaceuticals. *The American Economic Review*.

Dubois, P., de Mouzon, O., Scott-Morton, F., & Seabright, P. (2015). Market size and pharmaceutical innovation. *The RAND Journal of Economics*, *46*(4), 844–871.

Dubois, P., Gandhi, A., & Vasserman, S. (2022, May). *Bargaining and international reference pricing in the pharmaceutical industry* (tech. rep. No. 30053). Cambridge, MA, National Bureau of Economic Research.

Dugoua, E., & Gerarden, T. (2023). *Induced innovation, inventors, and the energy transition* (tech. rep.). National Bureau of Economic Research.

Durvasula, M., Hemphill, C. S., Ouellette, L. L., Sampat, B., & Williams, H. L. (2023). The NBER Orange Book Dataset: A user's guide. *Research Policy*, *52*(7), 104791.

Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies*.

European Commission. (2021). Proposal for a regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52021PC0206

Farrell, J., & Saloner, G. (1986). Installed base and compatibility: Innovation, product preannouncements, and predation. *The American economic review*, 940–955.

Feldman, R. (2022). Understanding 'evergreening': Making minor modifications of existing medications to extend protections. *Health Affairs (Project Hope)*, *41*(6), 801–804.

Foster, A. D., & Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of political Economy*, *103*(6), 1176–1209.

Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annu. Rev. Econ.*, *2*(1), 395–424.

Frank, R., Shahzad, M., Kesselheim, A., & Feldman, W. (2022). Biosimilar competition: Early learning. *Health Economics*, *31*(4), 647–663.

Frankel, A. P., Krieger, J. L., Li, D., & Papanikolaou, D. (2023). *Evaluation and learning in r&d investment* (tech. rep.). National Bureau of Economic Research.

Frei, E., Karon, M., Levin, R. H., Freireich, E. J., Taylor, R. J., Hananian, J., Selawry, O., Holland, J. F., Hoogstraten, B., Wolman, I. J., Abir, E., Sawitsky, A., Lee, S., Mills, S. D., Burgert, E. O., Jr, Spurr, C. L., Patterson, R. B., Ebaugh, F. G., James, G. W., 3rd, & Moon, J. H. (1965). The effectiveness of combinations of antileukemic agents in inducing and maintaining remission in children with acute leukemia. *Blood*, *26*(5), 642–656.

Galasso, A., & Luo, H. (2018). Punishing robots: Issues in the economics of tort liability and innovation in artificial intelligence. In *The economics of artificial intelligence: An agenda* (pp. 493–504). University of Chicago Press.

Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *The American Economic Review*, *97*(3), 713–744.

Gilbert, R. J., & Newbery, D. M. (1982). Preemptive patenting and the persistence of monopoly. *The American Economic Review*, 514–526.

Gilchrist, D. S. (2016). Patents as a Spur to Subsequent Innovation? Evidence from Pharmaceuticals. *American Economic Journal: Applied Economics*, *8*(4), 189–221.

Goede, V., Fischer, K., Busch, R., Engelke, A., Eichhorst, B., Wendtner, C. M., Chagorova, T., de la Serna, J., Dilhuydy, M.-S., Illmer, T., Opat, S., Owen, C. J., Samoylova, O., Kreuzer, K.-A., Stilgenbauer, S., Döhner, H., Langerak, A. W., Ritgen, M., Kneba, M., … Hallek, M. (2014). Obinutuzumab plus chlorambucil in patients with CLL and coexisting conditions. *The New England Journal of Medicine*, *370*(12), 1101–1110.

Goettler, R. L., & Gordon, B. R. (2011). Does AMD Spur Intel to Innovate More? *The Journal of Political Economy*, *119*(6), 1141–1200.

Goldman, B. (2003). For investigational targeted drugs, combination trials pose challenges. *Journal of the National Cancer Institute*, *95*(23), 1744–1746.

Goodenough, Whittingham, & Yoshino. (2019). The Nobel Prize in Chemistry 2019 [Accessed: 2024-9-15].

Gormsen, N. J., & Huber, K. (2023). *Corporate discount rates* (tech. rep.). National Bureau of Economic Research.

Gote, V., Bolla, P. K., Kommineni, N., Butreddy, A., Nukala, P. K., Palakurthi, S. S., & Khan, W. (2023). A comprehensive review of mrna vaccines. *International journal of molecular sciences*, *24*(3), 2700.

Gowrisankaran, G., Nevo, A., & Town, R. (2015). Mergers When Prices Are Negotiated: Evidence from the Hospital Industry. *The American Economic Review*, *105*(1), 172–203.

Green, J. R., & Scotchmer, S. (1995). On the division of profit in sequential innovation. *The RAND Journal of Economics*, 20–33.

Griliches, Z. (1979). Issues in assessing the contribution of research and development to productivity growth. *The Bell Journal of Economics*, 92–116.

Griliches, Z. (1991). The search for r&d spillovers. *National Bureau of Economic Research Working Paper Series*, (w3768).

Griliches, Z. (1998). *R&d and productivity: The econometric evidence*. University of Chicago Press.

Grossman, G. M., & Helpman, E. (1991). Quality ladders in the theory of growth. *The review of economic studies*, *58*(1), 43–61.

Hall, B. H., Jaffe, A., & Trajtenberg, M. (2005). Market value and patent citations. *RAND Journal of Economics*, 16–38.

Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). The nber patent citation data file: Lessons, insights and methodological tools.

Hall, B. H., Mairesse, J., & Mohnen, P. (2010). Measuring the returns to r&d. In *Handbook of the economics of innovation* (pp. 1033–1082, Vol. 2). Elsevier.

Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2013). Who creates jobs? small versus large versus young. *Review of Economics and Statistics*, *95*(2), 347–361.

Hedestam, G. K., & Sandberg, R. (2023). Discoveries concerning nucleoside base modifications that enabled the development of effective mrna vaccines against covid-19.

Heller, M. A., & Eisenberg, R. S. (1998). Can patents deter innovation? the anticommons in biomedical research. *Science*, *280*(5364), 698–701.

Henderson, R. (1993). Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry. *The RAND Journal of Economics*, 248–270.

Henderson, R., & Clark, K. B. (1990). Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative science quarterly*, 9–30.

Hill, A., Gotham, D., Fortunak, J., Meldrum, J., Erbacher, I., Martin, M., Shoman, H., Levi, J., Powderly, W. G., & Bower, M. (2016). Target prices for mass production of tyrosine kinase inhibitors for global cancer treatment. *BMJ Open*, *6*(1), e009586.

Hodgson, C. (2024). *Information externalities, free riding, and optimal exploration in the uk oil industry*. National Bureau of Economic Research.

Holbeck, S. L., Camalier, R., Crowell, J. A., Govindharajulu, J. P., Hollingshead, M., Anderson, L. W., Polley, E., Rubinstein, L., Srivastava, A., Wilsker, D., Collins, J. M., & Doroshow, J. H. (2017). The national cancer institute ALMANAC: A comprehensive screening resource for the detection of anticancer drug pairs with enhanced therapeutic activity. *Cancer Research*, *77*(13), 3564–3576.

Hopenhayn, H., Neira, J., & Singhania, R. (2022). From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share. *Econometrica*, *90*(4), 1879–1914.

Howitt, P. (1999). Steady endogenous growth with population and r. & d. inputs growing. *Journal of Political Economy*, *107*(4), 715–730.

Humphrey, R. W., Brockway-Lunardi, L. M., Bonk, D. T., Dohoney, K. M., Doroshow, J. H., Meech, S. J., Ratain, M. J., Topalian, S. L., & Pardoll, D. M. (2011). Opportunities and challenges in the development of experimental drug combinations for cancer. *Journal of the National Cancer Institute*, *103*(16), 1222–1226.

Ifrach, B., & Weintraub, G. Y. (2016). A Framework for Dynamic Oligopoly in Concentrated Industries. *The Review of economic studies*, *84*(3), 1106–1150.

Igami, M. (2017a). Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *Journal of Political Economy*, *125*(3), 798–847.

Igami, M. (2017b). Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *The Journal of Political Economy*, *125*(3), 798–847.

Igami, M., & Uetake, K. (2019). Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996–2016. *The Review of Economic Studies*.

Institute of Medicine, Board on Health Care Services, & National Cancer Policy Forum. (2012, February). *Facilitating Collaborations to Develop Combination Investigational Cancer Therapies: Workshop Summary*. National Academies Press.

Jaffe, A. B. (1986). Technological opportunity and spillovers of r&d: Evidence from firms' patents, profits and market value.

Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *the Quarterly journal of Economics*, *108*(3), 577–598.

Jones, B., & Summers, L. (2021). A calculation of the social returns to innovation. In *Innovation and public policy* (pp. 13–60). University of Chicago Press. https://doi.org/doi:10.7208/chicago/9780226805597-005

Jones, C. (2023). Recipes and economic growth: A combinatorial march down an exponential tail. *Journal of Political Economy*, *131*(8), 1994–2031.

Jones, C., & Williams, J. (1998). Measuring the social return to r&d. *The Quarterly Journal of Economics*, *113*(4), 1119–1135.

Jones, C., & Williams, J. (2000). Too much of a good thing? the economics of investment in r&d. *Journal of Economic Growth*, *5*, 65–85.

Jones, C. I. (1995). R & d-based models of economic growth. *Journal of political Economy*, *103*(4), 759–784.

Jones, C. I. (2016). Life and growth. *Journal of political Economy*, *124*(2), 539–578.

Jones, C. I. (2023). *The ai dilemma: Growth versus existential risk* (tech. rep.). Stanford GSB. Mimeo.

Jung, J. K., Feldman, R., & McBean, A. M. (2017). The Price Elasticity of Specialty Drug Use: Evidence from Cancer Patients in Medicare Part D. *Forum for Health Economics & Policy*, *20*(2).

Jurowetzki, R., Hain, D., Mateos-Garcia, J., & Stathoulopoulos, K. (2021). The privatization of ai research (-ers): Causes and potential consequences–from university-industry interaction to public research brain-drain? *arXiv preprint arXiv:2102.01648*.

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T., Lerner, J., & Tahoun, A. (2023). The diffusion of new technologies.

Kao, J., Ross, J., & Miller, J. (2023). Transparency of results reporting in cancer clinical trials. *JAMA Network Open*, *6*.

Karikó, K., Buckstein, M., Ni, H., & Weissman, D. (2005). Suppression of rna recognition by toll-like receptors: The impact of nucleoside modification and the evolutionary origin of rna. *Immunity*, *23*(2), 165–175.

Katz, M. L., & Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of political economy*, *94*(4), 822–841.

Kim, S. (2023). Shortcuts to Innovation: The Use of Analogies in Knowledge Production.

Kinch, M. (2018). *Between hope and fear*. Simon; Schuster.

Kitch, E. W. (1977). The nature and function of the patent system. *The Journal of Law and Economics*, *20*(2), 265–290.

Klette, T. J. (1996). R&d, scope economies, and plant performance. *The RAND Journal of Economics*, 502–522.

Klette, T. J., & Kortum, S. (2004). Innovating firms and aggregate innovation. *Journal of political economy*, *112*(5), 986–1018.

Klinger, J., Mateos-Garcia, J., & Stathoulopoulos, K. (2020). A narrowing of ai research? *arXiv preprint arXiv:2009.10385*.

König, M., Storesletten, K., Song, Z., & Zilibotti, F. (2022). From imitation to innovation: Where is all that chinese r&d going? *Econometrica*, *90*(4), 1615–1654.

Krieger, J., Li, D., & Papanikolaou, D. (2022). Missing novelty in drug development. *The Review of Financial Studies*, *35*(2), 636–679.

Kuhn, T. (1970). The nature of scientific revolutions. *Chicago: University of Chicago*, *197*(0).

Lee, R., Whinston, M., & Yurukoglu, A. (2021). Chapter 9 - Structural empirical analysis of contracting in vertical markets. *Handbook of Industrial Organization*, *4*(1), 673–742.

Lerner, J., & Seru, A. (2022). The use and misuse of patent data: Issues for finance and beyond. *The Review of Financial Studies*, *35*(6), 2667–2704.

Lerner, J., & Tirole, J. (2004). Efficient patent pools. *American Economic Review*, *94*(3), 691–711.

Levengood, T. W., Conti, R. M., Cahill, S., & Cole, M. B. (2024). Assessing the Impact of the 340B Drug Pricing Program: A Scoping Review of the Empirical, Peer-Reviewed Literature. *The Milbank Quarterly*, *102*(2), 429–462.

Lichtenberg, F. (2010). Has medical innovation reduced cancer mortality?, 135–177.

Lucas, R. E., Jr, & Moll, B. (2014). Knowledge growth and the allocation of time. *Journal of Political Economy*, *122*(1), 1–51.

Lucking, B., Bloom, N., & Van Reenen, J. (2019). Have R&D Spillovers Declined in the 21st Century? *Fiscal Studies*, *40*(4), 561–590.

Luttmer, E. G. (2007). Selection, growth, and the size distribution of firms. *The Quarterly Journal of Economics*, *122*(3), 1103–1144.

Maini, L., & Pammolli, F. (2020, April). *Reference Pricing as a Deterrent to Entry: Evidence from the European Pharmaceutical Market*.

Marconi & Braun. (1909). The Nobel Prize in Physics 1909 [Accessed: 2024-9-15].

Marshall, A. (1890). *Principles of economics*. Macmillan.

Martin, I. W., & Pindyck, R. S. (2015). Averting catastrophes: The strange economics of scylla and charybdis. *American Economic Review*, *105*(10), 2947–2985.

Martin, I. W., & Pindyck, R. S. (2021). Welfare costs of catastrophes: Lost consumption and lost lives. *The Economic Journal*, *131*(634), 946–969.

Maskin, E., & Tirole, J. (1988). A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica*, *56*(3), 549.

Meric-Bernstam, F., Ford, J. M., O'Dwyer, P. J., Shapiro, G. I., McShane, L. M., Freidlin, B., O'Cearbhaill, R. E., George, S., Glade-Bender, J., Lyman, G. H., Tricoli, J. V., Patton, D., Hamilton, S. R., Gray, R. J., Hawkins, D. S., Ramineni, B., Flaherty, K. T., Grivas, P., Yap, T. A., … Moscow, J. A. (2023). National Cancer Institute combination therapy platform trial with Molecular Analysis for Therapy Choice (ComboMATCH). *Clinical Cancer Research*, *29*(8), 1412–1422.

Milgrom, P., & Shannon, C. (1994). Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, 157–180.

Mokhtari, R., Homayouni, T. S., Baluch, N., Morgatskaya, E., Kumar, S., Das, B., & Yeger, H. (2017). Combination therapy in combating cancer. *Oncotarget*, *8*(23), 38022–38043.

Moore, T. J., Heyward, J., Anderson, G., & Alexander, G. C. (2020). Variation in the estimated costs of pivotal clinical benefit trials supporting the US approval of new therapeutic agents, 2015-2017: a cross-sectional study. *BMJ Open*, *10*(6), e038863.

Murphy, S. L., Kochanek, K. D., Xu, J., & Arias, E. (2021). Mortality in the united states, 2021.

Narang, A. K., & Nicholas, L. H. (2017). Out-of-pocket spending and financial burden among Medicare beneficiaries with cancer. *JAMA Oncology*, *3*(6), 757–765.

National Comprehensive Cancer Network. (2024). National Comprehensive Cancer Network Guidelines.

NCI. (2014, January). Off-Label Drug Use in Cancer Treatment [Accessed: 2024-6-29].

NCI. (2022a, May). 2023 NCI Budget Fact Book - Funding Trends [Accessed: 2024-7-16].

NCI. (2022b, May). 2023 NCI Budget Fact Book - Research Funding [Accessed: 2024-7-16].

Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. harvard university press.

Ochieng, N., Biniek, J. F., Freed, M., Damico, A., & Neuman, T. (2023, August). Medicare Advantage in 2023: Enrollment Update and Key Trends [Accessed: 2024-1-25].

Ochieng, N., Cubanski, J., & Neuman, T. (2024, September). A Snapshot of Sources of Coverage Among Medicare Beneficiaries [Accessed: 2024-10-9].

O'Shaughnessy, J. A., Vukelja, S., Marsland, T., Kimmel, G., Ratnam, S., & Pippen, J. E. (2004). Phase II study of trastuzumab plus gemcitabine in chemotherapy-pretreated patients with metastatic breast cancer. *Clinical Breast Cancer*, *5*(2), 142–147.

Pakes, A. (1986). Patents as Options: Some Estimates of the Value of Holding European Patent Stocks. *Econometrica*, *54*(4), 755–784.

Pakes, A., Ostrovsky, M., & Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics*, *38*(2), 373–399.

Parente, S. L., & Prescott, E. C. (1994). Barriers to technology adoption and development. *Journal of political Economy*, *102*(2), 298–321.

Podkonjak, T., Taylor, H., Taylor, A., Morgan, B., Mann, O., Roffe, E., Kaiser, M., Rabin, N., Hatswell, A., & Low, E. (2021). *Voluntary Arbitration Framework for Combination Therapies* (tech. rep.). Voluntary Arbitration Working Group.

Price, W., Nicholson, I., & Rai, A. (2015). Manufacturing barriers to biologics competition and innovation. *Iowa Law Review*, *101*, 1023.

Rao, A. (2020). Strategic Research and Development Investment Decisions in the Pharmaceutical Industry. *Marketing Science*, *39*(3), 564–586.

Reinganum, J. F. (1983). Uncertain innovation and the persistence of monopoly. *The American Economic Review*, *73*(4), 741–748.

Rietz, T. A. (1988). The equity risk premium a solution. *Journal of monetary Economics*, *22*(1), 117–131.

Robinson, J. C., Whaley, C., & Dhruva, S. S. (2024). Hospital Prices for Physician-Administered Drugs for Patients with Private Insurance. *New England Journal of Medicine*, *390*(4), 338–345.

Robinson, J. C., Whaley, C. M., & Brown, T. T. (2021). Price differences to insurers for infused cancer drugs in hospital outpatient departments and physician offices: Study examines differences in prices insurers pay for cancer drugs in hospital outpatient clinics compare with physician offices. *Health Affairs (Project Hope)*, *40*(9), 1395–1401.

Roin, B. N. (2013, October). *Solving the Problem of New Uses*.

Romer, P. (1992). Two strategies for economic development: Using ideas and producing ideas. *The World Bank Economic Review*, *6*(suppl 1), 63–91.

Romer, P. M. (1986). Increasing returns and long-run growth. *Journal of political economy*, *94*(5), 1002–1037.

Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy*, *98*(5, Part 2), S71–S102.

Sampat, B., & Williams, H. L. (2019). How do patents affect follow-on innovation? evidence from the human genome. *American Economic Review*, *109*(1), 203–236.

Sandler, A., Gray, R., Perry, M. C., Brahmer, J., Schiller, J. H., Dowlati, A., Lilenbaum, R., & Johnson, D. H. (2006). Paclitaxel-carboplatin alone or with bevacizumab for non-small-cell lung cancer. *The New England Journal of Medicine*, *355*(24), 2542–2550.

Sanofi. (2022). Combination treatments in oncology [Accessed: 2024-1-21].

Schilsky, R. (2013). Publicly funded clinical trials and the future of cancer care. *The Oncologist*, *18*(2), 232–238.

Schumpeter, J. (1934, January). *The Theory of Economic Development*. Harvard University Press.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*. harper.

Scotchmer, S. (2004). *Innovation and incentives*. MIT press.

Scott, P. (2014, May). Dynamic Discrete Choice Estimation of Agricultural Land Use. In *TSE Working Paper* (p. 56, Vol. 14-526). TSE Working Paper.

Scott Morton, F., & Kyle, M. (2011). Markets for pharmaceutical products. In *Handbook of health economics* (pp. 763–823, Vol. 2). Elsevier.

Sertkaya, A., Birkenbach, A., Berlind, A., & Eyraud, J. (2014). *Examination of Clinical Trial Costs and Barriers for Drug Development* (tech. rep.). Eastern Research Group.

Sertkaya, A., Wong, H.-H., Jessup, A., & Beleche, T. (2016). Key cost drivers of pharmaceutical clinical trials in the United States. *Clinical Trials (London, England)*, *13*(2), 117–126.

Shevlane, T., Farquhar, S., Garfinkel, B., Phuong, M., Whittlestone, J., Leung, J., Kokotajlo, D., Marchal, N., Anderljung, M., Kolt, N., Ho, L., Siddarth, D., Avin, S., Hawkins, W., Kim, B.,

Gabriel, I., Bolina, V., Clark, J., Bengio, Y., ... Dafoe, A. (2023). Model evaluation for extreme risks.

Siegel, R. L., Miller, K. D., Wagle, N. S., & Jemal, A. (2023). Cancer statistics, 2023. *CA: A Cancer Journal for Clinicians*, *73*(1), 17–48.

Smith, G. L., Banegas, M. P., Acquati, C., Chang, S., Chino, F., Conti, R. M., Greenup, R. A., Kroll, J. L., Liang, M. I., Pisu, M., Primm, K. M., Roth, M. E., Shankaran, V., & Yabroff, K. R. (2022). Navigating financial toxicity in patients with cancer: A multidisciplinary management approach. *CA: A Cancer Journal for Clinicians*, *72*(5), 437–453.

Solow, R. M. (1957). Technical change and the aggregate production function. *The review of Economics and Statistics*, *39*(3), 312–320.

Song, M., Nicholson, S., & Lucarelli, C. (2017). Mergers with interfirm bundling: a case of pharmaceutical cocktails. *The RAND journal of Economics*, *48*(3), 810–834.

Stern, A. D., Chen, J., Ouellet, M., Trusheim, M., El-Kilani, Z., Jessup, A., & Berndt, E. (2021). Biosimilars and follow-on products in the United States: Adoption, prices, and users: Study examines the early years of entry of biosimilar and related follow-on products including sales data, market share, and price trajectory. *Health Affairs (Project Hope)*, *40*(6), 989–999.

Strohbehn, G. W., Kacew, A. J., Goldstein, D. A., Feldman, R. C., & Ratain, M. J. (2021). Combination therapy patents: A new front in evergreening. *Nature Biotechnology*, *39*(12), 1504–1510.

Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica*.

Tu, S. S., & Sarpatwari, A. (2023). A" method of use" to prevent generic and biosimilar market entry. *The New England Journal of Medicine*, *388*(6), 483–485.

Unger, J. M., Nghiem, V. T., Hershman, D. L., Vaidya, R., LeBlanc, M., & Blanke, C. D. (2019). Association of national cancer institute-sponsored clinical trial network group studies with guideline care and new drug indications. *JAMA Network Open*, *2*(9), e1910593.

Wang, L. X. (2022). Global drug diffusion and innovation with the medicines patent pool. *Journal of Health Economics*, *85*(102671), 102671.

Wang, L. X. (2023). A cost-benefit analysis of the medicines patent pool". *Economics Bulletin*, *43*(3), 1298–1319.

Weintraub, G., Benkard, L., & Van Roy, B. (2008). Markov Perfect Industry Dynamics With Many Firms. *Econometrica*, *76*(6), 1375–1411.

Weitzman, M. L. (1998). Recombinant growth. *The Quarterly Journal of Economics*, *113*(2), 331–360.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The review of economics and statistics*, *91*(1), 1–19.

Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*.

Williams, H. L. (2013). Intellectual property rights and innovation: Evidence from the human genome. *The Journal of Political Economy*, *121*(1), 1–27.

Wong, C., Siah, K., & Lo, A. (2019). Estimation of clinical trial success rates and related parameters. *Biostatistics (Oxford, England)*, *20*(2), 273–286.

Yu, T. (2021). How scientists drew weissman (med'87, grs'87) and katalin karikó developed the revolutionary mrna technology inside covid vaccines. *Bostonia*. https://www.bu.edu/articles/2021/how-drew-weissman-and-katalin-kariko-developed-mrna-technology-inside-covid-vaccines/

Yurukoglu, A., Liebman, E., & Ridley, D. (2017). The role of government reimbursement in drug shortages. *American Economic Journal: Economic Policy*, *9*(2), 348–382.

Zacchia, P. (2020). Knowledge spillovers through networks of scientists. *The Review of Economic Studies*, *87*(4), 1989–2018.