

Essays on the Economics of Innovation

by

David Colino Llamas

Master in Economics and Finance, CEMFI (2011)

M.Sc. Mechanical Engineering, Ecole Polytechnique (2009)

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

June 2017

© David Colino Llamas, MMXVII. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature redacted

Author

Department of Economics

May 15, 2017

Signature redacted

Certified by...

Glenn Ellison

Gregory K. Palm Professor of Economics

Thesis Supervisor

Signature redacted

Certified by..

Heidi L. Williams

Class of 1957 Career Development Associate Professor

Thesis Supervisor


Signature redacted

Accepted by

Ricardo J. Caballero

Ford International Professor of Economics

Chairman, Departmental Committee on Graduate Studies



Essays on the Economics of Innovation

by

David Colino Llamas

Submitted to the Department of Economics
on May 15, 2017, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis consists of three chapters.

In the first chapter, I estimate the dynamic or inter-temporal knowledge spillovers resulting from corporate R&D in a setting with cumulative innovation, using a panel of US firms and a network of corporate patent citations. I show that the positive effect of dynamic spillovers on other firms' productivity is economically important, and at least as large as that of own R&D investments. Accounting for both static and dynamic spillovers, my estimates suggest that the social returns to corporate R&D are about three times as large as the private returns.

The second chapter, joint with Jean-Noel Barrot, studies the effect of patent term duration on the rate and direction of follow-on innovation, using a quasi-natural experiment that lengthened the term of existing patents in the US. Leveraging a kink in the patent term extension formula, we find no significant impact of extensions on subsequent innovation, neither locally around the kink using a sharp "Regression Kink Design" nor on average on the population of treated patents.

The third chapter, joint with Nicolas Caramp and Pascual Restrepo, studies how consumer durables amplify business cycle fluctuations on aggregate employment. We show that employment in durable manufacturing industries is more cyclical than in other industries, and that this cyclicality is amplified in general equilibrium. Our estimates suggest that consumer durables are responsible for up to 40% of aggregate employment volatility.

Thesis Supervisor: Glenn Ellison

Title: Gregory K. Palm Professor of Economics

Thesis Supervisor: Heidi L. Williams

Title: Class of 1957 Career Development Associate Professor

Acknowledgments

I am extremely grateful to my advisors, Glenn Ellison and Heidi Williams. Glenn has been a constant source of thoughtful advice throughout my peripatetic research journey. It took me too long to fully grasp the strength of his main mantra: "Just write the papers." Heidi was instrumental in orienting my research within the economics of innovation. I genuinely wish we had crossed paths earlier.

I have greatly learned from my interactions with many faculty at the Economics department and at Sloan. Likewise, and as a testimony of the wide range of interests spanned in these projects, my thesis was enriched by comments from participants at the IO, Finance, Applied Micro, Macroeconomics, and Labor lunches at MIT, as well as the NBER Productivity lunch. I would also like to thank my classmates and peers, who shaped my daily doctoral experience, as well as the many departmental staff who made my life easier. Finally, much of this thesis would not have been possible without my co-authors Jean-Noel, Nico, and Pascual.

Over almost six years at MIT, I have made many friends: at the department, through Spain@MIT, on the soccer pitch. Although listing them all would fill out pages, I would like to specifically thank Nico and Pascual, who were the first friends I made in Cambridge; Diego, who brought increased goofiness to our group; Enrico, with whom I navigated the most stressful times in grad school; DJ and Leo, for all those relaxing evenings at home; Elena, who was a strong source of support during my job market paper's infancy; Sara and Patty, for bringing joy to the office; all the committee members of Spain@MIT, for their time and dedication; Juan, for showing me the value of public service; Leo, Ray, Tite, Horacio, Iñigo, Quique, Enrique, and Roberto, for organizing all those soccer games. I am deeply honored and grateful to have shared your path.

To Caitlin and to my family: You are my rock. Thank you.

Contents

1	Cumulative Innovation and Dynamic R&D Spillovers	15
1.1	Introduction	15
1.2	Data and construction of variables	22
1.2.1	Measuring R&D spillovers	22
1.2.2	Data	25
1.2.3	Dynamic network proximity measures	26
1.2.4	Static proximity measures	28
1.3	Identification	30
1.3.1	Endogeneity of R&D	31
1.3.2	Endogenous network formation	32
1.4	Empirics	34
1.4.1	Productivity and market value equations	34
1.4.2	Baseline results	36
1.4.3	Robustness of dynamic spillovers	39
1.4.4	Effect on innovation outcomes	40
1.5	Heterogeneity across industries and technology types	42
1.6	Extensions	46
1.6.1	Second-degree spillovers	47
1.6.2	Geographic spillovers and customer-supplier relations	47
1.6.3	Depreciation of spillovers	49
1.6.4	Size of originating firm	50
1.6.5	Venture Capital spillovers	50
1.7	Estimates for the private and social returns to R&D	51
1.7.1	Own vs. others' spillovers	52
1.7.2	Calculating social and private returns	53
1.8	Conclusion	57

2	Patent duration and cumulative innovation: Evidence from a quasi-natural experiment (joint with Jean-Noel Barrot)	75
2.1	Introduction	75
2.2	Research setting	79
2.2.1	Uruguay Round	79
2.2.2	TRIPS reform	80
2.2.3	Data and variable construction	82
2.2.4	Empirical strategy	84
2.3	Impact of patent term extension	87
2.3.1	Panel evidence	87
2.3.2	Cross-sectional evidence	90
2.3.3	Regression Kink design	94
2.4	Heterogeneity	96
2.4.1	By technology class	96
2.4.2	By grant and filing year	97
2.4.3	By assignee size	98
2.4.4	By patent quality	98
2.5	Discussion	99
2.6	Conclusion	100
3	Durable Crises (joint with Nicolas Caramp and Pascual Restrepo)	125
3.1	Introduction	125
3.2	Data sources	130
3.3	Evidence from U.S. industries	132
3.4	Evidence from U.S. local labor markets	135
3.5	Mechanisms that amplify the shock to durables	140
3.5.1	Input-output linkages	141
3.5.2	Demand spillovers affecting non-tradables	144
3.5.3	Lack of reallocation	146
3.6	Quantitative implications and remarks	147
A	Appendices for Chapter 1	157
A.1	Variable construction	157
A.1.1	Static spillovers	157
A.1.2	Dynamic spillovers	160

A.2	Instruments	161
A.2.1	Tax instrument	161
A.2.2	Network instrument	164
A.3	Analytical framework	165
A.3.1	Cumulative innovation production model	166
A.3.2	Computing private and social returns without amplification	169
A.3.3	Optimal provision of R&D	171
A.4	Using alternative proximity measures	173
A.5	Additional results	174
A.5.1	Dynamic spillovers within and between industries	174
A.5.2	Dynamic spillovers within and between technologies	175

List of Figures

1-1	Patent counts across time in the NBER data	60
1-2	Correlation between static and dynamic knowledge spillovers, DYN- SPILL and CITSPILL	60
1-3	Heterogeneity across industrial sectors	61
1-4	Heterogeneity across product and technology types	62
1-5	Heterogeneity across technologies	63
1-6	Direct and indirect dynamic spillovers	63
1-7	Depreciation of spillovers	64
1-8	Heterogeneity across firm sizes	65
1-9	Spillovers from Venture Capital	65
2-1	Histogram of patent term extensions due to TRIPS	102
2-2	Patent applications and grants around TRIPS implementation	103
2-3	TRIPS treatment function	103
2-4	Histogram of processing time around 3-year TRIPS cutoff.	104
2-5	Difference in yearly citation counts per extra term day	105
2-6	Difference in yearly citation counts per extra term day, controlling for processing time	106
2-7	Difference in yearly citation counts per extra term year, RKD without controls	107
2-8	Difference in yearly citation counts per extra term year, RKD with controls	108
2-9	Difference in yearly citation counts by technology type	109
2-10	Difference in yearly citation counts by grant year	110
2-11	Difference in yearly citation counts by application year	111
2-12	Difference in yearly citation counts by assignee size	112
2-13	Difference in yearly citation counts by pre-TRIPS patent citation count	113

3-1	Cyclical component of U.S. employment and its counterfactual behavior if no U.S. region produced durables, nor supplied durable industries.	148
3-2	Unemployment slack and output gap in the U.S over time.	149
3-3	Average durability of goods and upstream exposure to consumer durables by commuting zone in 1988.	149
3-4	Employment and its counterfactual behavior if no U.S. region produced durables, there were no upstream propagation to industries in affected regions, or there were no upstream propagation to industries in other regions. All series are expressed in percent deviations from their 2007 level.	150
A-1	Correlation between CIT and SIC proximity measures	176
A-2	Correlation between CIT and TEC proximity measures	176
A-3	Correlation between SIC and TEC proximity measures	177
A-4	Timing of analytical model	177
A-5	Dynamic spillovers between and within industries and technology categories	178

List of Tables

1.1	Descriptive statistics, variables	66
1.2	Predictions for market value, productivity and R&D	66
1.3	First stages, effect of tax instruments on R&D expenses	67
1.4	First stages, pre-period network on spillover network	68
1.5	Baseline estimates	69
1.6	Robustness checks	70
1.7	Dynamic spillovers on R&D and citation-weighted patenting	71
1.8	Supply chain and geographic spillovers	72
1.9	Own vs others' spillovers, output	73
2.1	Descriptive statistics, variables	114
2.2	t-tests of covariates for large processing time differences	115
2.3	t-tests of covariates around the kink	116
2.4	Panel specifications	117
2.5	Cross-sectional specifications	118
2.6	Alternative cross-sectional specifications	119
2.7	Alternative cross-sectional specifications, citations by years	120
2.8	Extensive margin, assignee counts	121
2.9	Regression kink discontinuity on covariates	122
2.10	Regression kink discontinuity on outcome variables	123
3.1	Descriptive statistics of covariates at the commuting zone level.	151
3.2	Estimates at the industry level of the different response of durable industries to economic fluctuations.	152
3.3	Estimates at the commuting zone level of the different response of regions that host durable industries to economic fluctuations.	153

3.4	Estimates at the commuting zone level of the different response of net migration in regions that host durable industries to economic fluctuations.	154
3.5	Estimates at the industry level of the different response of durable industries and their suppliers to economic fluctuations.	155
3.6	Estimates at the commuting zone level of the negative spillovers created by the decline in employment in the durable industry on other sectors.	156
A.1	Effect of R&D intensity on average patent citations	179
A.2	Alternative static spillover measures, sales	180
A.3	Alternative static spillover measures, MTB	181

Chapter 1

Cumulative Innovation and Dynamic R&D Spillovers

Abstract

While much theoretical attention has focused on the important role of dynamic knowledge spillovers for economic growth, such spillovers have been difficult to empirically measure. Using a panel of US firms and a network of corporate patent citations, this paper estimates the dynamic spillovers of corporate R&D on firm productivity, value, and innovation activity. Causal effects are estimated with an instrumental variables strategy that exploits the persistence of the network as well as variation in tax incentives. The positive effect of dynamic spillovers on firm productivity is economically important, and at least as large as that of own R&D investments. Dynamic spillovers accrue mainly for so-called "complex" technologies that build cumulatively on multiple components, they exhibit little depreciation over time, are larger for established firms than for VC-backed startups, and do not decrease in magnitude with geographic distance. Around 40% of dynamic spillovers are re-absorbed by the original innovator, and –accounting for other spillovers– these estimates suggest that the social returns to R&D are about three times as large as the private returns.

1.1 Introduction

Innovation and technological progress are at the core of long-term economic growth. The endogenous growth literature attributes steady-state growth to the knowledge spillovers resulting from innovative activity.¹ As a result of their central importance, a substantial literature has attempted to estimate knowledge spillovers accruing be-

¹See e.g. Romer (1990), Aghion and Howitt (1992), and Jones (1995).

tween firms.² These measurements have traditionally focused on contemporaneous spillovers resulting among others from complementarities in the research and development (R&D) efforts of different firms.³ However, innovation is often cumulative in nature and the endogenous growth literature in particular depicts knowledge spillovers as intertemporal or dynamic, accruing when past ideas become the new foundation on which to build further innovation. We might thus expect a large share of knowledge spillovers to occur dynamically through the impact of a given technology on subsequent future innovation. By not fully accounting for this dynamic dimension, existing measurements of R&D knowledge spillovers may underestimate the wedge between the social and private rates of return to R&D. The contribution of this paper is to provide more robust estimates of the gap between social and private returns to R&D by flexibly estimating *dynamic knowledge spillovers* of US corporate R&D.

Since Jaffe (1986)'s seminal paper, the R&D spillover literature has assumed a rigid inter-temporal structure of knowledge spillovers in which both the private value of knowledge and its spillovers are assumed to depreciate at the same constant rate, which marks the progressive obsolescence of past knowledge as it is replaced by new innovations.⁴ In other words, R&D expenses build a knowledge stock that depreciates over time, in the same way that capital investments contribute to a capital stock, and this knowledge stock is responsible for both private value creation and contemporaneous knowledge spillovers on others. However, as pointed out by Griliches (1979) the social depreciation rate of innovation is likely to be lower than the private one, and depreciation rates may not be the same across innovative streams, as a result of cumulative processes in which specific past knowledge is used as a foundation upon which to build future innovation. Knowledge spillovers are thus likely to entail a complex dynamic or inter-temporal structure, and not accounting for dynamism is likely to lead to underestimating spillovers. Moreover, this bias can be especially acute for technologies that rely more on long and cumulative innovation processes. In what follows, and in order to make a sharp distinction between the traditional knowledge spillover measures and the more flexible spillovers I introduce, I will denote the former

²See Griliches (1979), Jaffe (1986), Bloom et al. (2013), and Manresa (2016).

³For example, Bloom et al. (2013) model knowledge spillovers as occurring through encounters between scientists or engineers of different firms. These encounters could be in person through meeting at conferences or local coffee shops, or they could be virtual through online exchanges of ideas.

⁴For more literature using this same obsolescence structure, see Bloom et al. (2013), Manresa (2016), and Schnitzer and Watzinger (2015).

as *static* knowledge spillovers and the latter as *dynamic* knowledge spillovers.⁵

In this paper, I estimate dynamic knowledge spillovers by constructing measures of cumulative knowledge proximity between US publicly-listed firms. My empirical estimates suggest that the R&D (carried out in the past and largely by other firms) upon which a given firm builds increases its current productivity by at least as much as current own R&D. I find that dynamic spillovers complement the traditional *static* knowledge and business stealing spillovers, as they were not picked up by these measures.⁶ However, the share of spillovers that are dynamic versus static is largely heterogeneous across technologies, with dynamic spillovers being more prevalent among industries and technologies characterized by their complexity, that is by commercializable products or processes being comprised of numerous separately patentable elements.⁷ Dynamic spillovers persistently accrue over a long period of time, are stronger when originating in R&D carried out by established firms rather than VC-backed startups, and do not decrease in magnitude with geographic distance, conditional on citation flows being observed. Including all the spillover measures results in estimates of the social returns to corporate R&D being about three times as large as the private ones.

The following example illustrates my empirical estimation of R&D spillovers between organizations. Motorola, IBM, and Apple are R&D-intensive companies that in the 1980s and 1990s are close in technological space (as revealed by both their patenting in similar technological areas and the large cross-citation patterns between them). However, only IBM and Apple compete in the PC market, with little product market competition between the other two pairs. Knowledge spillovers will therefore accrue between all pairwise relationships, but business stealing spillovers will only occur between IBM and Apple, so we can estimate both types of spillovers by comparing the different pairwise combinations. Moreover, traditional knowledge spillover measures assume that, for example, in 1996 Motorola and Apple are affected by IBM's 1996

⁵Although *static* spillovers of R&D expenditures still entail some dynamic component through building a longer-lived firm-specific knowledge capital, this slight abuse of denomination is appropriate if one considers that knowledge spillovers stem from the knowledge capital itself rather than the R&D spending flow.

⁶Static knowledge spillovers have been extensively studied by papers such as Griliches (1979), Griliches (1992), and Jaffe (1986). More recently, Bloom et al. (2013) separately measure both types of static spillovers and document that the knowledge spillovers dominate the business stealing effects, resulting in a social rate of return to R&D up to three times as large as the private returns.

⁷These technologies inherently involve cumulative innovation processes. See Levin et al. (1987), Cohen et al. (2000), Hall et al. (2005), and Galasso and Schankerman (2015).

stock of R&D, which includes past R&D spending depreciated at a constant rate.⁸ However, in 1996 Motorola was investing heavily in research on copper interconnect technology within the semiconductor chip industry and learning from IBM’s previously developed technology.⁹ In particular, Motorola developed and filed two patents that year building upon a 1985 IBM patent that proved to be a foundational innovation for copper interconnect.¹⁰ It is therefore hard to believe that knowledge IBM developed in 1985, upon which Motorola built in 1996, should be discounted by 83% in compounded terms when looking for knowledge spillovers between the two firms. My dynamic spillover measure takes this into account, and using patent citation patterns, selects IBM’s 1985 knowledge stock as a candidate for knowledge spillovers affecting Motorola in 1996. In my empirical application, I incorporate the traditional business stealing and static knowledge spillover measures between the three firms in my example, but also a dynamic spillover measure accounting more flexibly for inter-temporal cumulative channels.

Measuring innovation spillovers is challenging, because of the non-rivalry and non-excludability of ideas (Nelson, 1959; Arrow, 1962) and the fact that these spillovers are not usually formally recorded. Knowledge flows are largely unobserved, and can in principle affect a wide range of firms. In section 1.2, I propose an approach to overcome this challenge by using the patent citation network to construct measures of knowledge proximity between firm-year observations. I primarily rely on two data sources, Compustat and the NBER Patent Database, which allow me to observe accounting and financial data for all US publicly listed firms and their patents granted by the US Patent and Trademark Office (USPTO) between 1976 and 2006.¹¹ I use the patenting activity of firms and their citation patterns to construct a weighted network linking firm-year nodes. The network is then used as a proxy for the cumulative knowledge proximity between nodes, with network edges used to define proximity weights of the R&D pool upon which a firm is likely to draw in order to build subsequent innovations. These patent citation network measures are significantly distinct from the static technological proximity measures used previously in the literature,

⁸The most commonly used rate of depreciation in the literature is 15% annually. See Griliches (1979), Jaffe (1986), Hall et al. (2005), and Bloom et al. (2013).

⁹See Lim (2009) for more information about the copper interconnect technology, IBM’s technological advantage, and its competitors’ strategies on absorptive capacity.

¹⁰Patent 4789648: *Method for producing coplanar multilevel metalinsulator films on a substrate and for forming patterned conductive lines simultaneously with stud vias.*

¹¹See Hall et al. (2001) for more information about the patent data.

and I am thus able to separately estimate static and dynamic knowledge spillovers.

In section 1.3, I propose an identification strategy to estimate causal spillover effects on firm-level real outcomes such as productivity and market value. I tackle the endogenous network formation that underpins my knowledge spillover measures by exploiting the observed persistence in the network structure. In particular, I use the 1976-1984 network connecting firm-year observations to predict the 1987-2001 network connections between 340 R&D-intensive firms. I also tackle the endogeneity of R&D decisions using variation in the tax treatment of R&D expenditures at the state and federal level. I exploit state and federal changes in R&D tax credit rules and corporate taxes to construct R&D tax-price shifters, which are used to predict corporate R&D investment. I combine the predicted network and the tax-predicted R&D expenditures to construct instruments for the static and dynamic knowledge spillovers.¹² Finally, I use a set of firm and year fixed effects to take care of unobserved firm heterogeneity as well as any economy-wide shocks. I also include a full set of industry-times-year fixed effects in some specifications, to tackle industry-level shocks more flexibly and ensure that I compare similar firms within narrowly-defined industries and years.

I estimate R&D spillovers on firm output and market value in section 1.4. I document that the past R&D pool upon which any given firm builds increases its current output by as much as current own R&D, in both the OLS and the 2SLS specifications. The comparative effects of dynamic spillovers on market value are even larger. These spillovers are not picked up by traditional static spillover measures, and the dynamic spillover measures therefore increase the magnitude of total R&D spillovers. Static knowledge spillovers are about twice as large as the dynamic ones, and although there is evidence of negative business stealing effects, these are dominated by the positive knowledge spillovers as in Bloom et al. (2013). These results hold robustly across a battery of specifications and robustness tests. Among others, they are robust to including a set of industry-times-year fixed effects that flexibly control for time-varying industry-wide common shocks to firms. They are also robust to alternative specifications of the dynamic spillover measure, to controlling flexibly for own R&D, to including flexible polynomials in patent counts, and to considering only manufacturing firms. Dynamic spillovers are also associated to increases in innovative activity

¹²Instruments for business stealing spillovers are constructed following Bloom et al. (2013), using the tax-predicted R&D together with the actual proximity weights based on the firm's position in the product market space.

of spillover recipients, in terms of R&D expenditures and citation-weighted patent production. In particular, an increase in predicted dynamic R&D spillovers of 10% leads to an increase of 4.2% in citation-weighted patents granted.

Static knowledge spillovers are generally larger than dynamic ones, but when I analyze the heterogeneity in their relative importance across industries and technology types in section 1.5, I find that it varies greatly. Dynamic spillovers are larger within industries involving complex product types, that is commercializable products or processes being comprised of numerous separately patentable elements, whereas static spillovers are prevalent within more discrete product industries. Likewise, dynamic spillovers are large for electrical and mechanical technologies, whereas static spillovers dominate in chemical industries. Dynamic and static knowledge spillovers therefore look substantially different, and it seems dynamic measures are more adept at estimating knowledge flows in complex, more cumulative technologies. In other words, traditional spillover measures are too rigid to account for the extent of inter-temporal knowledge flows, and this limitation is shown to be especially acute for complex innovation processes. As a result, considering only static spillover measures for analyses of knowledge flows across technologies or industries can lead to misleading results.

In section 1.6, I study the heterogeneity of dynamic spillovers across a number of dimensions in order to understand them better. First, I study whether dynamic spillovers can accrue through second-degree connections within the patent citation network, and document evidence for indirect dynamic spillovers, which are about half as large as the direct ones in terms of productivity gains. I then analyze how geographic and customer-supplier relationships between firms influence the magnitude of dynamic spillovers. I show that, conditional on knowledge flows being observed, geographic distance between originator and receiver does not affect the magnitude of spillovers. However, if the two firms share a customer-supplier relation within a production network, the dynamic spillovers become insignificant. This is consistent with cited upstream innovation having to be incorporated in downstream production processes, rather than being used to further innovative output. I next study the depreciation of dynamic spillovers, and document evidence that the spillovers continue accruing over a long time period. R&D of over nine years of age still shows significant spillover effects. I also examine how the magnitude of dynamic spillovers varies depending on the size of the originating firm. Although dynamic spillovers are present across all firm sizes, corporate R&D carried out by larger firms generate

larger dynamic spillovers. Likewise, dynamic spillovers of R&D carried out by Venture Capital-backed startups are estimated to be substantially lower than the spillovers from established firms' corporate R&D.

In section 1.7, I quantify social returns to R&D by including static technology and business stealing spillovers together with dynamic spillovers, and estimate them to be about three times as large as private ones. Around 37% of dynamic knowledge spillovers are re-absorbed by the original innovators, for example in the context of longer and more complex projects involving separately patentable steps, and therefore probably do not constitute externalities.¹³ As a result, considering dynamic spillover measures increases both the private and the social returns to corporate R&D, as well as the wedge between both returns, but does not increase the ratio between returns. The estimates suggest a sizable under-investment in R&D in the decentralized equilibrium relative to the social optimum, with the social optimal level of R&D being three times as large as the decentralized level for a unitary elasticity of R&D to its user cost of investment.

This paper contributes mainly to the literature measuring R&D spillovers, with seminal studies by Griliches (1979) and Jaffe (1986).¹⁴ More recently, Bloom et al. (2013) separately measure static knowledge and business stealing spillovers following the framework set forth by Griliches and Jaffe. Manresa (2016) goes further in estimating spillovers without imposing a structure of knowledge flow interactions first. The existing literature, however, is mainly focused on static R&D spillover measures, with rigid inter-temporal structures, and this paper is, to the best of my knowledge, the first evidence of dynamic inter-temporal spillovers through knowledge flows associated to the cumulateness of innovation.

A large literature studies knowledge flows using citation patterns: in particular geographic determinants of knowledge flows (Jaffe et al., 1993; Furman et al., 2006), their age profile (Mehta et al., 2010), how they differ depending on whether the originating innovator is a VC-backed start-up or an established firm (González-uribe, 2012), and the effect of the institutional background (Furman and Stern, 2011; Murray et al., 2016). I build on this literature by using patent citations as evidence for knowledge flows, and estimate the spillovers associated to these flows. Moreover, I show

¹³This is consistent with Belenzon (2012), who shows that patents that are subsequently self-cited by the original patent-holding firm are positively related to its market value, evidence that the future absorption of spillovers is internalized *ex-ante*.

¹⁴There is also a large literature on absorptive capacity spurred by Cohen and Levinthal (1990), which focuses on firms' absorption of external knowledge spillovers.

heterogeneity in the magnitude of dynamic spillovers across product and technology types, depending on how prevalent their cumulative process is, and across firm sizes carrying out the original R&D. I also find evidence of compounding spillovers through indirect citations, and analyze how the spillover flows are affected by the relationship between the originating and receiving firm, by the age profile of citations, and by the size and type of originating firm. Williams (2013), Sampat and Williams (2015), and Galasso and Schankerman (2015) study the effect of intellectual property (IP) protection on subsequent innovation, building on a large theoretical literature started by Scotchmer (1991). In this paper, I document evidence of significant cumulative spillovers in the presence of patent protection for the original innovations.

1.2 Data and construction of variables

In this section, I discuss the challenges of measuring the R&D spillover flows, and how I tackle these challenges in constructing the dynamic and static spillover measures. I also describe the data used for the empirical analysis,

1.2.1 Measuring R&D spillovers

The economic intuition behind knowledge spillovers of corporate R&D is that a firm might benefit from the knowledge created by another firm. This knowledge absorption can take many forms: scientists and engineers from different firms might meet and exchange ideas; a firm can hire scientists and engineers previously employed by other firms; researchers might read publications or patents written by other researchers; a firm might reverse engineer its competitors' novel products. These knowledge flows are typically not recorded, and thus unobservable by the econometrician. An underlying challenge to the spillover literature is therefore to determine a way to measure these flows.

In order to estimate and measure R&D spillovers between different firms, in particular knowledge spillovers, one would ideally want to proceed in two steps. First, estimate a knowledge production function within each firm that takes into account own R&D activity and knowledge spillovers. Then, estimate the effect of that knowledge on firm real outcomes such as its market value, or its productivity. In other words, we would like to estimate the following system of equations:

$$\begin{aligned}
k_{it} &= g(RD_{it}, A_t, k_{it-1}, \text{other inputs}_t), \\
Y_{it} &= h(k_{it}, k_{-it}, \text{other inputs}_t),
\end{aligned}
\tag{1.1}$$

where k_{it} stands for knowledge specific to firm i at period t , the subscript $-i$ relates to all other non- i firms, RD is the R&D effort exerted by firms, and A_t is the stock of economy-wide knowledge and technology that is used to produce further innovation. The knowledge variable k can be thought of as firm-specific know-how or ability. In settings with product differentiation, increases in knowledge could increase product quality. Likewise, in simpler homogeneous good settings, increases in knowledge can be thought of as decreasing marginal costs of production. As a result of competition between firms i and j , the equilibrium production level of firm i will also depend on firm j 's ability. For example, in a Cournot competition setting with two firms and heterogeneous marginal costs, the production level of each firm depends on both its own marginal cost and its competitor's marginal cost.

The knowledge production function $g(\cdot)$ includes own R&D effort and past firm knowledge k_{it-1} as variables, but also others' R&D through the term A_t . A_t differs from k in that it is not firm-specific, rather economy-wide. It represents how the aggregate level of knowledge due to past and current innovations affect the productivity of current R&D effort, both due to the cumulative nature of innovation and to complementarities in the R&D effort that arise if scientists and engineers working in similar areas discuss their research, share ideas or publications, and increase their productivity as a result. In terms of the different denominations of spillovers I have used in the introduction, the inclusion of k_{-it} in the production function $h(\cdot)$ corresponds to business stealing spillovers. Meanwhile, the spillovers included in the innovation production function $g(\cdot)$ correspond to knowledge spillovers: both static and dynamic. Cohen and Levinthal (1990), in their seminal paper on absorptive capacity, discuss a number of strategies through which firms can capitalize on these knowledge spillovers.

In practice, knowledge is a public and non-rival good which is non-observable by the econometrician. Knowledge flows between firms are equally unobserved. Therefore, we are reduced to estimate the reduced form function $f(\cdot)$

$$Y_{it} = f(RD_{it}, RD_{-it}, \text{past } RD_i, \text{past } RD_{-i}, \text{other inputs}), \tag{1.2}$$

where I have combined the previous $g(\cdot)$ and $h(\cdot)$ functions, as well as posited that the stock of knowledge A_t is available because it was created as innovations by R&D in the past and present, and is thus itself a function of past and present R&D $_i$ and R&D $_{-i}$. In order to make my estimates comparable with the literature, I use a log-linear reduced form function $f(\cdot)$. I consider the effect of R&D spillovers mainly on two different outcomes of interest, firm value and productivity, and analyze how knowledge spillovers act through the innovation production function by studying their effect on R&D decisions and innovative output in section 1.4.

As the spillover knowledge flows between firms I am interested in are unobserved, I need to make assumptions on which external R&D investments enter into the reduced-form function 1.2 and how they do so, because of the curse of dimensionality.¹⁵ The existing literature has dealt with this unobservability by considering pools of external R&D from all possible neighboring sources, weighted by a measure of the likelihood that each R&D actually generates spillovers. That is, the relevant pool of spillovers for firm i would be the sum of all other firms j 's R&D, each weighted by the likelihood that it actually results in spillovers. For a set of firms $\{1, 2, \dots, N\}$, we have

$$Spill_{it} = \sum_{j=1}^N \sum_{t' \leq t} \omega_{ijt'} RD_{jt'}, \quad (1.3)$$

where $Spill$ is the spillover pool of R&D, and ω is a weight associated to the proximity between firm-years. Each $\omega_{ijt'}$ can be thought of as the probability that R&D carried out by firm j at time $t' \leq t$ affects firm i at time t . Notice that by convention $\omega_{iit} = 0, \forall i, t$, since own R&D enters directly into the equation 1.2. Also, the matrix Ω of weights does not necessarily have to be symmetric, as we can have R&D $_{jt'}$ more likely to influence outcome Y_{it} than vice-versa. The definition and construction of the proximity weights ω will differ depending on the type of R&D spillover that is parameterized. I discuss the proximity matrices used to construct the different measures of R&D spillovers in the following subsections 1.2.3 and 1.2.4.

¹⁵Even assuming time-invariant relationships between a panel of N firms' R&D, the number of directed pair-wise coefficients to estimate is of the order of N^2 , with $N * T$ observations. If the dynamic structure of spillovers is flexible, the coefficients increase with $N^2 T^2$.

1.2.2 Data

I use yearly firm-level data from Compustat for financial and accounting data. According to the NSF Science and Engineering Indicators,¹⁶ in 2013 71% of all R&D conducted in the US was performed by businesses, of which 80% in companies with over 1,000 employees and 71% in multi-national companies. This suggests that R&D conducted by large, publicly-listed companies represents a major share of the total. Compustat firms are matched to their granted patents using USPTO data from the NBER Patent Project and the provided link.¹⁷ This data contains detailed information on over three million patents granted between 1976 and 2006, as well as their citation patterns. I consider patents filed between 1976 and 2001 to avoid attrition in the patent data due to delays between filing and granting of patents by the USPTO.¹⁸ The R&D-related data in Compustat is regarded as reliable starting in 1974.¹⁹

I use the patenting data in order to build the proximity measures for the static and dynamic knowledge spillovers, so I keep firms with active patenting and research activity. In particular, I restrict the analysis to US firms with at least two patents filed between 1976 and 2001 and at least three years of strictly positive R&D expenses. I exclude observations with negative or missing net sales and total book value of assets, and I drop regulated utilities and financial firms (SIC codes 6xxx and 49xx) since financial and accounting variables are not strictly comparable for them.²⁰ In order to minimize measurement error of R&D stocks, I require firms to have four or more Compustat observations, of which at least three years in a row.

I use Bloom et al. (2013)'s business stealing proximity measures. These are constructed by matching patenting firms to firms in Compustat Business Segments data between 1980 and 2001, which disaggregates firm sales by the industry in which they are conducted. See appendix A.1 for more details. This leaves a sample of 715 firms. I use the net stock of property, plant and equipment (Compustat variable PPENT) for the value of capital, and employee counts (EMP) for labor. I use R&D expenses (XRD) to calculate R&D stocks using a perpetual inventory method with a 15% de-

¹⁶See the report at <http://www.nsf.gov/statistics/2016/nsb20161/report/front-matter>.

¹⁷See Hall et al. (2001).

¹⁸The number of patents in the NBER Patent dataset starts to fall after 2001 due to attrition. See Figure 1-1. More conservative restrictions in the end year do not affect the results.

¹⁹The SEC and the Financial Accounting Standards Board have required since 1972 and 1974 respectively that publicly-listed firms report all material R&D expenditures, in the year in which they were incurred. See Bound et al. (1984).

²⁰This is common in the corporate finance literature, see e.g. Giroud and Mueller (2010).

preciation rate.²¹ I use net sales (SALE) as a measure of output, that I deflate using industry-specific deflators from the NBER-CES Manufacturing Industry Database²² and the Bureau of Economic Analysis. As a measure of Tobin’s Q , I use the market-to-book ratio calculated following Davis, Fama, and French (2000) as market equity over book equity, where market equity is price times outstanding shares and book equity is the book value of common equity, plus balance sheet deferred taxes and investment tax credit (if available).²³ I deflate all monetary values to 2000 USD using a GDP index from the BEA, and winsorize all variables at percentiles 1 and 99.

1.2.3 Dynamic network proximity measures

In order to measure the dynamic knowledge spillovers according to the framework in equation 1.3, I must define the relevant dynamic proximity matrix. The matrix construction must take into account the cumulative mechanism through which dynamic spillovers are expected to accrue, when building upon past innovation. The dynamic spillovers of R&D are thus to be generated *through* the ideas it originally creates, and their posterior use in the creation of further innovation. I use the firms’ patenting activity, and their citation network, in order to proxy for the dynamic accretion of knowledge. In doing so, I am conceptually using patents as a metric of firm innovative output and citations as recorded knowledge flows. This patent citation network is described in more detail in Colino (2016).

I define the dynamic spillover proximity matrix as follows. For firms i and j and years t and t' , with $i, j \in \{1, 2, \dots, N\}$ and $t, t' \in \{1, 2, \dots, T\}$, I define each element of the $NT * NT$ matrix $Dyn = (d_{ijt'})_{1 \leq i, j \leq N, 1 \leq t, t' \leq T}$ as

$$d_{ijt'} = \sum_{p \in \mathcal{P}_{it}} \sum_{q \in \mathcal{P}_{jt'}} \frac{\#Citations_{p \rightarrow q}}{Outcitations_p}, \quad (1.4)$$

where \mathcal{P}_{it} is the set of patents filed by node it , i.e. firm i in year t . That is, the proximity between nodes it and jt' will be constructed by counting all the citations

²¹Following Bloom et al. (2013), and Hall et al. (2005) among others. The R&D stock each year is $RDS_{i,t} = XRD_{i,t} + 0.85 * RDS_{i,t-1}$. Accounting for a 5% growth rate in R&D expenses, I set the first observation of the R&D stock to be $XRD/(0.2)$. Following the literature, I set missing R&D expenses to zero. Chauvin and Hirschey (1993) empirically test this assumption and confirm that it is generally appropriate

²²See Bartelsman and Gray (1996).

²³In Compustat mnemonics, $PRCC_F * CSHO / (CEQ + TXDB + ITCB)$. Results are robust to alternative specifications of market value over assets, see Appendix A.4 for details.

from all patents filed by it to all patents filed by jt' . I normalize by the total number of outcitations for each patent in it in order to keep the spirit of a constant returns to scale innovation production function.²⁴ That way, a patent does not mechanically receive more spillovers solely by increasing the size of its bibliography. Notice that this defines an asymmetric proximity matrix, with $d_{itjt'} \neq d_{jt'it}$. Likewise, for any it and jt' , if $t' > t$ then $d_{itjt'} = 0$. That is, a given firm-year cannot build upon innovation that has not yet been created.²⁵ Finally, we set $d_{itit} = 0$ for all it , since own R&D already enters the innovation production function directly.

The dynamic knowledge spillovers are then defined as

$$DynSpill_{it} = \sum_{j=1}^N \sum_{t' \leq t} d_{itjt'} * \tilde{G}_{jt'}, \quad (1.5)$$

where $\tilde{G} = \frac{RDS}{AT}$ represents the R&D intensity, or R&D stock over total assets, of each firm-year node. For more details on the relevance of using R&D intensities for the construction of dynamic spillover measures, see Appendix A.1.2.

Using patent data to study innovation production presents a set of well-documented challenges. First, the patenting decision is a strategic one. Patenting is costly, and there are alternative forms of intellectual property (IP) protection available to firms. However, a recent survey by Arora et al. (2016) reports 64% of large firms having patented their most significant recent innovation, testimony to the importance of patenting as a form of IP. Second, the patent quality distribution is highly skewed, and a large number of patents have been found to have little or no innovative value (Pakes and Schankerman, 1984; Bessen, 2009). I correct for this by considering only patents receiving a minimum amount of forward citations, above a threshold of three.²⁶ Third, patent citations themselves are noisy proxies for knowledge flows. In fact, citations themselves can also be strategic decisions by different actors (inventors, lawyers, patent examiners).²⁷ However, in order to gain patent protection of claims and assert the necessary novelty, as well as to define the boundaries of the protection, patents are required to cite prior relevant work. Notwithstanding their drawbacks, patents and citation data are commonly used in the innovation literature as proxies for innovative output and knowledge flows. Moreover, the drawbacks discussed above

²⁴I also estimate specifications with a non-normalized dynamic knowledge proximity for robustness in subsection 1.4.3. The results are qualitatively similar.

²⁵Results are robust to also setting $d_{itjt'} = 0$ if $t' = t$.

²⁶Results are robust to considering all patents instead, with no minimum threshold.

²⁷See for example Alcacer and Gittelman (2006) on citations added by patent examiners.

only indirectly apply to my analysis as I do not use patent or citation metrics directly as variables. Instead, I use them to proxy for cumulative technological proximity between different innovative actors, and measure how likely a given R&D effort is of generating spillovers. Therefore, using possibly noisy measures of proximity should result in downward bias against finding significant dynamic spillovers.

1.2.4 Static proximity measures

When considering spillovers, the literature has traditionally considered a *proximity* matrix that is static, and does not vary with time. I follow the literature here in assuming a rigid inter-temporal structure for static spillovers: R&D depreciates constantly across time, and R&D stocks only influence current outcomes through a static proximity matrix. In that way, for every time period t , we have

$$Spill_{it} = \sum_{j=1}^N \omega_{ij} RDS_{jt}, \quad (1.6)$$

where RDS_{jt} is R&D stock of firm j at time t , and ω_{ij} defines time-unvarying proximities between firms i and j . As a result, static spillover pools are composed of only concurrent R&D stocks by firm i 's neighbors. Although the spillovers considered still entail a dynamic aspect through the influence of past R&D expenses on the R&D stock, the inter-temporal structure is rigidly defined and the proximity measures are static. Ultimately, it is an empirical question whether these traditional measures provide good characterizations of the inter-temporal complexity of knowledge spillovers associated to R&D.

There are a number of possibilities to define ω_{ij} depending on the nature of spillovers we are interested in. In looking for knowledge flows, the most direct evidence of a flow between firms j and i can be established when a patent filed by i cites another patent filed by j as prior relevant art. The idea is that a patent citing another previous patent directly builds on knowledge incorporated into this prior art.²⁸ Therefore, to measure the proximity between firms i and j I look at all the citations from patents applied for by firm i to patents filed by j , aggregate them at the firm level and standardize by the total number of citations of the citing company, excluding self-citations.²⁹

²⁸See Schnitzer and Watzinger (2015), Azoulay et al. (2015), and Jaffe et al. (1993).

²⁹See Appendix A.1.1 for more details.

$$cit_{ij} = \frac{\#Citations_{i \rightarrow j}}{Outcitations_i}. \quad (1.7)$$

The resulting *proximity* matrix *cit* is asymmetric, and each line defines a set of weights of unitary sum. Citations are one of the most direct and intuitive way to construct knowledge spillover proximities. Nonetheless, there are many others. In robustness tests, I also use the Jaffe (1986) technological proximity measure, as well as more complex measures such as the Mahalanobis distance introduced by Bloom et al. (2013), or an expanded Jaffe (1986) technological proximity measure incorporating patent citation patterns between technology classes.³⁰ More details on the construction of these proximity measures can be found in Appendix A.1.1. I choose to use the citation proximity matrix in my main empirical analysis for two reasons. First, this matrix defines an asymmetric spillover measure, which can be an important characteristic to capture less efficient producers attempting to replicate industry leaders's best practices.³¹ Second, this static proximity closely replicates the dynamic proximity intuition of detecting knowledge flows through patent citations, while keeping the rigid inter-temporal structure of the other static spillover measures. As a result, it ensures that differences in the estimates of dynamic versus static spillovers will arise from the complex dynamic structure of knowledge flows, rather than from possible differences in detection of knowledge flows through technological proximity as in Jaffe (1986) versus through citation patterns as in Schnitzer and Watzinger (2015).

Using the proximity weights cit_{ij} , I define a measure of static knowledge spillovers

$$SpillCit_{it} = \sum_{j \neq i} cit_{ij} RDS_{jt}. \quad (1.8)$$

I show in Figure 1-2 how the resulting static knowledge spillover measure exhibits significant variation with respect to the dynamic spillover measure that I defined in the previous subsection.

With respect to business stealing spillovers, the distances between firms are defined using the product market activity of the firms. The product market vector S_i is defined using each firm's average sales broken down into a total of 597 4-digit SIC industries, where S_{ik} is the share of sales of firm i in industry k . The product market

³⁰See He (2015) and Schnitzer and Watzinger (2015).

³¹See Syverson (2011).

closeness between i and j is then defined as the correlation between S_i and S_j :

$$sic_{ij} = \frac{S_i S'_j}{(S_i S'_i)^{1/2} (S_j S'_j)^{1/2}}. \quad (1.9)$$

As before, this defines a business stealing proximity matrix sic , which is now symmetric. Spillovers $SpillSic$ are then calculated as for the $SpillCit$ case but using the relevant proximity matrix sic . More details about the construction of these spillover pools, including discussions about the variation between the distinct static measures, can be found in Appendix A.1.1.

1.3 Identification

In this section, I discuss issues of identification related to the estimation of R&D spillovers. I take the previously-defined spillover measures into account into equation 1.2 and use a log-linear form for the relationship of interest, in order to make my estimates comparable with the literature.³² I study the reduced-form relationship

$$\ln A_{it} = \alpha_1 \ln RDS_{it} + \alpha_2 \ln SpillCit_{it} + \alpha_3 \ln SpillSic_{it} + \alpha_4 \ln SpillDyn_{it} + \alpha_5 X_{it}^A + \nu_{it}, \quad (1.10)$$

where A_{it} is the outcome variable of interest for firm i at time t ,³³ the main variables of interest are own R&D stock RDS_{it} and the spillover terms, X_{it}^A are controls for each outcome A , and ν_{it} is the error term. The main issues to address in estimating equation 1.10 relate to unobserved heterogeneity, common shocks, endogeneity of R&D decisions, and endogenous network formation.

First, I condition out a set of firm dummies and year fixed effects in all my specifications. This takes care of time-invariant firm heterogeneity, as well as any economy-wide time-varying heterogeneity and shocks. In some specifications, I also include a full set industry-times year fixed effects.³⁴ These take care more flexibly of any industry-level time-varying shocks, and ensure that I compare similar firms in my analysis, within narrowly-defined industries and years.

³²See Bloom et al. (2013), Jaffe (1986), and Manresa (2016).

³³The two main outcomes of interest considered in this paper are firm value and productivity.

³⁴Industry is defined at the two-digit SIC code for these fixed effects, and at a narrower three-digit level, depending on the specification.

1.3.1 Endogeneity of R&D

Second, R&D spending is endogenously determined, due for example to transitory shocks to the profitability of research activity also affecting the outcome variables directly. This endogeneity problem will affect both own R&D and the R&D in the spillover terms. These transitory shocks should be mostly dealt with through the industry-times-year fixed effects. Nonetheless, I also use supply-side shocks to the user cost of R&D capital brought upon by state-specific and time-varying R&D tax treatments to instrument for corporate research expenses. I obtain firm-year specific R&D tax prices from Bloom et al. (2013), who calculate them using Wilson (2009)'s data on state and federal tax treatment of R&D expenses across time in the US. These tax R&D prices are affected by research tax credits and corporate income tax rates, at state and federal levels, as well as by the changing definition of what constitutes qualified R&D for existing tax credits.

There are two main components to the firm-year specific tax prices of R&D. First, a given firm will be exposed to different R&D tax treatments depending on the distribution of its research activity across the different US states. This distribution is proxied using the distribution of its patenting activity.³⁵ A firm's exposure to a state-specific tax treatment is calculated by its 10-year moving average share of patents filed from that state.³⁶ The second component takes advantage of federal rules to define what counts as qualified R&D for the existing tax credits and how large these credits can be. First, the federal rules have varied across time and have entailed an incremental component, where only R&D above a base level is eligible for tax credits. From 1981 to 1989, the base was the maximum of a rolling average of the previous three years' R&D. From 1990 onward (except 1995 and 1996, when the tax credit lapsed), the base was fixed to be the average of the firm's R&D to sales ratio between 1984 and 1988, multiplied by current sales (up to a maximum of 16%). Second, if the credit exceeds the taxable profits of the firm, it cannot be fully claimed and must be carried forward. With discounting, this leads to a lower implicit value of the credit for tax exhausted firms.

I project the endogenous R&D expenses on the two components of the tax instru-

³⁵This will be a good proxy as long as the patenting activity of a firm across the different states is closely related to its distribution of R&D expenses across states.

³⁶Before September 2012, patents could only be issued to human inventors. As a result, patents are originally filed on behalf of the inventor employee, and subsequently assigned to the firm before being granted. The addresses of the inventors are recorded in the applications, which are then used to distribute patents across the different states.

ment for the firms in the final sample between 1980 and 2001, and show the results in Table 1.3. Column (1) shows the basic results, column (2) adds year fixed effects, column (3) additionally includes firm fixed effects, and column (4) adds industry-times-year fixed effects. The instruments have considerable power in all specifications, with all the F -statistics above 28. I discuss the suitability of these instruments and the relevant exclusion restrictions in Appendix A.2. From the specification in column (3) with firm and year fixed effects, I calculate the level of R&D expenses R_{it}^{Tax} predicted by the tax regressors and then generate the stock RDS_{it}^{Tax} using perpetual inventory methods as in the true R&D case. This predicted R&D stock is then used to instrument for R&D stocks.

1.3.2 Endogenous network formation

Finally, I tackle the endogeneity inherent to the construction of the dynamic proximity metrics. The knowledge flows identified by the citation network may be correlated with the underlying quality of the researchers, with more knowledgeable researchers both producing better patents and also being aware of better or more research-intensive relevant art to build on and cite. The structure of the citation network underpinning the dynamic proximity matrix is thus likely to be endogenous. I find a suitable instrument by taking advantage of the observed persistence in the network structure. For a given firm i , I calculate its propensity to cite any other firm j between 1976 and 1984. I then use that propensity to predict the likelihood of firm i to cite j up to 25 years later, between 1987 and 2001.³⁷

Within the 715 firms originally in the network for which the static spillovers can also be constructed, 340 are found to be originating citations in both the 1976-1984 and the 1987-2001 networks. For those firms, and using the pre-period network, I average over $t \in [1976, 1984]$ the citation propensity $d_{it,jt-\tau}$ for all i, j and τ . That is, I calculate the average propensity for firm i to cite j 's research τ years prior in the past and denote it $\tilde{d}_{ij\tau}$. I then project the 1987-2001 network unto the 1976-

³⁷In choosing the 1984 and 1987 cut-offs, a number of forces are at play. First, I aim to include as many years as possible in the pre-period network to obtain a more precise estimate of the propensity to cite between firms. Second, I want to leave a large enough gap between the pre-period network and the current network, so as to ensure that the exclusion restriction holds. Finally, I want to include as many years as possible in the current network, to leave time for dynamic spillovers to accrue and to have a convenient time dimension for the panel analysis. I have carried out a variety of robustness tests using different cut-offs, and the results hold. I report one such test in subsection 1.4.3.

1984 network for these 340 firms. That is, for all nodes it and jt' with $t \geq t'$ in the 1987-2001 network, I regress

$$d_{it,jt'} = \eta_{it} + \mu_{jt'} + \phi \tilde{d}_{ijt-t'} + \epsilon_{it,jt'}, \quad (1.11)$$

where $d_{it,jt'}$ is the weight of the edge $it \rightarrow jt'$, $\tilde{d}_{ijt-t'}$ is the propensity for i to cite j 's patents filed $t-t'$ years earlier in the pre-period network, and η and μ are node-specific (i.e. firm-year-specific) fixed effects. The results, with different sets of fixed effects, are shown in columns (1) to (3) of Table 1.4, all of them with a large F -statistic. This strong network persistence is also observed by Acemoglu et al. (2016) in a more aggregated network of patent citations across technology classes rather than firm-year nodes. I use the specification with both sets of node-specific fixed effects as in column (3) to construct a predicted 1987-2001 network with edges $\hat{d}_{it,jt'} = \eta_{it} + \mu_{jt'} + \hat{\phi} \tilde{d}_{ijt-t'}$.

Using these predicted weights, I can construct an instrument for the dynamic spillovers, calculated in the same way as the dynamic spillover metric but with the predicted dynamic proximity weights rather than the true weights. This instrument takes care of the endogeneity in the decision of what to cite, but does not account for a possible endogeneity in the past R&D spending decision. Since this R&D decision is, by construction of the dynamic spillover network, carried out mainly in the past, I expect endogeneity concerns to be lessened. Nonetheless, for additional robustness I construct a dynamic spillover instrument *DynTax* by combining the predicted dynamic proximity network with the tax-predicted R&D constructed above.

Likewise, I take into account the endogenous citation and R&D decisions in constructing an instrument for the static knowledge spillover *SpillCit*. For the 340 firms found in both the subsequent 1987-2001 network and the pre-period 1976-1984 network, I calculate the average proximity measure \tilde{cit}_{ij} of citation propensity between firms in the pre-period network. I use this measure, together with firm fixed effects, to predict the citation proximity in the subsequent network and construct \hat{cit} . I regress

$$cit_{ij} = \eta_i + \mu_j + \psi \tilde{cit}_{ij} + \epsilon_{ij}, \quad (1.12)$$

where cit_{ij} is the citation proximity measure $i \rightarrow j$ using the subsequent networks, and \tilde{cit}_{ij} is the same measure in the pre-period network. Coefficients of this regression are shown in column (4) of Table 1.4. Finally, I combine this predicted proximity with the tax-predicted R&D to construct a static knowledge spillover instrument *CitTax*.

Last, I follow Bloom et al. (2013) in using the actual *sic* proximity weights together with the tax-predicted R&D stocks to construct an instrument *SicTax* for the business stealing spillover measures.

1.4 Empirics

Now that the general shape of the functions to empirically estimate is established, and with the previously-defined identification strategy in mind, in this section I turn to writing down the specific equations that I estimate and to discuss the main empirical results, robustness tests, and the effects on innovation-related outcomes. In order to quantify the effects of R&D spillovers, I focus on their influence on multiple firm-level outcomes such as productivity and market value, as in Griliches et al. (1991), and Bloom et al. (2013).

1.4.1 Productivity and market value equations

In order to investigate the effects of R&D spillovers on productivity, I estimate a log-linearized Cobb-Douglas production function, extended to account for both own R&D and possible spillovers:

$$\ln Y_{it} = \nu_i^Y + \delta_t^Y + \phi_{rds} \ln RDS_{it-1} + \sum_D \phi_D \ln SpillD_{it-1} + \phi_5 X_{it-1}^Y + \epsilon_{it}. \quad (1.13)$$

In this regression, Y is output (sales), $SpillD$ are the three measure of dynamic, static knowledge, and business stealing spillovers, and RDS is R&D stock.³⁸ The key control variables in X^Y are logged fixed capital and employee counts. In order to mitigate endogeneity concerns due to the simultaneity of decisions on endogenous variables, I lag the regressor variables compared to the dependent variable. That is, I look at how the regressors affect output one year later. I also include firm and year fixed effects, as well as dummies for zero R&D stock and no spillovers, and a 3-digit SIC industry-specific price deflator in order to account for price effects. In robustness exercises, I also include a more flexible set of fixed effects through 2-digit SIC and even 3-digit SIC industry-times-year dummies. Finally, I control for total industry

³⁸R&D stock is calculated using yearly R&D expenses and a perpetual inventory method.

output and lagged output to account for industry-wide dynamics,³⁹ and I include the price deflator in the right-hand side to allow for a more flexible relationship between revenue and prices.⁴⁰

I also analyze how own R&D effort and spillovers influence market value, by investigating a linearization of the value function introduced by Griliches (1981), again augmented by the spillover terms:

$$\ln \text{MTB}_{it} = \nu_i^Q + \delta_t^Q + \psi_{rds} \ln (RDS/AT)_{it} + \sum_D \psi_D \ln \text{Spill}D_{it-1} + \psi_5 X_{it}^Q + \epsilon_{it}^Q, \quad (1.14)$$

where MTB is defined as the market-to-book ratio (or Tobin’s average Q).⁴¹ This equation seeks to explain deviations to Tobin’s Q from unity through the R&D intensity of firms, and the spillovers they receive stemming from others. I include firm and year fixed effects in the regressions, to account for unobserved heterogeneity, and current and lagged log of total industry output, in order to account for industry-level dynamics that might influence the market valuation of firms. In robustness specifications, I also include a sixth-order series expansion in $\ln (RDS/AT)$, to control flexibly for R&D intensity.

Both the dependent variable and the R&D on the right-hand side are normalized by measures of book value of assets, so we can be worried about bias if these measures are noisy proxies for the actual variables of interest. This worry is attenuated in my specification, because the book value of assets used in constructing the MTB ratio is different than the total assets used in normalizing the R&D intensity. Nonetheless, in order to avoid the bias, I separate the relevant log of R&D intensity in two separate terms, log of R&D stock and log of total assets. I also estimate robustness tests with

³⁹Total industry output measures are constructed as follows. For each industry, data on total shipments are collected from the NBER-CES Manufacturing Industry Database and the BEA and deflated appropriately. A firm-year-specific measure is then constructed as a sum over all industries in which it operates (from Compustat Segments), weighted by its share of sales in each industry. The results are robust to using only the main four-digit SIC-code reported for each firm by Compustat instead.

⁴⁰Results are robust to using deflated sales more straight-forwardly as outcome variable.

⁴¹Market-to-book ratio is defined as market equity over book equity, where market equity is price at year’s end times shares outstanding and book equity is the book value of common equity, plus balance sheet deferred taxes and investment tax credit (if available). See Davis, Fama, and French (2000), or Kenneth French’s website for a complete description of the variable construction. Results are robust to alternative specifications of market value over assets. See Appendix A.4 for a detailed discussion.

R&D normalized over total sales. The coefficients on the spillover variables of interest remain robust across specifications.

The data used for this analysis is described in subsection 1.2 above, and the relevant descriptive statistics are shown in Table 1.1 for the years 1990-2001 as used in the empirical analysis. I have a relatively balanced panel with an average of 10.7 observations per firm in a panel of 12 years. The firms in this sample are on average larger, more valuable, and carry out more research than the wider sample of Bloom et al. (2013). In Appendix A.3.1, I study the implications for a given firm of a non-tournament model of innovation in which the production of new ideas results from three inputs: R&D effort by the firm itself, R&D effort by neighboring firms, and past innovation upon which to build. The qualitative predictions on the spillover effect signs and magnitudes are shown in Table 1.2. Among others, business stealing spillovers are expected to have effects only through prices, and thus not affect quantity-based productivity. In practice I analyze revenue-based productivity,⁴² and because of the lack of data on firm-specific prices and the use of industry-level price deflators there may be measurement error in using deflated sales as a measure of output in the productivity equation. As a result, it is hard to disentangle the effects of R&D spillovers specifically through prices or through quantities. I therefore do not seek to differentiate effects of R&D spillovers through prices or through quantities, as I believe the data is not appropriate for such a fine-grained distinction. The predictions for the effects of R&D spillovers on productivity or on market value are qualitatively equivalent, except for the business stealing spillovers which should not influence productivity, but may still lead to depressed revenue through lower prices.

1.4.2 Baseline results

The baseline results for both the market value equation and the productivity equation are shown in Table 1.5. The regressions shown include the dynamic spillovers as well as the static ones, and the standard errors are clustered twoway at the firm and year level. Columns (1) to (5) correspond to the productivity equation 1.13, whereas columns (6) and (7) show the coefficients of the market value equation 1.14. The reported F-tests correspond to the Kleibergen and Paap (2006) rk Wald F statistic of weak instruments. When available, the Stock and Yogo (2005) 10% critical values

⁴²See Syverson (2011) for a detailed discussion on revenue- and quantity-based productivity measures.

for weak instruments based on size are below the reported F-tests, rejecting that the instruments be weak.⁴³

I first estimate an OLS specification without controls, regressing deflated or "real" sales as a measure of output on own R&D stock, the dynamic spillover measures, and dummy variables for no R&D and no spillovers. In column (1), coefficients on own R&D as well as on the static knowledge spillovers are positive and significant. The point estimate on business stealing is negative, although not significant. Incorporating dynamic spillovers in column (2) mostly does not affect the other coefficients. Dynamic spillovers are positive and significant at the 1% level, and about five times smaller in magnitude than the static knowledge spillovers. In terms of magnitudes, all these coefficients correspond to elasticities: an increase of 10% in past R&D upon which to build is associated to an increase in output of 0.5%. I include all the control variables, including fixed capital and employee count, in column (3). The resulting equation is then closer to estimate the effect of own R&D and spillovers on firm productivity. Including the controls, and in particular controls for the size of the firm in terms of capital and labor, reduces the coefficient on own R&D drastically. The coefficients on both static and dynamic knowledge spillovers are also reduced, but remain statistically significant. I incorporate 2-digit SIC industry-times-year dummies for my preferred OLS specification in column (4). These account more flexibly for unobserved industry-level shocks and ensure we compare similar firms within closely-related industries. The static knowledge spillovers are reduced with the inclusion of these fixed effects, which suggests that they may be picking up common shocks that affect all firms within an industry. The dynamic spillovers remain robust.

Column (5) reports the coefficients of the 2SLS specification. I use the combination of the past network and the tax-predicted R&D as instrument for the dynamic spillovers. I instrument own R&D stock with tax-predicted R&D stock, and the static spillovers are also instrumented. The coefficients on both knowledge spillovers are positive and significant, and are more than doubled relative to the OLS specification. The coefficient on own R&D also increases, and is now marginally significant. The coefficient on dynamic spillover is somewhat larger than that on own R&D, and slightly less than half that on static knowledge spillovers. The increase in coefficients is consistent with positive shocks to firm sales or profitability leading to reduced in-

⁴³This threshold indicates that the conventional 5%-level Wald test will have an actual size of maximum 10%. That means that we can reject the null hypothesis of weak instruments if we are willing to tolerate a rejection rate of 10%.

novative effort, possibly because of firms focusing their attention and resources on the current available opportunities rather than search for new ideas, processes or products. This would affect both own R&D and dynamic spillover coefficients. It can also be indicative of measurement error in both measures of own R&D and dynamic spillovers. Finally, in the 2SLS specification the coefficient on business stealing spillovers increases in magnitude and becomes statistically significant. Business stealing spillovers are nonetheless smaller in magnitude than both knowledge spillovers.

Column (6) of the same Table 1.5 reports the coefficients on the OLS specification of the market value equation including all the controls. The coefficient on the dynamic spillovers is once again positive and significant, while the coefficient on own R&D stock and other spillovers are statistically indistinguishable from zero. Instrumenting for spillovers and R&D in column (6) more than triples the coefficient on dynamic spillovers, just as in the output equation. The coefficient on static knowledge spillover also increases markedly and becomes significant, whereas business stealing effects become negative and significant once again. As for own R&D, even though the point estimate is markedly increased in the 2SLS specification, it remains statistically insignificant. Although the lack of strong relation between own R&D and measured productivity or market value is at first sight a surprising result, this relationship in the literature is only strong and consistent in the cross-section. Using longitudinal data with firm fixed-effects, the relationship has been found to be fragile in the past.⁴⁴ In my sample, the same robust relationship is found in the cross-section. When including firm fixed-effects, and especially in the 2SLS regressions, the relationship holds marginally for productivity, but breaks down for the market-to-book ratio.

The estimates on the spillover terms suggest that although the static measure does a good job of detecting knowledge spillovers, a sizable share of additional spillovers is still accounted for through the novel dynamic measure. In Appendix A.4, I use other traditional static spillover measures such as that introduced by Jaffe (1986), and find similar results. This indicates that the static measures are not sufficiently flexible to account for complex dynamic structures of knowledge flows and can prove especially problematic for technology streams that rely heavily on cumulative innovation processes, whose knowledge flows may be particularly prone to not be detected using static spillover measures. In section 1.5 below, I study how the differential dynamic structure of knowledge flows across industries and technology types is reflected in the

⁴⁴See e.g. Klette and Kortum (2004).

relative importance of static versus dynamic spillovers.

1.4.3 Robustness of dynamic spillovers

The results regarding dynamic spillovers of R&D are robust across a variety of different specifications. I report the coefficients on dynamic and own R&D across a battery of robustness checks in Table 1.6. Along these tests, I report both the OLS coefficients and, when applicable, the 2SLS coefficients using the dynamic spillover instrument as well as the own predicted R&D instrument. First, I introduce a more flexible set of fixed effects through 3-digit SIC industry-times-year dummies in columns (1) and (2). These fixed effects take care of unobserved shocks at a fine-grained industry-year level, and lead to a more specific comparison of similar firms within narrower industries. Both the OLS and 2SLS results remain robust, in both the productivity and the market value equation. However, the importance of own R&D relative to the dynamic spillovers increases in the 2SLS specifications. I then control flexibly for own R&D using a sixth-order polynomial in columns (3) and (4) to ensure that results are not driven by non-linearities.⁴⁵ Again, the coefficients on dynamic spillovers remain robust.

I also restrict the analysis to manufacturing firms, which comprise the largest portion of my sample, to ensure that the coefficients are not driven solely by outliers among non-manufacturing firms. I then control flexibly for patent counts using a fourth-order polynomial in columns (7) and (8),⁴⁶ in order to ensure that my results are not driven by non-linearities in the effect of patents on firm outcomes. I also include citation-weighted patent counts in the regression, to control for the observed quality of the patents filed by each firm in each year. Including this control variable reduces the effects of dynamic spillovers substantially. Nonetheless, I believe it constitutes a "bad control", as it corresponds in fact to an outcome variable. In subsection 1.4.4 below, I show that dynamic spillovers lead to increased citation-weighted patent counts. Therefore, controlling for citation counts blocks out a possible spillover channel, as dynamic spillovers may accrue precisely by leading to production of higher-quality subsequent innovation.⁴⁷ In columns (11) and (12), I use an alter-

⁴⁵I also estimate unreported specifications with larger order polynomials and find that a sixth-order polynomial is enough to control flexibly for R&D. I do not instrument for own R&D in the 2SLS regressions that include flexible polynomial controls in own R&D.

⁴⁶I also estimate unreported specifications with higher-order polynomials of log-patent count, and find that a fourth-order is enough to control flexibly for patent counts.

⁴⁷Where quality of ideas is understood as "the magnitude of inventive output associated with

native definition of R&D intensity as R&D stock over total sales rather than assets for the dynamic spillover construction. This way, I can ensure that results are not driven by the specific structure of the normalization parameter itself. Results remain robust. In unreported specifications, I also use a non-normalized proximity measure from the patent citation network to construct the dynamic spillover measures, and find qualitatively similar results.

In columns (13) and (14), I use a rate of 10% for the depreciation of R&D stocks and estimate a specification including the static spillovers. This lower obsolescence rate makes R&D and its static spillovers longer-lived, but does not affect the coefficient on dynamic spillovers significantly.⁴⁸ This is consistent with dynamic spillovers not just picking up a mis-specification in the actual obsolescence rate used for the depreciation of knowledge, but rather channeling spillovers that accrue through cumulative mechanisms. Or in other words, knowledge spillovers depreciate slower only when past innovation is actually used as a foundation on which to build upon further. Finally, in the last column, I use the full patent citation network between 1976 and 2001 to analyze dynamic spillovers, with an extended panel of 532 firms between 1980 and 2001. Due to my identification strategy, I cannot estimate any 2SLS specifications on this larger sample, but the OLS specifications show large and significant dynamic spillover effects, with the magnitude of these spillovers being similar to those in the reduced sample.

1.4.4 Effect on innovation outcomes

In order to shed light on the mechanism through which spillovers affect firm productivity and market value, in this subsection I analyze the effects of the spillovers on innovation-related outcomes such as R&D expenditures and citation-weighted patent counts. I posit that the dynamic spillovers' effect on productivity and market value occur through the innovation production function. Because spillovers increase returns to R&D, I expect dynamic spillovers to increase R&D expenditures and innovation production in terms of patents. The dynamic spillover measure itself is influenced by the degree of innovative activity taking place, because a given firm-year will be measured as receiving spillovers only if it files patents that are citing prior art. As a result, I do not use *SpillDyn* as a regressor and estimate instead a reduced-form IV

them" (Griliches, 1990), and is positively associated to citation counts.

⁴⁸Unreported coefficients on the static knowledge and business stealing measures are also similar to the baseline estimates.

specification in which I include the dynamic spillover instrument *DynTax* directly in the right-hand side. That is, I study how the *potential* for large dynamic knowledge spillovers, measured using the pre-period citation network and the tax-predicted R&D, affects R&D expenditure decisions and innovative output. If in the pre-period network, firm i used to build heavily upon firm j 's innovation τ years earlier, I analyze how a predicted increase in R&D spending by firm j in year t influences firm i 's innovation in year $t + \tau$.

I estimate the following equations:

$$\ln P_{it} = \mu_1 \ln \text{DynTax}_{it-1} + \mu_2 \ln RDS_{it-1} + \mu_3 X_{it}^P + \nu_i^P + \delta_t^P + \epsilon_{it}^P, \quad (1.15)$$

$$\ln RD_{it} = \gamma_1 \ln \text{DynTax}_{it-1} + \gamma_2 \ln Y_{it-1} + \gamma_3 X_{it}^R + \nu_i^R + \delta_t^R + \epsilon_{it}^R, \quad (1.16)$$

where P is citation-weighted patent count, RDS is R&D stock, RD is R&D expenses, and Y is total revenue to account for firm size,⁴⁹. Controls X^R for the R&D equation and X^P for the patent equation include current and lagged industry-wide sales to account for industry-level dynamics, and a dummy for no dynamic spillovers. Both equations also include firm and year dummies, and I cluster standard errors twoway at the firm and year level. The coefficients are shown in Table 1.7. I estimate equation 1.15 in column (1), and find that the potential for dynamic spillovers increases patent counts significantly. An increase of 10% in past R&D to potentially build upon is associated to an increase of 5.2% in the citation-weighted patents filed. The relationship is slightly reduced but remains strong and significant when including a lagged dependent variable. In column (3), I estimate a negative binomial count model with bootstrapped standard errors, and the results are largely unchanged.

I estimate equation 1.16 and show the coefficients of interest in columns (4) and (5). An increase of 10% in past R&D to potentially draw from is associated to an increase of 0.4% in the R&D expenses. There is strong persistence in R&D activity within firms, but the dynamic spillovers retain a strong and significant coefficient in column (5). The results shown in this table are consistent with dynamic spillovers being an integral part of the innovation production function, and leading to increases

⁴⁹Results are robust to controlling for total assets AT instead of revenue.

in R&D spending because of its increasing the returns on R&D activity.

1.5 Heterogeneity across industries and technology types

In the previous section 1.4, I found significant effects of both dynamic and static knowledge spillovers on productivity and market value. To this point, I have assumed that the magnitude of spillovers is constant across industries and technology types. However, knowledge spillovers can be affected by a number of factors, which may in fact influence the spillovers estimated through the dynamic and static measures differently. In this section, I examine how knowledge spillovers vary across industries and technology types, with a particular emphasis on the static versus dynamic knowledge spillover decomposition.

I first investigate how the dynamic and static knowledge spillovers vary across industries. Because of the limited size of my sample, it is difficult to separately quantify spillovers for every industrial sector. Nonetheless, it would be worrying if significant knowledge spillovers were not found within high-tech sectors, for example.⁵⁰ I first focus on the six sectors for which I have the largest number of observations within my sample:⁵¹ industrial and commercial machinery and computer equipment (SIC 35); electronic and other electrical equipment and components, except computer equipment (SIC 36); transportation equipment (SIC 37); measuring, analyzing and controlling instruments, photographic, medical and optical goods, watches and clocks (SIC 38); chemicals and allied products, except drugs (SIC 28, except for 3-digit SIC 283); and pharmaceutical firms (SIC 283).

Figure 1-3 plots the coefficients on dynamic and static R&D spillovers for the baseline productivity equation (across all industries) and for the particular industries mentioned above.⁵² The regression estimates industry-specific dynamic and static

⁵⁰58% of weighted edges in the patent citation network occur between firms belonging to the same 2-digit SIC code industry, so dynamic spillovers occur majoritarilly within industries. In Appendix A.5.1 I show that dynamic spillovers are statistically significant regardless of whether they accrue within or between industries, but the point estimates are larger in magnitude within.

⁵¹These six sectors happen to be high R&D intensive manufacturing sectors, in which I would expect knowledge spillovers to therefore be large.

⁵²In this section, I use the Jaffe (1986) technological proximity as a static knowledge spillover measure in order to make a more comparable distinction between the novel dynamic spillover measure and the measures used traditionally in the literature. Standard errors are clustered twoway at the firm and year level, and 90% confidence intervals are shown in the figure.

spillover coefficients on the subsample corresponding to the six industries mentioned above. I find that dynamic spillovers are larger than the baseline coefficient within Machinery, Electronics, Instruments, and Chemicals (without Drugs). However they are insignificant within the Transport sector, and the point estimate is even negative for Pharma. The pattern for the static knowledge spillovers seems to be reversed, with large and statistically significant coefficient in Pharma and Transport, while the estimates for Machinery, Electronics, Instruments, and Chemicals are statistically insignificant and the point estimates are generally smaller than the baseline level.

In order to understand these patterns better, I consider an industry classification from Coad and Rao (2008), which defines 2-digit SIC industries 35, 36, and 38 as high-tech industries that can be classified as *complex* product types; industry 37 as a mature complex technological sector; and industry 283 as a high-tech industry that is nonetheless classified as *discrete*. The distinction between complex and discrete technologies follows Levin et al. (1987), Cohen et al. (2000) and Galasso and Schankerman (2015) in differentiating "whether a new, commercializable product or process is comprised of numerous separately patentable elements versus relatively few". I discuss this characterization in more detail below. According to this classification, it seems that complex high-tech complex industries exhibit coefficients on dynamic spillovers that are larger than the baseline, whereas the mature Transport industry as well as the simple high-tech Pharma industry show small and insignificant coefficients. In particular, I can reject that the dynamic spillover coefficients in Pharma and Transport are equal to any of the complex high-tech coefficients with 10% confidence in the productivity equation. This is consistent with dynamic spillovers accruing primarily in high-tech industries, which are more reliant on innovation, but mainly within complex technologies that are intrinsically more cumulative.

The relative importance of dynamic versus static knowledge spillovers seems therefore to be affected by the industrial sector one considers. In particular, for the Chemicals sector, the point estimate on the elasticity of output to dynamic spillovers is even larger in levels than the coefficient on static spillovers, as compared to a baseline ratio of almost 14 to one in favor of the static spillovers. This is indicative that considering only static knowledge spillovers, as the R&D spillover literature has traditionally done, not only leads to underestimating R&D spillovers homogeneously across the productive sector. It actually will entail heterogeneous biases across industries, which can hinder our understanding of the mechanisms through which spillovers

accrue, and lead to faulty policy recommendations.

In order to further study the heterogeneity found above using a wider industry spread, I classify manufacturing industrial sectors according to their product type. First, I build upon Levin et al. (1987), Cohen et al. (2000), and Hall et al. (2005) in classifying industries into complex and discrete products. Discrete products rely on few patents and the importance of patents for appropriability of returns to innovative activity has traditionally been larger. In fact, "industries with discrete products tend to patent for the traditional reasons of excluding competitors and preventing litigation, whereas those in complex product industries are significantly more likely to patent for cross-licensing and trading/negotiation purposes, as well as to prevent litigation."⁵³ Industries that correspond to discrete products are in SIC codes 2000-3499, and include food and tobacco, textile, wood and paper, chemicals, plastic and metals. Discrete industries include high-tech sectors such as drugs and pharmaceuticals (SIC 283), or chemicals (SIC 28). On the other hand, industries with SIC codes between 3500 and 3899 correspond to complex products.⁵⁴ These include machinery, electronics, transportation equipment, and instruments.

A related way to classify the industrial sectors is according to their technology base, depending on whether they primarily innovate on electrical, mechanical, or chemical technologies.⁵⁵ This classification relates to the technology type, rather than the product, but is nonetheless closely related to the complex versus discrete partition. Electrical technology sectors are all included within the complex product type, and correspond to medical instruments, computing and electronic machinery, instruments, and aero-space equipment.⁵⁶ Chemical technology sectors are all included within the discrete product type, and correspond to food and tobacco, chemicals, oil and plastic, and stone, clay and glass.⁵⁷ Mechanical technology sectors are partitioned between discrete and complex product types, and include textiles, paper and wood, metals, machinery and engines, and transportation equipment other than aero-space.

⁵³Hall et al. (2005), page 9.

⁵⁴SIC codes 39xx correspond to "Miscellaneous Manufacturing Industries" and are not considered into any of the categories. The highest considered manufacturing SIC code is thus 3873. For more details about this classification into complex and discrete technologies, see Hall et al. (2005).

⁵⁵This classification is also based on Hall et al. (2005).

⁵⁶The SIC codes are 357x, 36xx, 372x, 374x, and 38xx. I also include codes 737x that are not included in manufacturing, but correspond to *Computer and Data Processing Services* firms such as IBM.

⁵⁷The 2-digit SIC codes are 20, 21, 28, 29, 30, and 32.

I show the coefficients on dynamic and static knowledge spillovers for the productivity equation across both classifications of the manufacturing industries in Figure 1-4. Dynamic spillovers are larger in complex products than in discrete products, where they are not statistically significant. As for the static knowledge spillovers, the opposite is true. When partitioning by technology type instead, dynamic spillovers are large and statistically significant for electrical and mechanical technologies, while smaller and statistically indistinguishable from zero in chemical technologies. For the static knowledge spillover estimates, once again the opposite pattern is visible.

It therefore seems like industries with complex products involve relatively large dynamic spillovers and relatively low static knowledge spillovers, as compared to the baseline levels. Meanwhile, the opposite is visible in discrete products and particularly so for chemical technologies, with relatively large static knowledge spillovers and small dynamic spillovers. An explanation consistent with these results is that the innovative process in complex products is inherently more cumulative, that the innovation production function depends more on past foundational knowledge to build upon. As a result, the spillover flows are inherently more inter-temporal in nature and are therefore picked up by the dynamic spillover measure rather than the static one. Another related explanation leans on the differential use of patents for appropriability reasons across product types. If the use of patents as exclusive mechanisms is more prevalent across discrete product types, intellectual property rights may hinder cumulative innovation.⁵⁸

Finally, in order to obtain more evidence that the heterogeneous dynamic spillovers are indeed related to the technology type, I investigate whether this heterogeneity holds at a more fine grained originating patent technology class level. That is, I classify patents according to their technological area into complex and discrete types, and study whether dynamic spillovers stemming from complex patents differ from those stemming from discrete patents. I follow Galasso and Schankerman (2015) in categorizing *Computer & Communication* (Hall et al. (2001) technological category 2), *Electrical & Electronics* (category 4), *Medical Instruments* (subcategory 32), and *Biotechnology* (subcategory 33) into complex technologies, comprising innovation that is "highly cumulative and requires the input of a large number of patented components held by diverse firms". Galasso and Schankerman (2015) find that decreases

⁵⁸See Scotchmer (1991) and Williams (2013).

in patent protection, which are likely to lead to increases in dynamic spillovers, are more conducive to new R&D investment and innovation in complex technologies. It is important to remark that I do not in my setting analyze heterogeneity in the levels of IP protection, and that my results on dynamic spillovers are found *conditional* on the original patent being granted. In other words, I find dynamic spillovers even though the initial innovation is protected by patents.

In order to compare dynamic spillovers in discrete versus complex technologies, I separate the dynamic spillover measures depending on the technology class of the originating patent. I estimate the productivity and market value equations 1.13 and 1.14 including both types of dynamic spillovers, and plot the resulting coefficients in Figure 1-5 for both the OLS and the 2SLS specifications. The dynamic spillovers originating in complex technologies are consistently positive and significant, whereas patents from discrete technologies show no statistically significant dynamic spillovers. Therefore, I find that complex technologies result in large dynamic spillovers whereas these do not appear significant among simple technologies. This is consistent with the results shown previously, and indicates that the prevalence of dynamic inter-temporal spillovers within complex products is robust. In terms of cross- versus within-technology class distribution, most spillovers accrue within technology class as 77% of weighted citations in the patent citation network occur within technology class. Moreover, I show in Appendix A.5.2 that dynamic spillovers are statistically significant only within technology classes.

1.6 Extensions

In this section, I present a number of extensions to my empirical analysis. First, I look for evidence of second-degree spillovers accruing through indirect network connections. Second, I estimate the importance of geographic distance as well as customer-supplier relations for spillovers. Third, I study whether dynamic spillovers depreciate across time. Fourth, I analyze how the size of originating firms influences dynamic spillovers. Finally, I examine dynamic spillovers stemming from start-ups funded by Venture Capital.

1.6.1 Second-degree spillovers

I consider the possibility that dynamic spillovers also accrue through indirect connections. That is, suppose that firm-year kt'' builds upon research by jt' , which itself is citing it 's research. But there are no direct citations between kt'' and it . As R&D from firm-year it contributes to the creation of subsequent knowledge by jt' , does it also contribute to the 1-step-removed production of ideas by firm-year kt'' ? In order to investigate, I build an indirect dynamic spillover measure using second-degree connections.⁵⁹ I estimate the productivity and market value OLS regressions including both the usual direct dynamic spillover measure and the indirect second-degree measure. Results for the two coefficients are shown in Figure 1-6. Both the direct and indirect spillovers are positive and significant, for both the productivity and the market value equation. In terms of market value, the coefficients on direct and indirect dynamic spillovers are of similar magnitude. As for the productivity equation, the coefficient on indirect spillovers is about half that on direct spillovers. Results are thus consistent with own R&D and past knowledge being used to produce subsequent knowledge (proxied by the productivity equation), which is itself used to produce future knowledge.⁶⁰ As such, dynamic spillovers may reverberate through the citation network and continue accruing as they do so.

1.6.2 Geographic spillovers and customer-supplier relations

In this subsection, I study how the relationship between the originating and absorbing firms influences the magnitude of dynamic spillovers. I analyze in particular two dimensions, geographic co-location and customer-supplier linkages within production networks. First, I study how dynamic spillovers vary with geographic distance. An extensive literature has studied the influence of geographic distance on knowledge flows and mainly found that knowledge flows decrease with distance and geographic or political frontiers.⁶¹ I go further in examining the extent to which geographic distance influences spillovers of R&D given that a citation is observed, i.e. given that a knowledge flow is observed. I use the address of the main inventor of each patent

⁵⁹If the patent citation network is defined by the adjacency matrix Dyn , then the matrix of second-degree connections is Dyn^2 .

⁶⁰For the OLS specification with controls used, as in column (3) of the baseline table 1.5, the elasticity of dynamic spillovers is about 64% of that on own R&D. The ratio between the elasticity on direct and indirect spillovers is similar.

⁶¹See Furman et al. (2006), and Jaffe et al. (1993).

to allocate the patent to one of the 50 US states or 195 non-US countries. Using this geographic allocation, I construct a measure of dynamic spillovers $SpillDyn^{GEO}$ similarly to the baseline dynamic spillover measure, except I only consider citations between patents within each of the 245 geographic locations. If geographic considerations were important for the magnitude of the dynamic spillovers, including $SpillDyn^{GEO}$ in the regressions should lead to significantly positive coefficients on this measure, and decreases in the coefficient on the baseline measure $SpillDyn$.

I estimate these regressions in Table 1.8, in columns (3) and (4) for the productivity equation and (7) and (8) for the market value equation. Columns (3) and (4) show that geographic considerations do not seem to influence the magnitude of dynamic spillovers, as the coefficient on $SpillDyn^{GEO}$ is small and insignificant, whereas the coefficient on the baseline dynamic spillover measure is not markedly reduced. Column (7) shows that geographic considerations seem to have a marginally significant effect on market value, but this disappears in the 2SLS specification. These results suggest that, conditional on citations occurring, the dynamic spillovers of R&D are not affected by geographic distance. On the other hand, static knowledge spillovers are found to decrease with distance by Bloom et al. (2013) and Lychagin et al. (2016). These apparently contradicting results can be explained by the fact that the Jaffe (1986) static spillover measures do not condition on knowledge flows being observed. Therefore, static spillovers may decrease in geographic distance because of a lower likelihood of interactions between scientists, while conditional on interactions the spillover magnitudes could still be sustained.

Second, I study how customer-supplier relations within a supply chain influence dynamic spillovers. I use data from Barrot and Sauvagnat (2016)⁶² on customer-supplier links for Compustat publicly-listed firms. Firms are required to disclose the identity of any customer that represents more than 10% of the total reported sales, and that information is used to construct a supplier-customer linkage network. A small share of around 1% of the citation network links occur between customers citing suppliers, identified through this dataset. For those connections, I construct an analogue $SpillDyn^{PROD}$ to the dynamic spillover measure using as the relevant weights the share of total sales that corresponds to the customer-supplier linkage.

I estimate in Table 1.8 in columns (2) and (6) regressions including both $SpillDyn$

⁶²I would like to thank the authors for sharing their data.

and $SpillDyn^{PROD}$. I find that for the subset of observations that cite their upstream suppliers, no dynamic spillovers accrue in terms of productivity as the sum of the two coefficients cancel out. In terms of market value, the sum of the coefficients is also not statistically different than zero. These results suggest that customers that innovate by building on top of their suppliers innovation do not absorb meaningful R&D spillovers. This can be due to citations between customers and suppliers being driven by customers having to adapt their technologies and production processes to the specifications of their suppliers, rather than innovating further. Citations would thus not be a result of cumulative innovation processes, but rather an adaptation to idiosyncratic product specifications on the part of the downstream customer-innovator.

1.6.3 Depreciation of spillovers

I next analyze how the dynamic spillovers vary across the span of time between the original innovation and the follow-on innovation. I am interested in examining the persistence of dynamic spillovers since the original innovation, to establish whether they are long-lived or only accrue for a short time. In order to do so, I separate the dynamic spillovers in three groups, depending on the time span they bridge. In terms of the subsequent 1987-2001 network, the first group is restricted to a time span up to three years, the second corresponds to four and five years' worth of span, and the third to spillovers that accrue six or more years later. The cutoffs correspond to the median time span of 4 years and the 75th percentile of 6 years. In terms of the full 1976-2001 network, the median is 5 years and the fourth quartile starts at 9 years.

The results are shown in Figure 1-7.⁶³ In the subsequent network, dynamic spillovers tend to decrease with time for the productivity equation, with the spillovers being insignificant after six years. One explanation consistent with these results is that using older and possibly more out-of-date prior knowledge reduces the quality of the innovative output. Nonetheless, this depreciation of dynamic spillovers does not hold in the full network, which shows spillovers accruing with a gap of more than 8 years to be still statistically significant and of similar magnitude as the baseline estimate. Moreover, it seems the decreasing relationship of spillovers with age is reversed for the market value equation for older spillovers. In fact, the majority of market value spillovers stem from older spillovers. A tentative explanation could involve the de-

⁶³The regression panel for the subsequent network is 1990-2001, whereas the full network panel estimates the equations on 1980-2001.

cision on which past innovation to build upon having product market competitive effects. Building upon more recent innovations might lead to increased competition between companies in the product market, therefore exhibiting larger effects in terms of productivity rather than firm value.

1.6.4 Size of originating firm

A number of R&D subsidies are explicitly targeted to small and medium-sized enterprises (SMEs), through programs such as the Small Business Innovation Research program in the US. As a result, it is particularly interesting to assess how the magnitude of dynamic spillovers may vary depending on the size of originating firms. I examine the heterogeneity in the dynamic spillovers across this dimension by dividing the 340 firms in my sample in two halves depending on their average size in terms of employment, over the period 1980-2000. I then separate the dynamic spillovers based on the size category of the firm in which they originate. I show in Figure 1-8 that larger firms are responsible for R&D that generates larger dynamic spillovers. These spillovers are on average about 2.8 times larger than the spillover from smaller firms.⁶⁴ When looking at static knowledge spillovers, Bloom et al. (2013) also find that the R&D carried out by larger firms is responsible for larger spillovers. Of course, technology spillovers are not the only justification for directed government intervention. Moreover, the median firm in the smaller half of my sample has an average of 2,026 employees over the time period, which hardly makes it a SME. Nonetheless, taken at face value my estimates suggest that larger firms should receive more generous subsidies if the government was interested in providing incentives for organizations to internalize technological spillovers.

1.6.5 Venture Capital spillovers

The research activity carried out by Venture Capital (VC)-backed start-ups has been found to be more effective at producing innovation than traditional corporate R&D (Kortum and Lerner, 2000; Hirukawa and Ueda, 2013). Moreover, Schnitzer and Watzinger (2015) show that static knowledge spillovers from this VC-backed R&D are also larger than corporate R&D. I examine here if this VC-backed R&D also results in increased dynamic cumulative spillovers on the publicly-listed firms that

⁶⁴Although the differences in coefficients between small and large firms are only significant for the 2SLS market equation, all four regressions show the same pattern.

build upon it. Because start-ups are small private companies, no data is available on their R&D expenditures. Instead, the amount of VC investment that these firms receive is used as a proxy (Kortum and Lerner, 2000) for their research expenditures.⁶⁵ I use VentureXpert data on VC funding and link start-ups to the NBER patent data by company name. I match 5157 start-ups with at least one patent between 1976 and 2001, of which 3418 are subsequently cited by a publicly-listed firm in my sample. I construct a citation network between publicly-listed citing firms and cited start-ups and use it as a weighting proximity matrix. The VC-backed dynamic spillovers are then the weighted sum of the R&D stocks of the cited start-ups.⁶⁶

I estimate the productivity and market value OLS equations including both the VC-backed dynamic spillovers and the baseline publicly-listed dynamic spillovers. I do so using both the subsequent 1987-2001 network and the full 1976-2001 network.⁶⁷ The resulting coefficients are shown in Figure 1-9, and suggest that VC-backed R&D does not exhibit any significant dynamic spillovers on publicly-listed firms' productivity. In terms of market value, the coefficients are significant and positive, with the full-network regression showing dynamic spillovers from start-ups and larger corporations to be similar. The results suggest that VC-backed innovative activity does not accrue dynamic spillovers in terms of increased productivity on larger corporation that build upon it. These corporations do however seem to receive spillovers on market value, which suggests that the spillovers may occur through prices only. Nonetheless, within my sample I can reject that dynamic spillovers from VC-backed innovation on Compustat firms are larger than those from traditional corporate R&D.

1.7 Estimates for the private and social returns to R&D

In this section, I estimate the private and the social returns to R&D using previously-estimated coefficients as well as estimates from subsection 1.7.1 below, where I sepa-

⁶⁵In order for VC funding to be a good measure for the R&D expenditures of start-ups, one must assume that start-ups have no other sources of funding and that they use the funding primarily on R&D activities. To the extent that start-ups also have other expenses such as marketing, VC investments may be over-estimating their research budget.

⁶⁶Where the weights are obtained from the Compustat-to-VC citation network, and the R&D stocks are constructed using VC funding as a proxy for yearly R&D expenses.

⁶⁷The panel using the subsequent network is 1990-2001, whereas the full network estimates the equations on 1980-2001.

rately analyze dynamic knowledge flows within and between firms.

1.7.1 Own vs. others' spillovers

In subsection 1.4.2, I analyzed the influence of dynamic spillovers irrespective of where they originate, and found that they are large and significant. However, these spillovers do not necessarily correspond to externalities. In fact, even though the patent citation network does exclude same-node citations (that is, citations within any given firm-year node), it will still include citations originating and ending in nodes belonging to the same firm at different points in time. That is, a citing node it can be connected through an edge to another node it' as long as $t' < t$. If innovating firms are rational, they will internalize the share of spillovers that accrue dynamically within themselves and will take them into account when choosing their optimal R&D investment level. Belenzon (2012) shows that patents that are subsequently self-cited by the original patent-holding firm are more strongly positively related to its market value, consistent with the view that future absorbed spillovers are internalized *ex-ante*. On the other hand, spillovers that originate in other firms are more likely to be externalities, and not be internalized by the original innovator.

Within the subsequent 1987-2001 network, around 37% of weighted edges are self-cites, with the remaining 63% being citations to other firms. I compute a measure of own dynamic spillovers just as the main dynamic spillover metric, but considering only own-cites as possible connections (intra-firm), and likewise for others' spillovers (inter-firm). I estimate the same baseline regressions except including both own and others' spillover terms rather than the main one. The results are shown in Table 1.9. All the regressions include firm and year fixed effects, as well as the preferred controls discussed in subsection 1.4.2. The OLS regression in column (1) considers only own R&D and the inter-firm dynamic spillovers, and the coefficients are very similar to the baseline case. Including within-firm dynamic spillovers does not change the existing coefficients, and the intra-firm coefficient is positive albeit not statistically significant. Nonetheless, the elasticity estimate on within-firm dynamic spillovers is likely to entail a substantial OLS downward bias, as the plausible endogeneity due to firms abandoning research programs in order to concentrate on the immediately available opportunities in the presence of positive revenue shocks is likely to affect profitable albeit more complex and uncertain projects, with a longer maturation period, and that are based on more in-house research of more separately patentable steps until

the final commercializable product. The coefficients on both inter- and intra-firm spillovers are not altered with the inclusion of static spillovers and industry-times-year fixed effects in columns (3) and (4). In the 2SLS specification in columns (5) and (6) all the coefficients become significant, and there is some suggestive evidence of the intra-firm spillover decreasing the returns on own R&D when included. Notice that in column (6), the coefficient on intra-firm spillovers is slightly larger than half of that on inter-firm spillovers. This is consistent with the share of dynamic spillovers that accrues within versus between firms, due to between-firm citations being twice as prevalent as within-firm cites. I include static spillovers in the 2SLS specification in column (7) to obtain a specification with all the spillover estimates, in order to use it for the social return calculation.

1.7.2 Calculating social and private returns

Using the coefficients estimated in the previous section on the three relevant types of spillovers of R&D, I calculate the private and social rates of returns to R&D for the sample of firms considered. This exercise assumes that the previously estimated coefficients can be used for policy purposes. That the estimated functional forms are correct, that the proximity measures correctly quantify knowledge flows, and that the coefficients are causal. These are strong assumptions, and the discussion below should therefore be considered more speculative.

First, and following Bloom et al. (2013), I define the *marginal social returns* (MSR) to R&D of firm i as the marginal increase in aggregate output due to a marginal increase in the R&D stock of firm i . I also define the *marginal private return* (MPR) as the respective marginal increase in firm i 's output to its R&D stock. The detailed derivation of the rates is discussed in Appendix A.3.2. The main assumptions used to derive the formulas are that all firms are fully symmetric and that there are no amplification effects of R&D. That is, I assume that firms all have the same size in terms of output and R&D stocks, and are identically connected to other firms in terms of their dynamic and static proximity measures. Likewise, I do not consider how an increase in the R&D stock of firm i can affect returns to the R&D of other firms, or own returns in the future, leading to amplification through additional increases of R&D effort throughout the economy. Under these assumptions, the social returns of R&D can be written as

$$\text{MSR} = \frac{Y}{RDS} (\phi_{rds} + \phi_{sic} + \phi_{cit} + \delta\phi_{dyn}), \quad (1.17)$$

where the coefficients ϕ refer to the coefficients in equation 1.13 for own R&D stock, dynamic spillovers, business stealing and static knowledge spillovers, and δ is the discount factor of the social planner that accounts for the fact that dynamic spillovers accrue in the future. I consider the coefficients from the production function rather than the value equation, because the latter will not only capture the productivity increase due to R&D. Instead, it also captures resulting changes in other input variables such as employment and capital, leading to a larger compounded effect. The MSR takes into account how increases in own R&D stock yield increases in own productivity, but also how they affect other firms' outputs through all the spillover terms.

Under the same assumptions, the private returns to R&D can be written as

$$\text{MPR} = \frac{Y}{RDS} (\phi_{rds} + \delta'\phi_{intra}), \quad (1.18)$$

where ϕ_{intra} represents the coefficient on within-firm dynamic spillovers that are internalized by the original innovator, and δ' is the discount factor of the firm.⁶⁸ The differences between the MSR and the MPR are as follows. First, both take into account the direct increase in a given firm's output due to increases in its R&D stock (ϕ_{rds}), but the social planner also considers the business stealing spillovers (ϕ_{sic}) and the static knowledge spillovers (ϕ_{cit}) generated by R&D. Second, the private returns only internalize part of the dynamic spillovers generated, equal to the share of spillovers that are re-absorbed by the firm (ϕ_{intra}). Third, the discount factors of the social planner and of the firm (δ and δ') may not be equal. In fact, due to possible financial frictions in firms, it is likely that $\delta > \delta'$.

The wedge between private and social rates of return to R&D depends on the magnitudes of the spillover terms, as well as the magnitudes of the parameters δ , δ' , and that of $\frac{Y}{RDS}$. In order to evaluate the expressions, I use the value of $\frac{Y}{RDS}$ for the median observation in my sample, at 5.8.⁶⁹ As for δ , within the subsequent network

⁶⁸The equations derived here diverge from the formulas of Bloom et al. (2013) in two main aspects. First, I include both dynamic spillovers and re-internalized dynamic spillovers. Second, I include the effect of business stealing spillovers on sales directly, rather than indirectly through the market value equation.

⁶⁹Another possibility is to consider the value of the average output over the average R&D stock, which is 4.5. These different values will affect rates of return, but not their ratio or elasticities of output to R&D.

the average number of years between origination and accrual of dynamic spillovers is 4.2 years. Using a discount rate of 6% for the social planner yields a discount factor $\tilde{\delta} = 0.783$. But because spillovers continue reverberating across indirect network connections, and their magnitude is halved at each step, I obtain $\delta = 1.28$.⁷⁰ I set $\delta = \delta'$, and thus do not consider possible financial frictions or myopia in firms that could drive a larger wedge between social and private returns.

I use my estimated coefficients from columns (4) and (7) in Table 1.9 together with the selected values for the parameters above.⁷¹ The estimated value of the MSR is 91% ($= 5.8 * (0.028 + 0.124 - 0.027 + 0.025 * 1.28)$) using the OLS coefficients, and 23.7% ($= 5.8 * (0.028 + 0.010 * 1.28)$) for the MPR. Using the 2SLS estimates, the values are 235% for the MSR and 80.8% for the MPR. This means that an increase of \$1 in a firm's R&D stock increases its own output by 81 cents and total output by \$2.35. In order to quantify the effect of an extra \$ of R&D investment, these values need to account for the durable effect this investment will have. If we continue assuming that R&D investment depreciates at a 15% rate and the discount rate is 6%, the above values must be multiplied by a factor of about 5.⁷² That is, an increase of \$1 in R&D investment leads to an increase in own output of \$4 (in discounted terms) and up to \$12 in terms of total output.

These estimates are informative, but are also dependent on the choice of values for the output-over-R&D stock ratio. Considering elasticities instead of the marginal returns reduces this issue. The elasticity of own output to own R&D stock is equal to $\phi_{rds} + \delta' \phi_{intra}$, whereas the elasticity of total output to a firm's R&D stock is $\frac{1}{N} (\phi_{rds} + \phi_{sic} + \phi_{cit} + \delta \phi_{dyn})$. Therefore, using the 2SLS estimates and assuming symmetric firms, increasing the R&D stock of a given firm by 1% leads to an increase in that firm's output of 0.14% and an increase in the total output of the 340 firms of $\frac{0.405}{340}$ %.

Therefore, for the sample of firms considered, the social returns to R&D are about three times as large as private returns. Interestingly, the inclusion of dynamic knowledge spillovers in the calculation of the rates of returns does not substantially affect the ratio between social and private returns to R&D. Even though dynamic spillovers increase the MSR to R&D, the MPR, and the resulting wedge between these two, if

⁷⁰ $\tilde{\delta} + \frac{\tilde{\delta}}{2} + \frac{\tilde{\delta}^2}{2^2} + \dots = \frac{\tilde{\delta}}{1-\tilde{\delta}/2}$

⁷¹For ϕ_{dyn} I use the coefficient from columns (4) and (5) in Table 1.5 instead. This is because dynamic spillovers for the social value should incorporate both within and between firm spillovers.

⁷²This factor is given by the discounted infinite sum of effects, with discount factor $\frac{1-0.15}{1+0.06}$.

anything they even slightly reduce the ratio between both returns. This is because the rate of internalization of within-firm dynamic spillovers slightly smaller than the ratio between private and social returns. As a result, even though dynamic spillovers are large and significant, including them in the returns calculations does largely not affect this ratio. This is comforting in that it reinforces the relevance of past estimations of the ratio between social and private returns to R&D.

In order to obtain estimates of the over- or under-provision of R&D in the decentralized equilibrium relative to the social optimum, I need to make additional assumptions. First, I assume that the marginal social and private returns defined above are appropriate to measure returns to R&D in terms of social and private surplus. The MSR will be a good proxy for social surplus returns to R&D if increases in social surplus (both profits and consumer surplus) are associated to increases in equilibrium market quantities. The MPR will be a good proxy for private returns to R&D in terms of profits if increases in profits are associated to increases in own output.⁷³ For a more detailed discussion, see Appendix A.3.3. Under these assumptions, estimates of the provision of R&D can be obtained combining the estimates of the social and private returns to R&D with the elasticity of R&D with respect to its cost. The first stage regressions in Table 1.3 show an elasticity of between -0.7 and -2.6, depending on whether one considers the federal component of R&D tax price or the state-level one. Previous estimates in the literature range. Hall and Van Reenen (2000) find a unitary elasticity of corporate R&D to its tax price using Compustat data, Chang (2014) finds an elasticity of state-level R&D of -2.8 to -3.8 in the US, and Bloom et al. (2002) report a long-term unitary elasticity of aggregate R&D using data from a panel of countries. Using a unitary price elasticity of R&D of -1, the ratio of returns is also equal to the ratio of the optimal social provision of R&D relative to the decentralized economy's provision.⁷⁴ Therefore, according to my estimates the under-provision of corporate R&D in the decentralized equilibrium is sizeable, with the social optimum level of corporate R&D being about three times as large as the decentralized levels.

⁷³This includes assuming that re-absorbed within-firm spillovers are indeed internalized, while spillovers between firms are not, through for example unobserved licensing agreements.

⁷⁴For a more detailed derivation of this result, see Appendix A.3.3.

1.8 Conclusion

In this paper, I have used a panel of publicly-listed US firms and their patenting activity to examine dynamic spillovers of corporate R&D. I have constructed measures of cumulative knowledge proximity between firm-year observations through a patent citation network, and used these measures to construct pools of R&D upon which firms are likely to build further. Causal effects have been estimated exploiting the persistence of the patent citation network, and variation in the federal and state tax treatment of R&D expenditures. I have found sizable dynamic spillover effects on firm output, with an elasticity at least as large as that of own R&D investment. The effect of dynamic spillovers on market value are comparatively even larger. I have also found that dynamic spillovers have sizable effects through the firms' innovation production function, and that they accrue primarily in complex product types, which are more reliant on cumulative innovation. Lastly, because the dynamic spillovers are partly internalized by the original innovator, their exclusion does not greatly influence the estimation of the ratio between social and private returns to R&D. I have found that social returns to R&D, incorporating knowledge spillovers (dynamic and static) as well as business stealing spillovers, are about three times as large as private returns. Using a unitary elasticity of R&D to its user cost, this results in a sizable underprovision of R&D in the decentralized equilibrium, with the social optimum level of corporate R&D being about three times as large as the decentralized levels.

I have also found that the relative importance of the knowledge spillovers estimated through the static and the dynamic measures is highly heterogeneous. It varies greatly across industrial sectors, and in particular across product and technology types, such as complex and discrete products, and chemical, electrical and mechanical technologies. First, this result is informative on the innovation processes involved within each product or technology category. It highlights that complex products and technologies do indeed seem to rely more on inter-temporal spillovers, on a cumulative innovation process. Moreover, it also underlines the importance of accounting for flexible dynamic structures in the analysis of knowledge spillovers in cross-industry studies, particularly with an eye for policy recommendations. For example, I show that certain industries, due to their complex inter-temporal knowledge flow structure, would not exhibit significant knowledge spillovers if only static measures were used, even though their dynamic spillover estimates are large.

The dynamic spillovers I have estimated are measured between patented inno-

vations, since the firm citation network uses patent data. This means that I have observed dynamic spillovers under the presence of patenting IP protection. Whether these spillovers would have been smaller or larger under an alternative IP system is outside of the scope of this paper. Studies such as Williams (2013), Galasso and Schankerman (2015), or Mezzanotti (2015) suggest that strong IP rights inhibit the rate and direction of innovative activity, and that lowering or ultimately scrapping patent protection could lead to stronger dynamic spillovers of R&D. However, on top of providing incentives for innovative activity, one of the aims of the patent system is also to further the diffusion of knowledge and innovation. In fact, patent exclusion rights are only awarded in exchange for the publication of innovation in a comprehensible and codified format. As such, the patent system could be furthering dynamic spillovers by providing for more important knowledge flows. Again, studies such as Furman and Stern (2011), Murray et al. (2016) or Furman et al. (2006) suggest that settings that are conducive to stronger knowledge flows will encourage subsequent innovation. The possible role of patents as codifiers and distributors of innovation deserves to be investigated further.

The estimates I have found in this paper can be informative of the shape of the aggregate innovation production function. In particular, I have found that the magnitude of dynamic spillovers from past innovations are about as large as that of own R&D and close to being unitary. Likewise, the magnitude of static knowledge spillovers suggests large complementarities of R&D activity. Following the framework in Jones and Williams (2000), let us consider an innovation production function of the type $k_{it} = \tilde{\delta}_t R_{it}$ at the individual firm level, with R = research effort. At an aggregate level, however, we have $\tilde{\delta}_t = \delta R_t^\lambda A_t^\phi$, as the productivity of R&D varies with the level of aggregate R&D effort and the existing stock of ideas. The results in this paper are consistent with such a function,⁷⁵ with parameters $\phi \approx 1$ and $\lambda \approx \phi_{cit} - \phi_{sic} = 1.9 > 1$.⁷⁶ This exercise is highly speculative, and only intended as an illustration of how the estimates found in this paper could be applied in macroeconomic models. Whether these are good estimators for the relevant parameters in an aggregate innovation

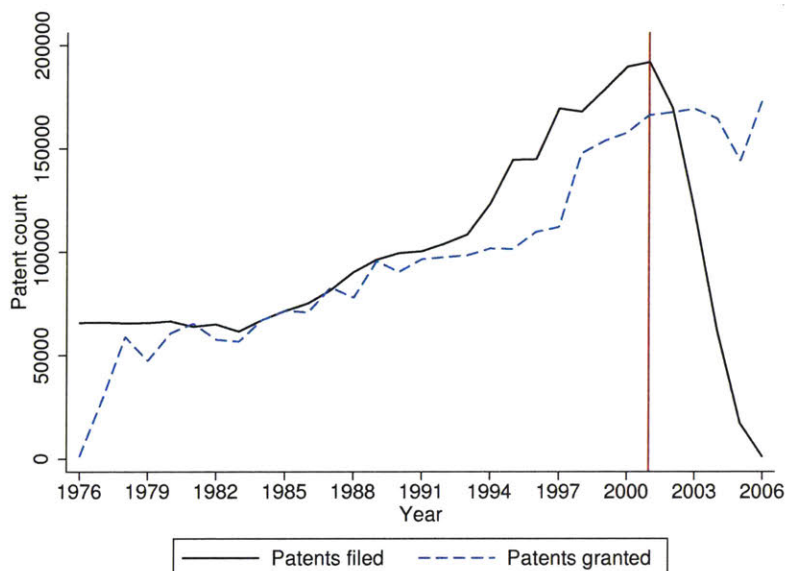
⁷⁵The elasticity of output with respect to own R&D in the empirical specifications is about 0.1 rather than unitary. Since the empirical specifications measure elasticity of output with respect to its arguments rather than the elasticity of knowledge creation, the coefficients need to be scaled by the unobserved elasticity of output with respect to firm-specific knowledge. In order to make the knowledge production proportional to the R&D effort, I multiply the empirical coefficients by 10.

⁷⁶ $\phi = 1$ would correspond to the case in Romer (1990). The estimate for λ would include both types of static spillovers, and implies large complementarities of R&D.

production function at the macro level is a difficult question. On the one hand, my estimation considers the influence of a selected amount of past and present R&D on output, chosen precisely because of its potential to have an influence. If one is considering instead how the entire existing body of knowledge affects the creation of new ideas, the influence might be smaller as this body includes possibly irrelevant technologies and ideas. On the other hand, the setting I consider is restricted to innovations for which patents have been filed and are thus likely to be protected by IP rights that may hinder or reduce subsequent innovation. In fact, Galasso and Schankerman (2015) find that reductions in patent IP protections lead to increased downstream innovation in complex technologies. One would thus expect that science and innovation carried out in more open settings with institutions devoted to the certification and dissemination of knowledge,⁷⁷ such as basic science in universities or public research institutes, may lead to larger cumulative spillovers.

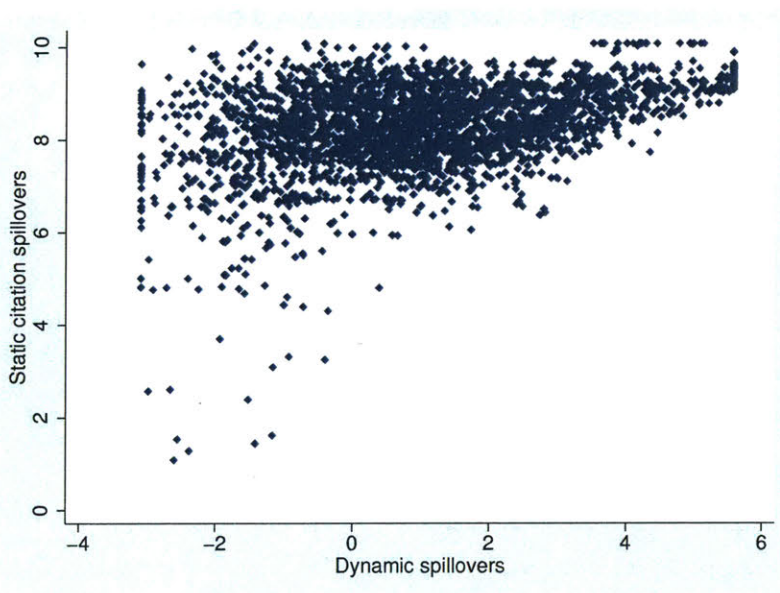
⁷⁷See Furman and Stern (2011).

Figure 1-1: Patent counts across time in the NBER data



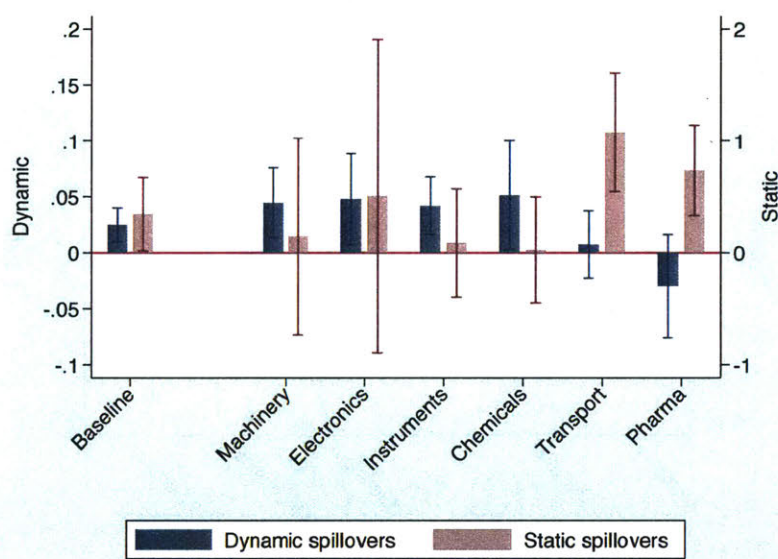
Notes: This figure shows the number of patents filed for (in a black continuous line) and granted (in a blue dashed line) in each year between 1976 and 2006. A vertical line marks 2001, when noticeable attrition starts.

Figure 1-2: Correlation between static and dynamic knowledge spillovers, DYNSPILL and CITSPILL



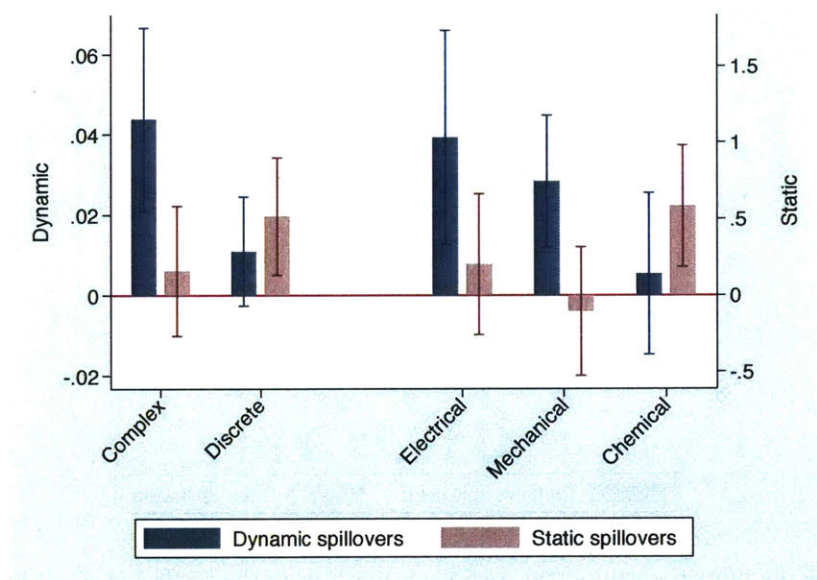
Notes: This figure plots the values of the dynamic knowledge spillovers $\ln(\text{SpillDyn})$ and the static knowledge spillovers $\ln(\text{SpillCit})$ for all firm-year observations with positive dynamic spillovers in my sample.

Figure 1-3: Heterogeneity across industrial sectors



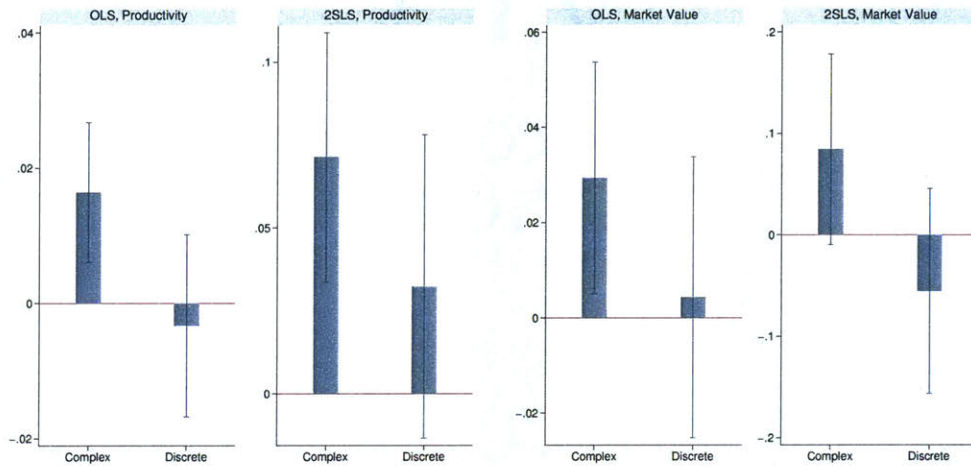
Notes: This figure plots the values and confidence intervals of the coefficient on dynamic and static knowledge spillovers, separated by the industrial sector of the receiving firm, for the productivity equation and OLS specification. The baseline coefficients corresponds to the baseline specification in Table 1.5; Machinery corresponds to 2-digit SIC code 35; Electronics corresponds to SIC code 36; Instruments corresponds to SIC code 38; Chemicals corresponds to SIC code 28, except drugs SIC code 283; Transport corresponds to SIC code 37; Drugs corresponds to 3-digit SIC code 283. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Figure 1-4: Heterogeneity across product and technology types



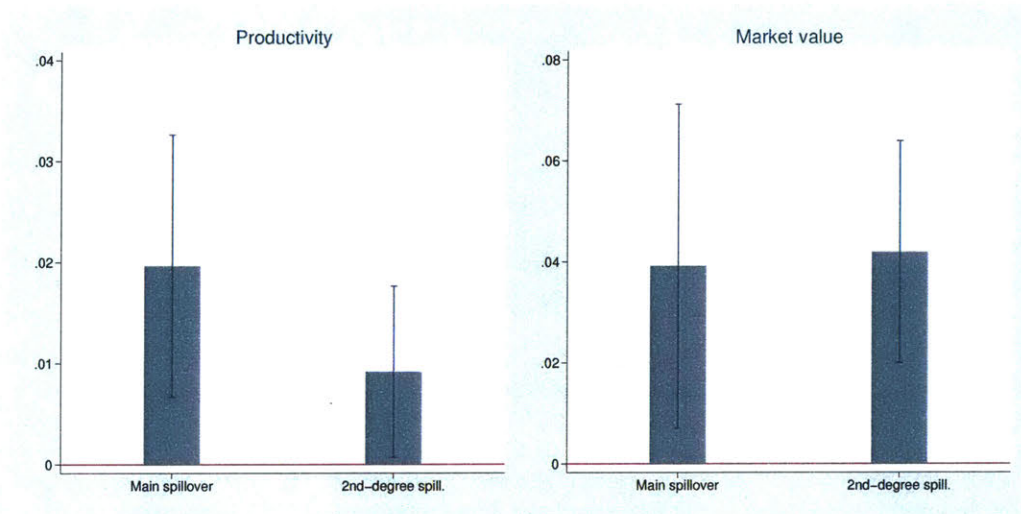
Notes: This figure plots the values and confidence intervals of the coefficient on dynamic and static knowledge spillovers, separated by the industrial sector of the receiving firm, for the productivity equation and OLS specification. The classification of SIC codes into complex and discrete, electrical, mechanical, and chemical are discussed in subsection 1.5. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Figure 1-5: Heterogeneity across technologies



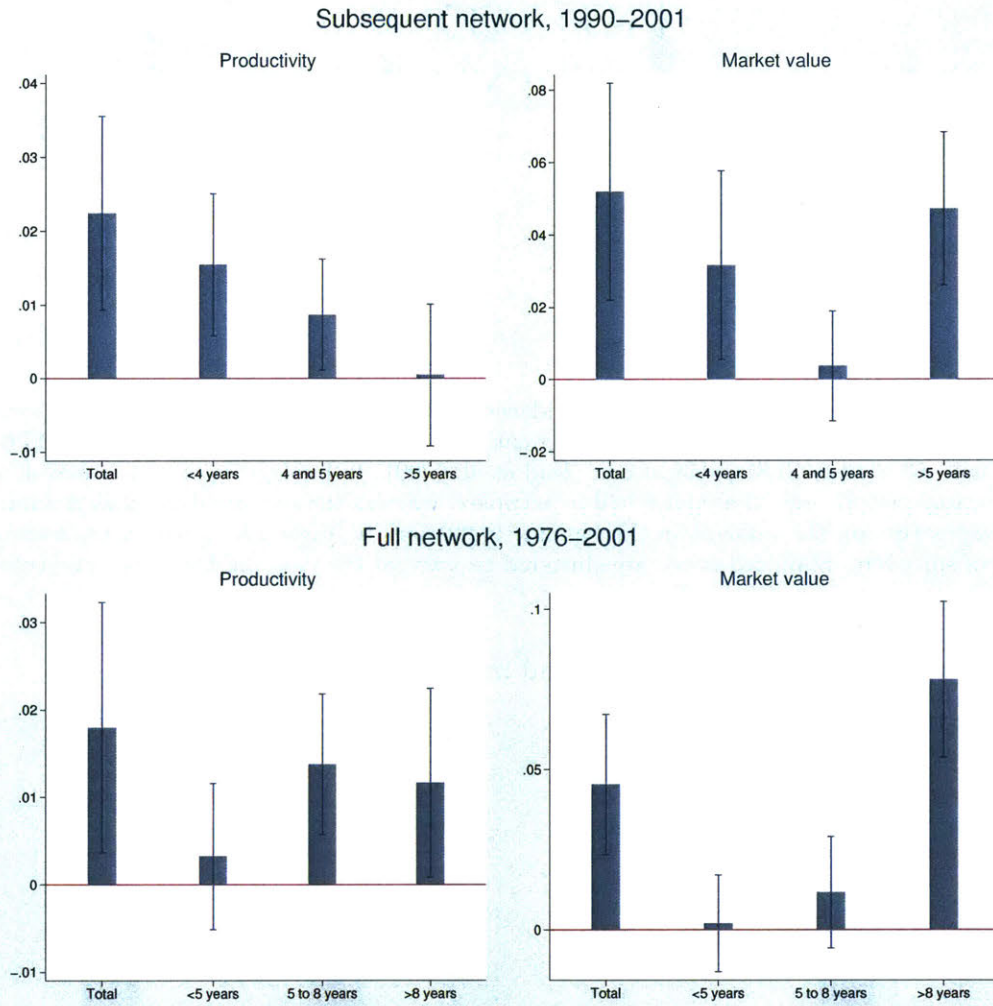
Notes: This figure plots the values and confidence intervals of the coefficient on dynamic spillovers, separated by the technology class of the originating patent, for both productivity and MTB equations, and OLS and 2SLS specifications. Hall et al. (2001) technology categories 2 and 4, as well as subcategories 32 and 33 are classified as complex, whereas the rest are defined as discrete. The regressions run are the same as in the baseline regressions in Table 1.5, albeit incorporating both types of spillovers. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Figure 1-6: Direct and indirect dynamic spillovers



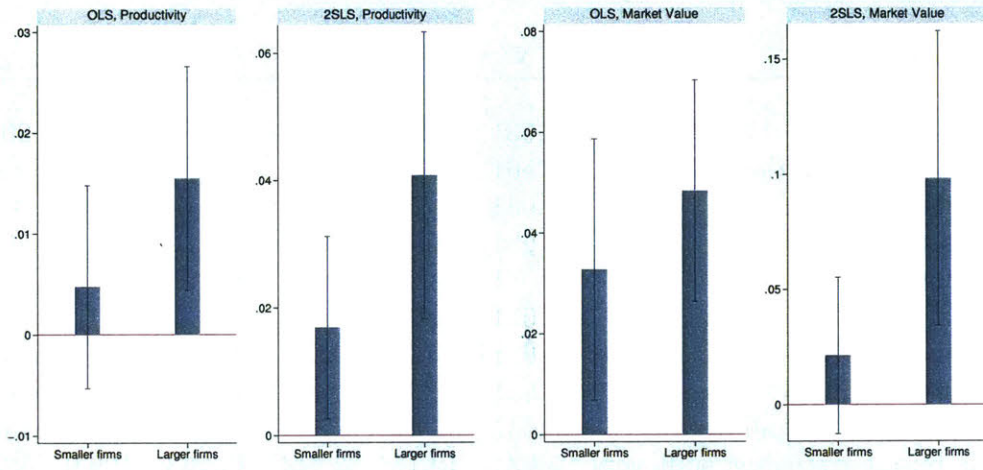
Notes: This figure plots the values and confidence intervals of the coefficients on dynamic spillovers, separated by whether they are constructed using direct or indirect second-degree patent citation network connections, for both productivity and MTB equations, and OLS specifications. The regressions run are the same as in the baseline regressions in Table 1.5, albeit incorporating both types of spillovers. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Figure 1-7: Depreciation of spillovers



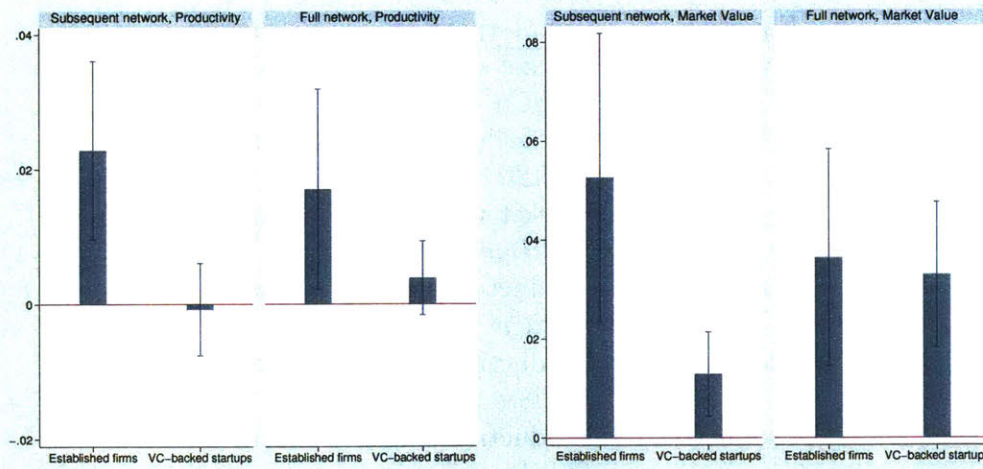
Notes: This figure plots the values and confidence intervals of the coefficient on dynamic spillovers, separated by the time bridged between the application year of the originating patent and that of the citing patent. It does so for both productivity and MTB equations, and using the restricted 1987-2001 network as well as the full 1976-2001 network. Spillovers are divided into three groups depending on how their time bridged falls in the corresponding distribution with respect to the median and 75th percentile within each network. The regressions run are the same as in the baseline regressions in Table 1.5, albeit incorporating all three types of spillovers. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Figure 1-8: Heterogeneity across firm sizes



Notes: This figure plots the values and confidence intervals of the coefficient on dynamic spillovers, separated by the size of the originating firm, for both productivity and MTB equations, and OLS and 2SLS specifications. Firms are classified as smaller if their average employee count between 1980 and 2001 lies in the bottom half of the distribution, and larger if it lies in the top half. The regressions run are the same as in the baseline regressions in Table 1.5, albeit incorporating both types of spillovers. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Figure 1-9: Spillovers from Venture Capital



Notes: This figure plots the values and confidence intervals of the coefficient on dynamic spillovers, separated by whether they originate in VC-backed startups or in established publicly-listed firms, for both productivity and MTB equations, and the full and restricted network. The regressions estimated are the same as in the baseline regressions in Table 1.5, albeit incorporating both types of spillovers. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Table 1.1: Descriptive statistics, variables

VARIABLES	N	mean	p50	sd	p5	p95
Sales	3,631	4,313	1,106	8,773	45.78	20,040
Market-to-book ratio	3,561	3.022	2.201	2.899	0.713	8.519
R&D flow	3,631	175.4	18.25	514.7	0	1,047
R&D stock	3,631	917.8	108.0	2,394	0	6,233
R&D intensity	3,631	0.232	0.150	0.242	0	0.698
No R&D	3,631	0.128	0	0.334	0	1
Dynamic spillovers	3,631	11.54	0.999	36.33	0	54.50
No dynamic spillovers	3,631	0.261	0	0.439	0	1
Static knowledge spillovers	3,631	5,411	4,127	4,629	163.6	14,314
Static business stealing spillovers	3,631	10,176	3,714	13,763	126.0	39,296
Fixed capital	3,631	1,514	257.0	3,560	8.200	7,197
Employment	3,631	20.61	6.740	39.51	0.352	90
Patent count	3,631	38.22	5	105.5	0	184.3

Notes: The statistics are taken over all non-missing observations between 1990 and 2001. All monetary values are measured in 2000 \$ in millions, employment measured in thousand employees. Missing observations are set to zero for R&D expenditures.

Table 1.2: Predictions for market value, productivity and R&D

Equation	Empirical Counterpart	Prediction
Market value	Market value with <i>SpillCit</i>	Positive
Market value	Market value with <i>SpillSic</i>	Negative
Market value	Market value with <i>SpillDyn</i>	Positive
Productivity	Productivity with <i>SpillCit</i>	Positive
Productivity	Productivity with <i>SpillSic</i>	Zero
Productivity	Productivity with <i>SpillDyn</i>	Positive

Notes: See theoretical framework in Appendix A.3.1 for the derivation of the predictions. The predictions are shown under the assumptions of positive technology spillovers and strategic complementarity between product market competitors' knowledge stock. *SpillCit* is the static technology-distance weighted sum of all other firms' R&D stocks. *SpillSic* is the static product market-distance weighted sum of all other firms' R&D stocks. *SpillDyn* is the dynamic technology-distance weighted sum of all other firms' R&D stocks.

Table 1.3: First stages, effect of tax instruments on R&D expenses

	(1)	(2)	(3)	(4)
ln(Federal component of tax)	-1.659** (0.677)	-1.874** (0.750)	-0.789*** (0.257)	-0.714*** (0.237)
ln(State component of tax)	-2.114*** (0.370)	-6.547 (4.142)	-3.571** (1.442)	-2.584* (1.460)
F-statistic	170.1409	172.5993	28.20957	29.29535
Observations	7372	7372	7372	7372
Year FE		✓	✓	✓
Firm FE			✓	✓
Industry-year FE				✓

Notes: Dependent variable is ln(R&D expenses) between 1980 and 2001. Regressions include a dummy for no R&D expenses and for no firm-specific federal tax instrument (about a third of observations). Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 1.4: First stages, pre-period network on spillover network

	Dynamic			Static
	(1)	(2)	(3)	(4)
Pre-period weight	0.068*** (0.013)	0.068*** (0.013)	0.067*** (0.013)	
Inter-firm weight				0.584*** (0.032)
F-statistic	356.4111	322.9695	325.4914	332.5909
Observations	2.8e+07	2.8e+07	2.8e+07	11615
Citing firm and year FE	✓	✓	✓	
Citing node FE		✓	✓	
Cited node FE			✓	
Citing and cited firm FE				✓

Notes: Dependent variable in columns (1) to (3) is weight of network edge in 1987-2001 citation network. Regressions include citing firm and citing year dummies in column (1), citing firm-times-year dummies in column (2), and a saturated set of citing firm-times-year and cited firm-times-year dummies in column (3). All regressions also include a dummy when both observations are zeros. The dependent variable in column (4) is the average weight of the 1987-2001 network between two firms, averaged across all years. Regressors include citing and cited firm dummies, and the average weight between the two firms in the pre-period 1976-1984 network. Standard errors in brackets are clustered two-way at the citing and cited firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 1.5: Baseline estimates

	Output				Market value		
	OLS				2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own R&D stock	0.348*** (0.049)	0.328*** (0.046)	0.036 (0.027)	0.029 (0.026)	0.092* (0.049)	0.024 (0.050)	0.144 (0.142)
Knowledge spill.	0.229** (0.101)	0.225** (0.100)	0.193** (0.076)	0.122** (0.062)	0.280** (0.130)	0.019 (0.162)	0.277** (0.133)
Business steal.	-0.034 (0.071)	-0.035 (0.065)	-0.040 (0.032)	-0.026 (0.030)	-0.087** (0.038)	0.116 (0.103)	-0.112* (0.067)
Dynamic spill.		0.052*** (0.012)	0.023*** (0.008)	0.025*** (0.009)	0.107** (0.047)	0.050*** (0.019)	0.171* (0.089)
First stage F-test					13.68		11.275
Observations	3631	3631	3631	3631	3631	3561	3561
Firm and year FE	✓	✓	✓	✓	✓	✓	✓
Controls			✓	✓	✓	✓	✓
Industry-year FE				✓			

Notes: Dependent variable is lead $\ln(\text{Sales})$ in columns (1) to (5), and $\ln(\text{Market-to-book ratio})$ in columns (6) to (7). All regressions include dummies for no R&D and no spillovers as well as a full set of firm and year FEs, output regressions also include a lead industry-specific price deflator. Controls include: industry-wide log-sales and lagged log-sales; log counts of patents filed; a sixth-order polynomial in $\ln(\text{R\&D stock})$, only the first term is shown for brevity; dummies for no patents filed. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively. Reported F-tests correspond to the Kleibergen and Paap (2006) rk Wald F statistic of weak instruments.

Table 1.6: Robustness checks

	3-digit SIC industry-year FEs		R&D polynomial		Only manufacturing		Patent count polynomial		Citation-weighted patent count		Alternative R&D intensity		Lower rate of depreciation		Full network
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)	OLS (7)	2SLS (8)	OLS (9)	2SLS (10)	OLS (11)	2SLS (12)	OLS (13)	2SLS (14)	OLS (15)
Panel A: Productivity															
Own R&D stock	0.037 (0.024)	0.156*** (0.053)	0.045* (0.023)	0.040* (0.023)	0.041* (0.023)	0.114** (0.045)	0.040* (0.023)	0.109** (0.044)	0.042* (0.023)	0.118*** (0.045)	0.042* (0.023)	0.101** (0.045)	0.036 (0.031)	0.102* (0.057)	0.047*** (0.011)
Dynamic spill.	0.023* (0.012)	0.072** (0.034)	0.023*** (0.008)	0.103** (0.046)	0.022*** (0.008)	0.106** (0.047)	0.023*** (0.008)	0.104** (0.047)	0.016** (0.007)	0.030* (0.017)	0.014** (0.006)	0.107** (0.048)	0.024*** (0.008)	0.103** (0.044)	0.020** (0.009)
First stage F-test		13.623		51.154		27.862		25.25		125.49		19.494		14.554	
Observations	3631	3631	3631	3631	3541	3541	3631	3631	3631	3631	3631	3631	3631	3631	9011
Panel B: Market Value															
ln(R&D stock)	0.072 (0.056)	0.130 (0.121)	0.019 (0.062)	0.009 (0.064)	0.013 (0.054)	0.042 (0.135)	0.008 (0.053)	0.035 (0.132)	0.016 (0.054)	0.059 (0.129)	0.017 (0.054)	0.028 (0.136)	0.022 (0.060)	0.015 (0.170)	-0.119*** (0.032)
ln(Dynamic spill.)	0.058** (0.025)	0.223** (0.094)	0.050*** (0.018)	0.200** (0.095)	0.051*** (0.019)	0.178* (0.100)	0.052*** (0.018)	0.165 (0.101)	0.033** (0.017)	0.045 (0.043)	0.037** (0.017)	0.180* (0.102)	0.053*** (0.019)	0.182* (0.094)	0.041*** (0.013)
First stage F-test		12.094		39.754		22.685		20.55		112.46		17.062		11.767	
Observations	3561	3561	3561	3561	3471	3471	3561	3561	3561	3561	3561	3561	3561	3561	8814
Firm and year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Industry-year FE	✓	✓													

Notes: In panel A, dependent variable is $\ln(\text{Sales})$ and regressions include the log of an industry-specific price deflator; in panel B dependent variable is $\ln(\text{MTB})$. All regressions include industry-wide log-sales and lagged log-sales; log counts of patents filed; dummies for no R&D, for no dynamic spillover and for no patents filed, as well as a full set of firm and year FEs. Columns (1) and (2) include narrow 3-digit SIC-code industry-times-year FEs, columns (3) and (4) include a sixth-order polynomial in $\ln(\text{R\&D stock})$ (only the first term is showed for brevity), columns (5) and (6) restrict the sample to manufacturing firms, SIC codes 2000-3999, columns (7) and (8) include a fourth-order polynomial in log patent counts, columns (9) and (10) include log of citation-weighted patent counts, columns (11) and (12) use R&D/Sales instead over R&D/Assets as a measure of research intensity, columns (13) and (14) use a lower depreciation rate of 10% for all R&D stocks, and column (15) estimates the OLS estimation on a 1980-2001 sample using the whole 1976-2001 network. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively. Reported F-tests correspond to the Kleibergen and Paap (2006) rk Wald F statistic of weak instruments.

Table 1.7: Dynamic spillovers on R&D and citation-weighted patenting

	Citation-weighted patents			R&D activity	
	(1)	(2)	Neg. Bin.	(4)	(5)
			(3)		
Own R&D stock	-0.034 (0.092)	-0.040 (0.088)	0.006 (0.041)		
Dynamic spill.	0.519*** (0.084)	0.420*** (0.080)	0.344*** (0.035)	0.037*** (0.012)	0.024** (0.010)
R&D/Sales _{t-1}					0.538*** (0.047)
Patents _{t-1}		0.122**	0.141***		
Observations	3289	3289	3289	2872	2863
Firm and year FE	✓	✓	✓	✓	✓

Notes: Dependent variable is lead ln(Citation-weighted patent count) in columns (1) and (2), lead citation-weighted patent count in column (3), and lead ln(R&D expenditures) in columns (4) and (5). Column (3) estimates a conditional negative binomial model with panel fixed effects, and the other columns estimate OLS regressions. Regressions include controls for industry-wide log-sales and lagged log-sales; dummies for no dynamic spillover, as well as a full set of firm and year FEs. Columns (1) to (3) also include log R&D stock and a dummy for no R&D stock; columns (2) and (3) also include logs of citation-weighted patent counts; columns (4) and (5) include the log of sales in order to normalize by firm size, and column (5) includes current ln(R&D expenditures). Standard errors in brackets are bootstrapped in column (3), and clustered two-way at the year and firm level in all other columns. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 1.8: Supply chain and geographic spillovers

	Productivity				Market value			
	Base- line	Supply chain	Geographic spillovers		Base- line	Supply chain	Geographic spillovers	
	OLS (1)	OLS (2)	OLS (3)	2SLS (4)	OLS (5)	OLS (6)	OLS (7)	2SLS (8)
Own R&D stock	0.042* (0.023)	0.040* (0.022)	0.041* (0.023)	0.116*** (0.044)	0.016 (0.054)	0.014 (0.054)	0.013 (0.053)	0.048 (0.129)
Dynamic spill.	0.022*** (0.008)	0.023*** (0.008)	0.018** (0.008)	0.049** (0.020)	0.052*** (0.018)	0.053*** (0.018)	0.037** (0.017)	0.111* (0.058)
ln(Supply chain spill.)		-0.025** (0.010)				-0.093*** (0.033)		
ln(Geographic spill.)			0.008 (0.008)	-0.001 (0.020)			0.026* (0.015)	0.008 (0.033)
First stage F-test				20.133				19.94
Observations	3631	3631	3631	3631	3561	3561	3561	3561
Firm and year FE	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Dependent variable is lead ln(Sales) in columns (1) to (3), and ln(Market-to-book ratio) in columns (3) and (4). Regressions include controls for industry-wide log-sales and lagged log-sales; log counts of patents filed; dummies for no R&D, for no dynamic spillover and for no patents filed, as well as a full set of firm and year FEs. Columns (1) to (4) also include the log of an industry-specific price deflator. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively. Reported F-tests correspond to the Kleibergen and Paap (2006) rk Wald F statistic of weak instruments.

Table 1.9: Own vs others' spillovers, output

	OLS				2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own R&D stock	0.045*	0.043*	0.034	0.028	0.135***	0.119**	0.083*
	(0.023)	(0.023)	(0.027)	(0.026)	(0.044)	(0.047)	(0.049)
Intra-firm dynamic		0.009	0.010	0.010		0.042*	0.044*
		(0.006)	(0.006)	(0.006)		(0.022)	(0.023)
Inter-firm dynamic	0.027***	0.029***	0.029***	0.029***	0.091***	0.081**	0.085**
	(0.007)	(0.008)	(0.007)	(0.008)	(0.034)	(0.033)	(0.034)
Knowledge spill.			0.194**	0.124**			0.279**
			(0.077)	(0.063)			(0.129)
Business steal.			-0.040	-0.027			-0.094**
			(0.032)	(0.029)			(0.038)
First stage F-test					12.704	13.25	8.7923
Observations	3631	3631	3631	3631	3631	3631	3631
Firm and year FE	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Industry-year FE				✓			

Notes: Dependent variable is lead ln(Sales). All regressions include dummies for no R&D, no spillovers and no patents filed; industry-wide log-sales and lagged log-sales; log counts of patents filed; a sixth-order polynomial in ln(R&D stock), only the first term is shown for brevity; a lead industry-specific price deflator, as well as a full set of firm and year FEs. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively. Reported F-tests correspond to the Kleibergen and Paap (2006) rk Wald F statistic of weak instruments.

Chapter 2

Patent duration and cumulative innovation: Evidence from a quasi-natural experiment (joint with Jean-Noel Barrot)

Abstract

Cumulative innovation is at the core of economic growth, but the impact of patent policy on it is not well understood. This paper investigates whether patent term duration affects the rate and direction of follow-on innovation. We use a quasi-natural experiment that lengthened the term on existing patents in the US, and leverage a kink in the extension formula to identify the effects of patent term increases. We find no statistically significant impact of patent extensions on subsequent innovation, neither locally around the kink using a sharp "Regression Kink Design" nor on average on the population of treated patents. We further analyze whether the null average effect could be masking important heterogeneous effects, and find no such evidence.

2.1 Introduction

Innovation is at the core of long-term economic growth. However, the non-rivalry and non-excludability of ideas (Nelson, 1959; Arrow, 1962) can lead to lessened incentives for innovation production. Recognizing this, intellectual property (IP) rights have been established by most governments in order to provide incentives to innovate.¹

¹For example, the U.S. constitution explicitly links IP rights to incentives to innovate in its article I, section 8: *The Congress shall have power... To promote the progress of science and useful*

Most evaluations of IP rights have focused on the trade-off between *ex-ante* incentives for innovation and the *ex-post* inefficiencies associated with the increased market power derived from the IP exclusionary rights.²

A more recent empirical literature adds a layer of complexity to the analysis by recognizing that innovation is often cumulative in nature,³ and that IP rights on existing technologies can also have implications for the intensity and direction of follow-on innovation. It largely finds that IP rights reduce technology use, as well as subsequent research and innovation, and complements a more extensive previous theoretical literature that provided no clear predictions on the impact of IP on cumulative innovation (Kitch, 1977; Scotchmer, 1991; Green and Scotchmer, 1995). However, this recent literature has mostly concentrated on the extensive side. That is, it has compared follow-on outcomes under IP protection to outcomes without it, or to outcomes when IP protection drops markedly. This research is useful to determine whether particular technologies should be awarded IP rights or not. However, it does not necessarily help determine policy parameters of interest on the intensive side, such as what is the optimal duration of IP protection?

The impact of IP duration on follow-on innovation could depend both on the impact of the actual lapse in protection at the end of the term, but also on the strategic response of agents before the term lapses. In terms of the direct impact, if patent protection reduces follow-on innovation⁴ due to bargaining or transaction costs, increases in patent terms can depress subsequent innovation for longer. Moreover, longer patent terms result in the release into the public domain of older and more obsolete technologies, which may therefore be less valuable to build upon when the patent lapses.⁵ Furthermore, the strategic response of agents before patent lapse can strengthen this impact. For example, Li et al. (2016) study the UK Copyright Act of 1814 and find that additional years of copyright protection increased prices by improving publishers' ability to practice intertemporal price discrimination. Similarly,

arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries.

²See for example Nordhaus (1969); Klemperer (1990); Gilbert and Shapiro (1990); Budish et al. (2015).

³In fact, the endogenous growth literature in particular depicts knowledge spillovers as dynamic, accruing when past ideas become the new foundation on which to build further innovation. See e.g. Romer (1990), Aghion and Howitt (1992), and Jones (1995).

⁴As in Murray et al. (2016); Williams (2013); Galasso and Schankerman (2015); Biasi and Moser (2017); Nagaraj (2016).

⁵Mehta et al. (2010) study age profiles of patent citations and find that citations peak at around two to three years after grant date, and then decrease with age.

longer patent protection could slow technology diffusion, as well as hinder follow-on innovation in the presence of transaction costs by strengthening the bargaining position of the upstream innovator. Nonetheless, empirical research on the effect of patent protection term duration on subsequent innovation is lacking.

The contribution of this paper is to provide, as far as we know, the first formal empirical analysis of whether longer patent terms hinder follow-on innovation. We leverage a quasi-natural experiment in 1995 with the passing of the Trade-Related Aspects of Intellectual Property Rights (TRIPS). This reform package, negotiated during the Uruguay Round of trade agreements that led to the creation of the World Trade Organization, lengthened patent terms for existing outstanding U.S. patents. Specifically, it moved patent terms from a maximum length of 17 years after patent *grant* to 20 years after patent *application* for new patents, and included a clause of retro-activity that also awarded existing patents a lengthened term equal to the more generous of the two regimes. As a result, most outstanding patents received a term boost that depended negatively on their processing time, that is on the difference between application and grant date. Meanwhile, outstanding patents with a processing time longer than three years saw no change in their term. In our empirical analysis, we identify the impact of patent term extensions on follow-on innovation by taking advantage of the kink in the TRIPS-induced term extension function at a processing time of 3 years. Under the assumption that the impact of processing time on follow-on innovation is the same on both sides of the kink, we can identify the impact of the TRIPS treatment by comparing the estimates on either side.

In section 2.2, we describe the research setting in more detail, including relevant institutional information about the U.S. patent law and the implementation of TRIPS. The TRIPS reform has been exploited as a quasi-natural experiment by a number of recent papers. Among the most relevant for our analysis, Abrams (2009) analyzes the impact of longer expected patent term on innovation after the implementation of TRIPS, and finds differential effects by technology class. More recently, Lemus and Marshall (2017) as well as Sukhatme and Cramer (2014) find that TRIPS increased incentives for assignees to shorten patent processing times. Hshieh (2017) studies the response by firms to windfall profits due to patent term extensions because of TRIPS, and finds no evidence of increased internal R&D expenditures. We contribute to this literature by analyzing how TRIPS impacted incentives for follow-on innovation on existing outstanding patents.

We present our empirical strategy and results in section 2.3. We follow an established empirical literature in using in-citations by later patents to the focal patent of analysis to trace knowledge spillovers on follow-on innovation.⁶ First, we use panel and cross-sectional specifications to estimate the impact of processing time on in-citation counts on both sides of the TRIPS-induced 3-year kink. These specifications allow us to identify the *average treatment effect on the treated* under the assumption that the underlying impact of processing time on follow-on innovations is the same on either side of the kink. We analyze a number of empirical specifications, and study different outcome variables in terms of the span of time considered for follow-on innovation, as well as depending on whether the follow-on innovation is carried out by the original patent assignee or not. Across all specifications, we consistently find no evidence of an impact of longer patent terms on follow-on innovation.

We then carry out a more local analysis around the 3-year kink, using regression kink design specifications. On the one hand, the underlying identifying assumption of equal impact of the processing time running variable on follow-on innovation is more likely to be satisfied in a regression kink design, as it focuses on differences within a localized bandwidth around the cutoff. On the other hand, the regression kink estimates correspond to a local *treatment effect on the treated conditional on processing time being equal to 3 years*, which may or may not present external validity further away from the kink. We first show that there is no discontinuity in patent covariates around the cutoff in either levels or slopes, nor any evidence of bunching. We then estimate regression kink specifications and consistently find no significant impact of longer patent terms on follow-on innovation. Although our regression kink estimates are noisier, they lend weight to the null average result estimated over the population of treated patents.

In section 2.4, we analyze whether the zero average treatment effects mask heterogeneity of treatment impacts across patent characteristics. We differentially estimate impacts of TRIPS on follow-on innovation depending on the technology class of the focal patent, and show null results across each of Hall et al. (2001) technology classes, as well as for both complex and discrete product types. We then separately estimate the impact by focal patent grant year and filing year, and show homogeneously null results. We also study heterogeneity across focal patent assignee sizes, and across the distribution of patent quality or value, as proxied by pre-TRIPS citation counts.

⁶Although using this proxy is not perfect, it is "*the only feasible approach if one wants to study the impact of patent rights across diverse technology fields*" (Galasso and Schankerman, 2015).

Once again, we find no evidence of meaningful heterogeneity being masked by the average null treatment effect.

Finally, in section 2.5 we discuss the implications of our empirical results and their place in the extant empirical literature. Murray et al. (2016) and Williams (2013) find that non-patent IP protection on genetically-modified mice and on the human genome respectively reduces subsequent research and innovation. On the other hand, Sampat and Williams (2015) find no effect on follow-on innovation of patent protection on human genes. Meanwhile, Galasso and Schankerman (2015) do find that patent invalidations by courts increases follow-on innovation for a broad range of technologies. Moser and Voena (2012) find that compulsory licensing of German patents during WWI led to increased follow-on innovation in the US. Finally, Biasi and Moser (2017) and Nagaraj (2016) find that copyright protection decreases re-use and may hinder creation of follow-on innovation. In our panel specifications, we can reject reductions in citation counts of 0.6% during the first 11 years post-reform per additional year of patent term protection. Meanwhile, our cross-sectional estimates are noisier, and allow us to reject reductions of 6% in post-TRIPS citations. We discuss our contribution to the literature, and conclude in section 2.6.

2.2 Research setting

2.2.1 Uruguay Round

The Uruguay Round was a set of multilateral negotiations spanning from 1986 to 1994, conducted within the framework of the General Agreement on Tariffs and Trade (GATT). It involved 123 countries and led to the creation of the World Trade Organization. The main goals of the Round were to expand GATT rules to new areas such as agriculture, textiles, and services; to reduce restrictions on foreign direct investment; and to set international minimum standards for IP rights.

The Uruguay Round culminated with the signing of the Marrakesh Agreement in April 1994 by 124 countries, which included an agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) to unify minimum IP standards. The United States implemented the Marrakesh Agreement into U.S. law through the Uruguay Round Agreements Act (URAA). The URAA bill was introduced in the House on September 27, 1994, and was passed on November 29 of that same year. It then passed the Senate on December 1, and was signed into law by President Clinton

on December 8, 1994. Although the bill was submitted under special procedures that prohibited either chamber from modifying it, its passage was uncertain until late November. The bill faced strong opposition in the Republican-controlled Senate, where the soon-to-be Majority Leader Robert Dole demanded a capital-gains tax cut in exchange for support. It also faced strong protectionist forces, and formidable opponents in Ralph Nader and Ross Perot.

Most of the uncertainty regarding the bill was lifted on November 24, 1994, after a meeting between President Clinton and Senator Dole. Dole agreed to drop his demand in exchange for the White House to back legislation contemplating future U.S. withdrawal of the WTO if its membership was deemed to hurt the United States (see Fletcher (1994)). Given extended Democrat support, and with Dole poised to become the Senate Majority leader in the coming January, the bill's approval was largely anticipated after this announcement. The contents of the TRIPS reform were therefore known throughout 1994, although the uncertainty on whether the reform would be implemented or not was not lifted until the end of November.

2.2.2 TRIPS reform

The TRIPS reform was implemented in the U.S with the signing of the URAA. TRIPS affected a number of aspects of IP policy, not all of them related to patent policy, but the most salient change for patents was to make their protection extend for a minimum of twenty years. U.S. patent law at the time provided for 17 years of protection after granting, which TRIPS reformed to 20 years of protection after filing. This change in patent term duration was the largest in the U.S. since 1861, and was implemented in June 8, 1995. It affected outstanding and incoming patents in the following way:

- For all patents filed on or after June 8, 1995, the patent term becomes 20 years from the filing date.⁷
- For all patents filed before June 8, 1995, but still outstanding on that date,⁸ the patent term becomes the longer of the two following options: either 17 years from the grant date of the patent, or 20 years from the earliest filing date.

⁷In fact, from the earliest U.S. or international filing date to which priority is claimed.

⁸That is, patents granted after June 7, 1978 that were not expired.

- All patents granted on or before June 7, 1978 or other patents that were expired by June 8, 1995 were not affected by the TRIPS reform.

That is, the TRIPS reform not only affected patents filed for after its implementation in June 1995, but also increased the patent terms of outstanding patents that had been previously applied for and possibly already granted. For the latter, the term extension depended on the processing time, or time between the earliest filing date and the subsequent grant date. For outstanding patents, the term extension was equal to 3 years minus the processing time, with no reduction in term, i.e. a floor at zero if processing time was greater than three years.

This patent term reform significantly affected the remaining maturity of outstanding patents. In Figure 2-1 we plot the histogram of extensions awarded to outstanding patents due to TRIPS. The average extension granted is almost 14 months, with a median extension of just under 15 months, and with a share of patents that receive no term extension of just under 10%. Moreover, there is suggestive evidence that patent assignees believed this patent term increase to be valuable. In Figure 2-2 we plot the number of weekly patent applications⁹ and patent grants between 1993 and 1997. It shows a clear spike in applications before June 8, 1995.¹⁰ The figure shows the equilibrium effect of two countervailing forces: filing before the TRIPS deadline results in the option value of receiving the greater of the two patent term options, and this option value actually increases with the number of patent filings, as more patent filings will likely lengthen the backlog at the U.S Patent and Trademark Office (USPTO) and slow down the process. On the other hand, applying for a patent earlier could lead to a looser application file which might face more difficulties in being approved or result in less comprehensive protection. It also decreases the option value of waiting if the value of the patentable technology is uncertain.

The TRIPS reform also included other elements for the patent system, although with less salience and significance. The main changes, other than the patent term modification, involved allowing for foreign activity to prove a date of invention and the creation of provisional patent applications.¹¹ However, these additional reforms affected new patent applications rather than outstanding patents previously granted.

⁹This count only includes applications to patents that are ultimately granted by December 16, 2016.

¹⁰For more evidence of this spike, see Abrams (2009), Sukhatme and Cramer (2014), Hshieh (2017).

¹¹See Van Horn (1995) and Sukhatme and Cramer (2014).

In order to avoid conflating concerns, we focus our empirical analysis on patents that were granted years before the implementation of TRIPS. Since the Uruguay Round would plausibly have led to increased trade and cross-country patenting activity that could affect citation patterns, we control directly for whether the original assignee is a foreign or US entity in our empirical specifications.

2.2.3 Data and variable construction

We use utility patent data from the USPTO PatentsView platform, current up to December 16, 2016. The data includes information on granted patents since 1976 and published patent applications since 2001. Variables include filing and grant date, number of claims, technology class and subclass, citation patterns, and patent assignee information. We combine this data with administrative information of maintenance fee payments from the USPTO bulk downloads. Utility patents issued on or after December 12, 1981 have to pay renewal fees in the sixth months prior to their 4th, 8th, and 12th year of protection in order to maintain validity. Failure to pay these fees in a timely manner results in patent expiry. We use the maintenance fee payment profile of each patent to more closely ascertain the running validity of each patent.

We keep patents granted in the 1980s for our empirical analysis. This allows us to focus on patents that are still outstanding (that is, as long as they paid their corresponding fees) by the time of the TRIPS reform in 1995 and that were not set to lapse immediately after. We restrict the analysis to patents granted before 1990 in order to observe enough of a pre-treatment citation profile.¹² Moreover, notice that our analysis focuses on patents that were filed for and granted years before the TRIPS reform took place. We can safely assume that no TRIPS-induced consideration would affect strategic patenting decisions in our sample.

For each patent i in the sample, we construct the following variables that are determined at grant date.

- treat_i is the TRIPS-induced extension of patent term, in years. This treatment intensity is defined as $\text{Max}(0, 3 - \text{Grant date} + \text{Filing date})$. That is, it is equal to three calendar years minus the patent's processing time, with a floor at zero.

¹²Mehta et al. (2010) find that the usual citation profile of patents starts at grant date in a sample of patents granted up to 2002. Most of their dataset corresponds to patents granted before 2001, when the USPTO only published patent grants and not applications. This is consistent with technology diffusion largely starting at publication date.

- $proctime_i$ is patent i 's processing time, equal the Grant date - Filing date. The underlying running variable that defines the treatment intensity is $procrun_i$ and is equal to the patent's processing time minus three calendar years; the treatment intensity is then $treat_i = \text{Max}(0, procrun_i)$.
- We use the categorization of patent technology classes by Hall et al. (2001) to define 6 HJT technology classes and 36 subclasses.
- We measure the HHI (from 0 to 1) of HJT technology subclasses among patent i 's out-citations. This index, defined by Hall et al. (2001) as *originality* is meant to capture how diverse the set of patents cited by the focal patent are.¹³
- Other control variables include counts of out-citations by the focal patent and counts of the number of patent claims. We also include dummies for the number of maintenance fees paid for, between 0 and 3. We also include dummies on whether the country of origin of the original patent assignee is the US or is foreign, in order to control for the internationalization associated with the Uruguay Round.

As our outcome variable, we follow a well-established literature using patent citations to follow knowledge spillovers resulting in follow-on innovation (Galasso and Schankerman, 2015). We distinguish between in-citations by subsequent patents that are filed by other assignees from those filed by the original innovator.¹⁴ In the first part of our empirical analysis, we aggregate a focal patent's citations at the year of application of the subsequent patent, and use the resulting panel data variation. In the cross-sectional analysis, we aggregate the citation data to pre-TRIPS counts as control variables¹⁵ and post-reform counts as outcome variables. Finally, we drop the 1st and 99th percentile in terms of the processing time running variable in order to avoid outliers, and we restrict the patent sample to patents that paid the relevant maintenance fees and were still outstanding by the TRIPS reform. The latter restriction allows us to maintain comparable patents in the treated and untreated subsamples, since less valuable patents would be more likely to lapse prior to the reform and therefore would be disproportionately represented among untreated patents.

¹³In the case of a patent with zero out-citations, we set its HHI to 1.

¹⁴In order to distinguish both, we rely on the PatentsView assignee disambiguation. See their website for a discussion of the algorithm used.

¹⁵We include both all citations before 1994, as well as a more restricted count from 1990 to 1993 as controls.

We address potential selection into the sample by also including lapsed patents and by analyzing intent-to-treat in robustness checks, and find that our results are robust to these variations. Descriptive statistics in our sample of patents are shown in Table 2.1.

2.2.4 Empirical strategy

We exploit the TRIPS reform in 1995 as a natural experiment that increased the patent term length of enforceable patents. The extent of patent term increase was highly heterogeneous, as shown in Figure 2-1, ranging from no increase to a lengthening of the term by close to three years. By focusing on the effect of TRIPS on patents granted well before the implementation of the reform, we avoid concerns of endogenous or strategic selection into the treatment. Abrams (2009) finds an heterogeneous impact across technologies of TRIPS on patent applications after its implementation, depending on their average expected patent term increases. In our analysis, we refrain from such considerations and include flexible fixed effects to absorb technology-wide effects in some specifications. Instead, we are interested in how, within each technology class,¹⁶ the duration of effective patent protection impacts follow-on innovation. The contents of the TRIPS reform were known throughout 1994, although the uncertainty on whether the reform would be implemented or not was not lifted until the end of November. As a result, we expect any impact of TRIPS on follow-on innovation to fully start taking place in 1995, although it is possible that some impacts could start as early as late 1994.

Although our empirical setting reduces concerns of strategic selection into treatment, we still worry that the treatment intensity is not randomly assigned. In fact, patent term extensions are a deterministic function of the patent processing time, defined as the span of time between the first application date and the eventual patent grant date. This processing time has an element of randomness, as it will be affected by backlog at the USPTO, and can depend on the efficiency of the individual patent officer quasi-randomly assigned to the application.¹⁷ However, processing time will also likely depend on the complexity of the application. More complex patent

¹⁶As shown later on in the empirical analysis in section 2.3, we focus on even more restrictive within-group variation. Our groups are generally defined as narrow technology subclasses times application year, in order to keep patents highly comparable within groups.

¹⁷For more information about the patent grant process, see Sukhatme and Cramer (2014). For more information about the quasi-random allocation of patents to examiners, see Lemley and Sampat (2012).

applications, e.g. with more claims, might take longer to evaluate. Likewise, more obscure applications with uncertain novelty could also lead to longer processing times.¹⁸ Finally, processing time can also be directly affected by the filing party requesting extensions to their allotted time for responses.¹⁹ Again, since the patents we focus on were granted well before the terms of the TRIPS agreement were public, concerns of strategic processing time modifications should be absent. Nonetheless, in Table 2.2 we analyze differences in patent observables for our sample of patents between untreated patents with a processing time slightly longer than 3 years (between 3 years and 3 years and 50 days) in column (1), compared to treated patents with a processing time of around 2 years in column (2). Patents with longer processing times tend to have more out-citations, more claims, and tend to remain valid for longer by paying more maintenance fees. They also receive significantly more citations prior to the TRIPS reform. Even though we can control explicitly for these variables in our empirical specifications, processing time is likely also correlated with some unobservable measure of patent quality which may skew our results.

We account for selection into treatment by taking advantage of the shape of the treatment intensity function. As a reminder, the effect of the TRIPS reform on outstanding patents was to change their duration from 17 years after grant date to the greatest of the former and 20 years after filing date. As a result, if we define an underlying running variable as 3 years minus the processing time, the extent of patent term increase due to TRIPS is equal to that underlying running variable as long as the variable is positive. If the running variable is negative (for processing times above 3 years), the treatment intensity is zero. Figure 2-3 plots the relationship between the TRIPS treatment intensity as a function of the processing time. We see that the function is flat at zero for values of processing time above 3 years (negative values of the running variable), and has a negative slope of -1 for values below three years (positive slope of 1 with respect to the running variable for positive values).

In order to identify the TRIPS treatment effect, we separately estimate the impact of an extra processing day on citations on both sides of the kink at zero. For negative

¹⁸In general, patent examiners will issue successive rounds of non-final rejections that can be responded to with arguments and/or amendments to the patent claims. This back-and-forth can last for many rounds (Sukhatme and Cramer, 2014). In fact, even faced with "final" rejections, applicants can still request extensions to make their case, or register an appeal.

¹⁹In particular, Sukhatme and Cramer (2014) and Lemus and Marshall (2017) show that the implementation of TRIPS does change the incentives to the filing party to hasten or slow down the patent processing to take advantage of new patent term rules after the implementation of TRIPS.

values of the running variable (that is, for processing times longer than 3 years), this will incorporate effects due to unobservable patent quality differences.²⁰ For positive values, the effect will incorporate both the direct effect of processing time on follow-on innovation (as in the previous case of negative values) as well as the effect of patent term extension. If we assume that the impact of processing time absent treatment is the same on both sides of the kink, any differential effect we would find can be attributed to the TRIPS term extension. In the empirical analysis in section 2.3, we follow this empirical strategy to identify the impact of patent term extension by including both a treatment variable and a processing time variable.²¹ This allows to separate the impact of patent term extension from selection-into-treatment concerns.

The identifying assumption in our analysis is that the impact of processing time on follow-on innovation, after controlling for observable covariates, is the same on both sides of the kink. We address a number of concerns about this homogeneity assumption. First, we suspect that the assumption is more likely to hold close to the kink rather than further away, because of possible nonlinearities or because patents with very different processing times are not comparable. Suggestive evidence for this can be found in Table 2.2, which compares covariates for two groups of patents that are quite diverse in term of their processing time. Most t-tests of the covariates show a statistically significant difference between the two groups. We address this first by considering sequentially narrower cut-off thresholds around the kink on which we conduct the analysis. Moreover, we also estimate regression kink discontinuity specifications in subsection 2.3.3, which correspond to a very local analysis of the differential slope of the impact on either side of the kink. Second, for the regression kink design specifications we require that the only differential impact of processing time on follow-on innovation around the kink is due to the TRIPS treatment. We therefore verify that there is no bunching of patents on either side of the kink by plotting a histogram of the processing time around the 3 year mark in Figure 2-4. We also study whether there are differences in observable covariates around the kink at zero in Table 2.3, and find no significant differences in most covariates.²² And finally,

²⁰Mehta et al. (2010) find that patent citation profiles start at grant date for patents mainly granted prior to 2001. To the extent that some technology diffusion could start before grant date, and could affect citation profiles in a correlated way to processing time, this will also be incorporated in the measured coefficient on processing time.

²¹The treatment variable is equal to the underlying running variable for positive values, and equal to zero for negative values. The processing time variable is equal to 3 minus the running variable. We define the variables in more detail in subsection 2.2.3.

²²We believe that the significant difference in filing years stems from our sample selection. Since

in subsection 2.3.3 we estimate a regression kink design using patent covariates as outcome variables, and again find no significant differences in covariates around the kink.

2.3 Impact of patent term extension

In this section, we analyze the main impact of patent term extensions on follow-on innovation along three dimensions. First we employ a panel dataset, we then focus on the cross-sectional variation to investigate longer-lived TRIPS impacts, and finally we implement a regression kink (RK) design.

2.3.1 Panel evidence

We follow a diff-in-diff framework with variable treatment intensity, and include years between 1990 and 2005: four calendar years before the TRIPS reform and up to 10 years later. We think of 1995 as the first year of the reform. The URAA was signed in December 1994, and there was considerable uncertainty as to whether it would be approved until late November. Likewise, even though the implementation of the act starts in June 1995, all of its provisions and subsequent impact on the term of existing patents were known by the time of its approval.

In order to investigate the effect of patent term extensions on citation counts, we estimate the following model for calendar years 1990 until 2005:

$$\begin{aligned} \ln(1 + \text{citations}_{it}) = & \alpha_i + \delta_t + \beta_1 \text{treat}_i \times \text{post}_t + \beta_2 \text{proctime}_i \times \text{post}_t & (2.1) \\ & + \eta_1 \text{treat}_i + \eta_2 \text{proctime}_i + \eta_3 \mathbb{1}_{\{\text{treat}_i > 0\}} \times \text{post}_t + \epsilon_{it}, \end{aligned}$$

where the outcome citations_{it} is a yearly count of citations to the focal patent i by subsequent patents filed at t , treat_i is the patent's TRIPS treatment variable in years, proctime_i is its processing time, post_t is an indicator variable equal to one in years 1995 and up, and α_i is a patent-level fixed effect, and also included is a dummy variable for positive treatment interacted with post_t . We include a number of different specifications of fixed effects in the estimations. In the baseline specification, we

we select patents based on their grant year, patents with a longer processing time (i.e., untreated) will mechanically have an earlier filing date.

include a set of fixed effects for calendar year times HJT technology subclass. These fixed effects absorb common technological shocks that affect patent citations due to time-varying attractiveness of the different technologies. We also include a dummy variable for years in which the patent is lapsed, to account for differential citation patterns under and without patent protection. Finally, standard errors are clustered at the patent level to allow for serial auto-correlation over time, and we multiply the outcome variable by 100 to express coefficients in terms of log points.

In Table 2.4, we show the coefficients on patent term extension and processing time interacted with the post-TRIPS dummy. By including both treatment and processing time variables, the treatment effect is identified as the differential effect on citations of having a shorter processing time on either side of the three year cut-off mark. That is, by comparing the effect of shorter processing time when the processing time was under three years (patents that are treated) versus the effect of shorter processing time when the processing time was above three years (patents that are not treated). Panel A shows results using citations stemming from patents filed by other assignees, while Panel B shows results using own citations. We analyze both outcome variables, since we believe that the effects of patent protection on follow-on innovation could differ depending on whether the subsequent innovator owns the focal patent or not. Differences in the number of observations between columns (1) to (3) are solely driven by singletons within FE groups being dropped.²³ In column (1), we show the coefficients for the baseline specification and find no statistically significant effect of term extension on post-treatment citation counts by others. With 95% confidence, we find an impact of an extra year of patent term on citation counts by others between -0.55% and 0.63%.²⁴ Meanwhile, we find that patents with longer processing times do exhibit larger citation counts by others post-TRIPS, by about 0.5% per year of processing time, although the effect is only marginally significant. In column (2), for our preferred specification, we add a set of more flexible fixed effects: we interact HJT technological subclass with application year and calendar year, as well as HJT subclass interacted with grant year and calendar year. This accounts fully for the age profile of patent citation depending on their application year times HJT subclass, as well as depending on their grant year times HJT subclass. In this specification, both coefficients on patent term extension and processing time remain essentially unchanged. In terms of the impact of an extra year of patent term on

²³In order to maintain consistent standard errors singletons should be dropped, see Correia (2015).

²⁴For such small coefficients, log points approximately correspond to percentage points.

others' citation counts, we can reject with 95% confidence a reduction larger than 0.56%. In column (3), we add an even more flexible set of fixed effects by interacting HJT subclass with application, grant, and calendar year. This set of fixed effects accounts for flexible age profiles of citations depending on a given patent's technology subclass, grant year, and application year. Note that with this set of fixed effects, the only variation in processing time and hence in term extension that can be used is capped. Within each group, the difference between the patent with the longest processing time (filed for on January 1 of filing year, and granted on December 31 of grant year) and that with the shortest time (filed for on December 31 of the same filing year, and granted on January 1 of the same grant year) must be under two years. Using that smaller variation leads to still insignificant albeit noisier estimates.

In terms of in-citations by subsequent patents filed by the original patentee in Panel B, the specification in column (1) shows a significantly negative treatment impact. However, once we add a set of interacted fixed effects to account for differential citation age profiles for our preferred specification in column (2), the point estimate is largely reduced and becomes insignificant. Adding a fully interacted set of fixed effects in column (3) increases standard errors, but does not vary the point estimates by much.

Because the prior specifications conflate the impacts of patent term increases on both citations after patent lapse, and possible strategic responses prior to patent lapse, we restrict the analysis to years during which the patent is still outstanding in column (4). We find no evidence of strategic response to patent term increases by either other agents or the original focal patentee. Finally, in column (5) we investigate possible behavioral responses due to the salience of the patent term increase in the year of the reform by focusing exclusively on the impact of the TRIPS reform in 1995. That is, we restrict the analysis from years 1990 to 1993 in the pre-period and only 1995 as the post-period in column (7).²⁵ Once again, we find no evidence of response by citations in 1995.

We examine further whether we are missing out on yearly variation in the impact of TRIPS by interacting the treatment and processing time variables with calendar

²⁵We leave out 1994 from the pre-period because of the TRIPS reform passing in 1994. If there was an effect on citations that started already in 1994, keeping it in the pre-period sample would lead to underestimating the true magnitude of the effect.

year dummies rather than the post-treatment dummy.²⁶ The estimation results are shown in Figures 2-5 and 2-6, with 1994 used as the baseline year. Figure 2-5 plots the treatment coefficients β_k of equation 2.1 without controlling for processing time, as well as their 95% confidence intervals.²⁷ By not controlling for processing time, we are conflating here the impact of TRIPS treatment as well as omitted variables related to extended processing time in the point estimates. The figure suggests that in 1995, the first year of treatment, receiving a term extension of one year is associated with about 0.75% fewer citations. Moreover, it shows no impact between 1996 and 1998, followed by significantly negative impacts of around 1% in the 2000s. However, we also observe positive and statistically significant coefficients prior to 1994, when the TRIPS agreement was signed. This is suggestive of selection into treatment, as discussed previously in subsection 2.2.4.

In order to take selection concerns into account, we control directly for processing time interacted with year dummies in Figure 2-6. By including both treatment and processing time variables, the treatment effect is now identified as the differential effect on citations of having a shorter processing time on either side of the three year cut-off mark. That is, by comparing the effect of shorter processing time when the processing time was under three years (patents that are treated) versus the effect of shorter processing time when the processing time was above three years (patents that are not treated). Including these controls results in no differential pre-trend. However, although the point estimates do not vary much relative to Figure 2-5, the standard errors increase, leading to no significant treatment effect on the outcome variable after treatment until 2001. Starting in 2002, we find some evidence of a negative impact of patent term extensions.

2.3.2 Cross-sectional evidence

In order to investigate the impact of patent term extensions further, we use a cross-sectional analysis where each observation is a patent. This allows us to analyze the impact of the TRIPS treatment on all citations until the end of our sample period,²⁸ rather than restrict the analysis to a fixed set of years post-reform. It is especially useful considering Figure 2-6, which seems to show significant impacts for later sample years. Just as before, we restrict the analysis to patents granted between 1980 and

²⁶Also interacted with calendar year is a dummy for positive treatment.

²⁷Standard errors here are clustered at the patent level.

²⁸December 16, 2016

1989 affected by TRIPS, i.e. still outstanding in June 1995,²⁹ and restricted to patents with a processing time running variable between -1 year and +1 year. That is, patents whose processing time lasted from 2 years to 4 years. This restriction allows us to analyze the differential impact of processing time more narrowly on both sides of the kink between more comparable patents, where the identifying assumption is more likely to hold. We also consider alternative restrictions on the application and grant years, as well as processing time in some specifications in order to check for robustness. Our baseline specification is

$$\ln(1 + \text{postcites}_i) = \alpha_{g(i)} + \beta \text{treat}_i + \gamma \text{proctime}_i + \delta \ln(1 + \text{precites}_i) + \eta X_i + \epsilon_i, \quad (2.2)$$

where the outcome variable postcites_i is a count of in-citations starting in 1995,³⁰ $\alpha_{g(i)}$ are a set of fixed effects defined in greater detail later on, treat_i is the TRIPS patent term extension treatment variable, proctime_i is the underlying processing time running variable, and precites_i is a count of in-citations by patents filed between 1990 and 1993. Covariates X_i include controls for focal patent originality, logs of out-citation counts plus one and number of claims, type of assignee, and dummies for the number of maintenance fees paid for in order to proxy for patent quality. The baseline fixed effects consist in a dummy for positive treatment, as well as a set of HJT technological subclass times application year, and a set of HJT subclass times grant year. These fixed effects absorb differential post-treatment citation counts across year times technology subclass groups of patents, and ensures that we compare similar patents. In some specifications, we fully interact the previous dummies to provide for more flexible fixed effects, and we also include cubic controls in the processing time running variable. Finally, we multiply all the coefficients by 100 to interpret them in terms of log points, and cluster standard errors at the application year times HJT subclass level.

The first set of coefficients are shown in Table 2.5, looking at both citations by other assignees in Panel A and at own citations in Panel B. The coefficients represent the impact in log points of the patent term extension (and underlying processing

²⁹Although the term of any patent granted after 1980 would extend to at least 2002, utility patents filed after December 11, 1980 have to pay renewal or maintenance fees prior to years 4, 8, and 12 in order to remain valid. If these fees are not paid, the patent expires automatically.

³⁰The citation data we use is current up to December 2016. Truncation concerns are likely small since the considered focal patents are all granted before 1990.

time respectively) on the citation outcome. Column (1) shows coefficients from our preferred baseline specification, with the coefficient of interest being statistically insignificant in both panels. In terms of magnitude of the point estimate, an extra year of patent term protection leads to a reduction in citations of around 1.5% post-reform. We can reject with 95% confidence that the actual impact is a reduction larger than 6%. Adding more flexible cubic controls for the processing time running variable in column (2) increases the standard errors, and results in a more negative point estimate on others' citations, but the effect remains insignificant in both panels. We add a set of interacted fixed effects for application year times grant year times HJT subclass in column (3), and the coefficients and standard errors do not vary much from the baseline. In column (4), we restrict the sample more narrowly around the kink, within a band of 6 months of processing time above and below the 3-year cutoff. When comparing these more similar patents, and although the point estimate in Panel A is more negative than the baseline, the impact of increased patent terms is still insignificant. Because of the reduced sample size, the estimates are noisier, but we can still reject a reduction of 19% due to the a one-year patent term increase among this sample of highly comparable patents.

Because the decision on whether to pay maintenance fees can be influenced by a number of variables, including the expectation of patent term length, we restrict the analysis to patents that are not required to pay a fee during 1994. In 1994, the contents of the TRIPS reform were known but it was unclear whether it would pass or not, nor which patents it would affect. As a result, any fee payment decision would have incorporated a possible option value of keeping the patent valid until the reform is implemented, which could introduce selection into the set of renewed patents. As a result, we restrict the analysis to patents granted in 1987-89, which paid their 4th year maintenance fee by 1993 and do not need to pay their 8th year fee until at least 1995. Likewise, with patents granted between 1983 and 1985 for the 8th and 12th respectively. Because of the reduced sample size, the estimates are noisier than for the baseline, but we still find them to be statistically insignificant and in the same ballpark of magnitude. We investigate possible selection concerns further by including all lapsed patents in the analysis and estimating a 2SLS specification in column (7) in which we instrument for actual treatment intensity (including zeros for lapsed patents) using the treatment patents would have received had they not

lapsed.³¹ The results, similar to those in the baseline specification, comfort us in that there does not seem to be selection concerns.

In Table 2.6, we analyze the effect of patent term extensions on other outcome measures. The same baseline specification as in Table 2.5 is shown in column (1), followed by a specification in which the outcome variable is the count of citations by patents filed in 1995 only. Figure 2-5 suggested that, when not controlling for the underlying running variable, most of the effect of the TRIPS treatment was localized in 1995. We confirm here further that, with the proper controls for the processing time running variable, there is no negative impact of term extensions. If anything, the impact on others' citation is marginally significantly positive. We then look at citation counts after patent expiry as well as after the TRIPS reform but prior to patent expiry. We find that the patent term increase leads to significant decreases in citation counts by others after expiry, and to similar increases before lapsing. However, this symmetric effect is mechanical: if the citation rate per year is constant, as the patent term increases it encompasses more years and thus more citations. Symmetrically, less citations are included in counts after expiry. In fact, the effect breaks down once we include a set of dummies for the effective year of patent expiration in unreported specifications. In column (5), we restrict the set of considered patents to patents granted between 1985 and 1989, and returning to the baseline outcome variable of post-treatment citations. The impact of an extra year of patent protection on others' citations among these patents is still insignificant, and we can reject reductions in citation counts larger than 5.3%. For additional robustness tests, in Table 2.7, we estimate the impact of patent term extensions on citation counts after the TRIPS reform, three years at a time. Across all of the considered gaps of time, the effects are remarkably consistent, and very close in magnitude to the baseline impact in Table 2.5.

Across the empirical analysis, we have consistently found no impact of TRIPS-induced patent term extension on follow-on innovation. Nonetheless, we have so far focused on citation counts and citation counts by other assignees as our preferred measures of subsequent innovation. Galasso and Schankerman (2015) show that patent invalidation leads to increases in patent citations, driven by entry of assignees into the technological field, which they measure using counts of assignees that cite the focal

³¹The unreported Kleibergen and Paap (2006) rk Wald F-stat is around $8.5 * 10^3$ in both panels.

patent rather than patent counts. We study this question in Table 2.8, where we estimate our usual specifications from Table 2.5 but using counts of citing assignees rather than citing patents post-TRIPS, pre-reform, and between 1990 and 1993. For the five specifications we show here, which correspond to the same specifications as in Table 2.5, the coefficients and their significance are very similar to those on citation counts. Therefore, there does not seem to be a differential impact on follow-on innovation, be it in the extensive or intensive margin.

2.3.3 Regression Kink design

We analyze the impact of patent term length on follow-on innovation further using a regression kink (RK) design. This method exploits discontinuous variation in the slope of a policy variable and determines whether there is a subsequent discontinuous change in the slope of the outcome. RKs are analogous to regression discontinuity (RD) designs, but analyze slopes instead of levels. In our setting, the treatment (i.e. the patent term extension) is a function of the patent's processing time with a kink or discontinuous slope at 3 years. By comparing how the outcome varies in the underlying processing time running variable locally on either side of the kink, we can identify the average treatment effect of increasing the patent term conditional on the processing time being equal to 3 years. This corresponds to the *treatment on the treated* for patents at the kink point (Ganong and Jaeger, 2016).

The RK design provides an estimate of treatment that is interpretable causally, provided there are no other discontinuities at the kink point that also affect the slopes of the outcome variables locally. We are not aware of any other institutional quirks that would occur at a processing time of 3 years. In Figure 2-4, we show a histogram of the distribution of processing times around the 3 year kink for patents in our sample. We find no evidence of bunching on either side of the kink, which would be indicative of patent assignees or the USPTO strategically responding to other institutional incentives.³² Moreover, we investigate whether patent covariates vary discontinuously around the kink by estimating RK designs in Table 2.9 using patent covariates determined prior to the TRIPS reform as outcome variables. We evaluate a specification without any controls nor fixed effects in Panel A, and in Panel B a specification in which we estimate the RK design on the residual of a

³²The whole process of patent granting for the patents in our sample occurred over 5 years prior to TRIPS, so we expect no strategic bunching in prevision to the reform.

regression of the covariate on dummies for: HJT subclass times grant year, HJT subclass times filing year, and count of maintenance fees paid for the patent. In each column we indicate the standard RK treatment effect estimate at the 3-year cutoff, p-values and 95% confidence intervals based on the procedure developed by Calonico et al. (2014).³³ Optimal bandwidths are selected based on Calonico et al. (2014) with cluster-robust variance estimator at the application year times HJT technology subclass level. Finally, the observation count corresponds to the effective number of observations considered within the bandwidth span. We find no statistically significant difference in the slopes of our covariates on either side of the kink in Table 2.9, with or without controls.

Table 2.10 reports the RK estimates, defined as in the previous paragraph, for the outcomes of interest. In Panel A, we estimate the RK specifications without controlling for covariates nor fixed effects, and find only a marginally significant treatment effect on post-TRIPS citation counts, but no significance in the estimates on citation counts between 1995 and 2000, on citation counts after patent lapse, on post-TRIPS citation counts before patent lapse, nor on 1995 citation counts. In Panel B, we first regress the outcome variables on a set of controls³⁴ and then estimate the RK specification on the residuals.³⁵ We find that the estimates with controls are smaller in absolute value than without controls, and all of them remain insignificant.

We carry out a final RK analysis in Figures 2-7 and 2-8. We return to the panel data and estimate RK specifications for each of the calendar years between 1990 and 2005, with or without controls.³⁶ We plot the RK estimate and robust 95% confidence intervals³⁷ without controls in Figure 2-7, and with controls in Figure 2-8. We find in both figures that the patent term extension treatment effects are statistically indistinguishable from zero for all calendar years, except for 2002. We interpret the results for that specific year as not very informative: first, because as

³³Ganong and Jaeger (2016) finds that this procedure leads to an adequate test size. Notice that the confidence intervals are not necessarily centered around the point estimate.

³⁴The controls include: log of pre-TRIPS citation count plus one, originality, log of out-citations plus one, log of number of claims, and a set of dummies for HJT subclass times grant year, HJT subclass times filing year, maintenance fee counts, and foreign assignee.

³⁵This is not ideal, since the Calonico et al. (2014) optimal bandwidth selection also depends on the covariates. However, because of the number of covariates, and in particular the fixed effects, running the full procedure is computationally problematic.

³⁶The controls include originality, log of number of claims, log of out-citations plus one, and a set of dummies for: HJT subclass times grant year, HJT subclass times filing year, foreign assignee, maintenance fee counts, and a dummy for years in which the focal patent is lapsed.

³⁷Notice that the Calonico et al. (2014) robust C.I. are not necessarily centered around the point estimate.

we are estimating 16 different regressions, the fact that one turns significant at 5% is close to the size of the test; second, because we find in unreported specifications that the significant effect in 2002 is highly dependent on the size of the selected bandwidth.

All in all, the coefficients of interest found in the RK specifications in this subsection are noisily estimated. As a result, it is difficult to meaningfully talk of the magnitudes of effects rejected by these tests. Nonetheless, we find comfort in the fact that the estimated coefficients in this analysis are largely statistically insignificant, which is consistent with our previous findings in the panel and cross-sectional specifications.

2.4 Heterogeneity

The null results we have found in the previous section are average impacts, and could be masking large heterogeneous effects. For example, a strong positive impact of patent term extension on follow-on innovation in mechanical technologies could cancel out with a large negative impact on drugs patents. In this section, we investigate possible heterogeneous effects of patent term extension along a number of observable dimensions.

2.4.1 By technology class

The impact of a patent term duration could be highly heterogeneous across technology classes for a number of reasons. First, the strategic use of patents as an exclusionary mechanism varies by industry and technology. Second, theoretical models of follow-on innovation tell us that the impact of IP rights on subsequent innovation will be more marked in instances where bargaining is more complex (Ziedonis, 2004; Galasso and Schankerman, 2015). This could lead to heterogeneous impacts depending on the fragmentation of patent ownership in the technology class. Finally, the intrinsic cumulative nature of innovation also varies by technology, and could also affect the impact of of patent term extensions.

In order to investigate possible heterogeneity by technology class, we allow the coefficients on the treatment and on the underlying processing time running variable to vary by HJT technology class. The coefficients of interest are shown in Figure 2-9, with 95% confidence intervals,³⁸ and show that the impact of patent term on

³⁸The confidence intervals are constructed from robust standard errors, rather than clustered.

follow-on innovation is statistically insignificant and relatively homogeneous across all technology classes. We analyze heterogeneity further by separating technologies into *complex* and *discrete*.³⁹ The coefficients, also shown in Figure 2-9, remain small and statistically insignificant in both instances. Moreover, in unreported specifications we also analyze the effect of patent-holder concentration on the impact of term extensions and find no heterogeneous effect. We measure concentration as the share of patents applied for by the four largest patent-filers in each technology subclass between 1990 and 1993, prior to the TRIPS reform.

2.4.2 By grant and filing year

We further analyze heterogeneous impacts by patent grant years, using our baseline specification for years ranging between 1980 and 1989 but allowing for the coefficients on treatment and running variable to vary by grant year. The resulting coefficients are shown in Figure 2-10, with their corresponding 95% confidence intervals, again using robust rather than clustered standard errors. The figure shows no significant coefficients and relatively homogeneous effects across grant years. This is consistent with our previous finding of no impact of the rescaled treatment in order to proxy for patent value increase. Rescaling the treatment by the amount of patent term left after the TRIPS reform is akin, albeit less flexibly so, to allowing for the treatment effect to vary by grant year. We find no significant impact in either specification.

We then study the heterogeneity by application year, for patents granted between 1980 and 1989 by allowing the coefficients to vary by application year. We restrict the data to patents filed between 1977 and 1986, since they are the only ones that provide variation on both treatment intensity and the processing time running variable for zero treatment (processing time longer than three years). The coefficients of interest on the treatment intensity are shown in Figure 2-11, together with the corresponding 95% confidence intervals for robust standard errors. We see that there is only one statistically significant coefficient at the 5% level for 1980, and all magnitudes are relatively homogeneous. We do not see the 1980 coefficient as strong evidence of an impact of patent term extensions on follow-on innovation. Instead, we believe it is a

This tends to go against finding insignificant effects.

³⁹This distinction follows Levin et al. (1987); Cohen et al. (2000), in differentiating “whether a new, commercializable product or process is comprised of numerous separately patentable elements versus relatively few”. Following Galasso and Schankerman (2015), complex technologies comprise HJT classes 2 and 4, as well as subclasses 32 and 33. Discrete technologies comprise all other technologies.

by-product of estimating a large number of different specifications, which should lead to some of them providing significant estimates.

2.4.3 By assignee size

As discussed before, fragmentation of the distribution of patent-holders can lead to more costly bargaining, itself resulting in decreases follow-on innovation. We therefore investigate the heterogeneous impact of patent term extensions on subsequent innovation depending on the size of the original assignee of the focal patent. We characterize assignees depending on the size of their portfolio of patents, among all patents filed between 1990 and 1993. We classify assignees with a portfolio larger than the 95th percentile (15 or more patents) as *large*, assignees with one or less patents filed between 1990 and 1993 are classified as *small*, and all others in-between are categorized as *medium*. The resulting coefficients are shown in Figure 2-12, and are all small and statistically insignificant.

2.4.4 By patent quality

Finally, we analyze heterogeneity in the effect of patent term extensions based on underlying patent quality.⁴⁰ The distribution of patent citations is highly skewed, and by analyzing average effects over the whole population, we could be missing significant effects on the most important patents. We thus separately estimate the impact of TRIPS-induced patent term extensions on four categories of patents, depending on their pre-TRIPS citation counts. For each HJT class times grant year we calculate the distribution of citation counts between 1990 and 1993, and classify patents depending on the percentile of the distribution they fall in within their group.⁴¹ We plot the coefficients of interest as well as their 95% confidence intervals for standard errors robust to heteroskedasticity in Figure 2-13 separately for patents in the bottom 10% of the within-group citation distribution, patents between the 10th percentile and the median, patents between the median and the 90th percentile, and patents in the top 10% of the distribution. We see that all four coefficients are statistically indistinguishable from zero, and the point estimates fall within a relatively homogeneous range of -3 to +8 log points by additional patent term year. Moreover, in

⁴⁰Where patent quality is understood as "the magnitude of inventive output associated with them" (Griliches, 1990), and is positively associated to citation counts.

⁴¹This allows us to compare standardized patent citation counts.

additional unreported tests, we carry out a yearly RK analysis as in Figures 2-7 and 2-8, separately by patent quality category. The resulting estimates are consistently statistically insignificant across all four quality groups.

2.5 Discussion

The predictions on the effect of patent protection on follow-on innovation established by an ample body of theoretical literature are ambiguous. Kitch (1977) argues that patents allow for coordination of subsequent innovation by the focal patentee. By fostering information sharing and thus avoiding duplication of research costs, this coordination can lead to increased innovation downstream. Moreover, patents can facilitate technology transfers across firms and foster a market for ideas that can increase efficiencies (Arora et al., 2001). On the other hand, Scotchmer (1991) and Green and Scotchmer (1995) emphasize how bargaining failures between the first and subsequent innovators could depress sequential innovation. These bargaining failures stemming from asymmetric information can be exacerbated by the existence of transaction costs between the parties (Bessen, 2004; Anand and Khanna, 2000).

More recently, a literature has sought to evaluate the impact of IP protection on follow-on innovation empirically. It has consistently found that non-patent IP protection hinders reuse and in general reduces subsequent research and innovation (Williams, 2013; Murray et al., 2016; Moser and Voena, 2012; Biasi and Moser, 2017; Nagaraj, 2016). Regarding patent protection, the evidence is more mixed. Galasso and Schankerman (2015) show that patent invalidations by the Court of Appeal for the Federal Circuit, with exclusive jurisdiction over patent-related appeals since 1982, lead to increases in citation counts in the order of 50%. This effect is very heterogeneous, and is concentrated on patents with high unobserved quality as well as in technological areas characterized by their complexity⁴² and patent-holder fragmentation. Moreover, they find the effect to be primarily driven by invalidations of patents held by large firms that lead to entry of small subsequent innovators. On the other hand, Sampat and Williams (2015) show that patent protection on human genes have no impact on follow-on research and innovation.

This paper contributes to the literature by analyzing the intensive side of patent

⁴²Following a distinction between complex and discrete by Levin et al. (1987); Cohen et al. (2000), in differentiating “whether a new, commercializable product or process is comprised of numerous separately patentable elements versus relatively few”.

protection, and providing evidence consistent with no significant impact of patent term extensions on follow-on innovation, measured in the form of subsequent patent citations. In our baseline panel specifications we can reject decreases in citation counts larger than 0.6% for every extra year of patent term awarded for the initial 11 years after reform, while in our cross-sectional specifications encompassing all citations until 2016 we can reject decreases in citation counts larger than 6%. These results contribute to our understanding of how the intensive margin of patent protection impacts cumulative innovation, which to the best of our knowledge has not been previously investigated empirically.

Moreover, to the extent that the extensive and intensive impact of patent protection on cumulative innovation are related, this paper can also shed light on the impact of patent protection on follow-on innovation. Our results, together with Sampat and Williams (2015) suggest that patent protection has an insignificant impact on follow-on innovation. Meanwhile, Galasso and Schankerman (2015) find that the impact of patent invalidations decreases with patent age, with no significant effect on patents with over 15 years of age. The average age of their patent sample is 10 years, with 10 years of protection remaining. Therefore, their results could be read as a 5% citation reduction per extra year. We cannot reject an effect of such size in the long term. Although it falls well wide of the 95% confidence interval in our preferred panel specification, it remains close to the edge of the confidence interval in our baseline cross-sectional specification.⁴³

2.6 Conclusion

In this paper, we study the impact of patent term length on follow-on innovation. We leverage a quasi-natural experiment in 1995, the TRIPS reform, that extended patent terms on established still outstanding patents. The patent term increase was a step-wise linear function of the patents' processing time, and presented a kink at a processing time of 3 years. We leverage this kink to identify the causal effect of term lengthening on subsequent citation counts, under the assumption that the underlying effect of processing time on citation counts is the same on both sides of the kink.

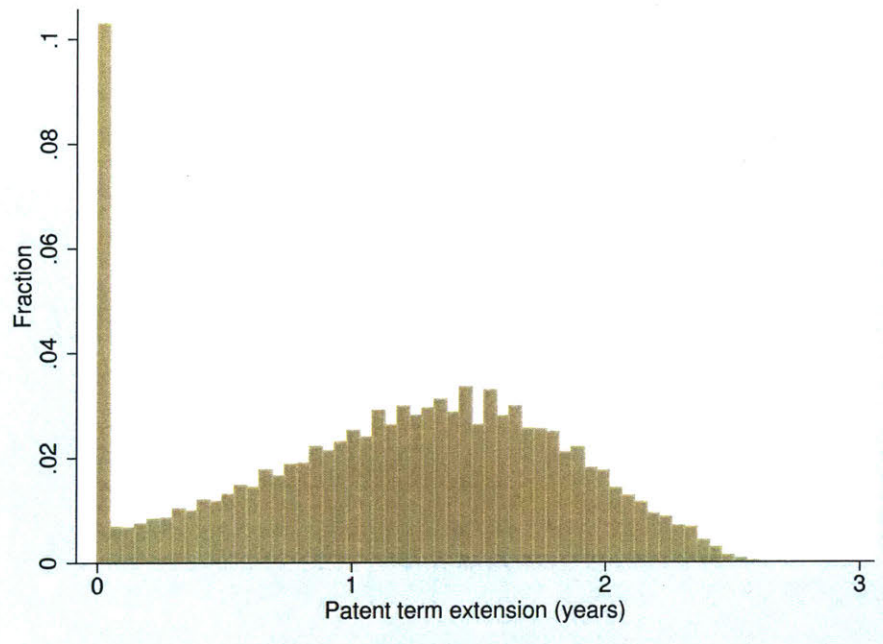
First, we present results from panel and cross-sectional specifications that show no average effect of patent term increases on follow-on citations on treated patents.

⁴³More specifically, it falls at the 6th percentile of the normal distribution of the true parameter centered around the baseline point estimate.

Our panel evidence allows us to reject citation count decreases larger than 0.6% per additional patent term year, and our noisier cross-sectional estimates reject impacts below -6%. We then restrict to a more local analysis using regression kink designs. These designs ensure that the identifying assumption is more likely to hold, as it only has to hold locally, but at the expense of possible lack of external validity. Our regression kink design specifications provide corroborate the absence of a statistically significant impact of patent term extensions on cumulative innovation, although the standard errors are too noisy to provide meaningful inference.

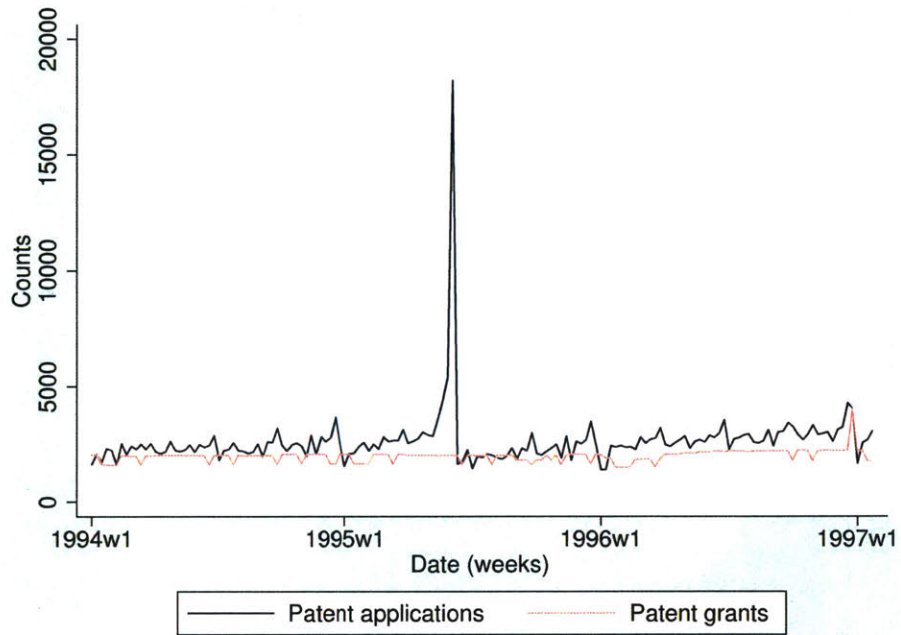
This paper is, to the best of our knowledge, the first formal empirical study on the impact of patent term length on follow-on innovation. It contributes to an extensive theoretical literature on the optimal breadth and length of patent protection, as well as a smaller and more recent empirical literature focusing on the impact of IP protection on cumulative innovation.

Figure 2-1: Histogram of patent term extensions due to TRIPS



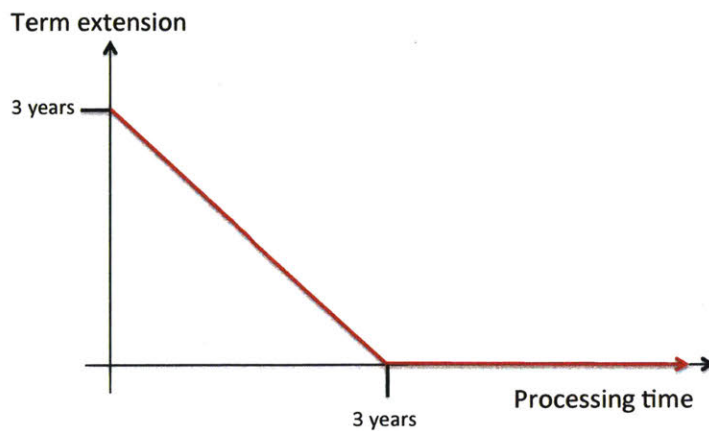
Notes: Histogram of patent term extensions in years due to the TRIPS reform for patents affected by TRIPS.

Figure 2-2: Patent applications and grants around TRIPS implementation



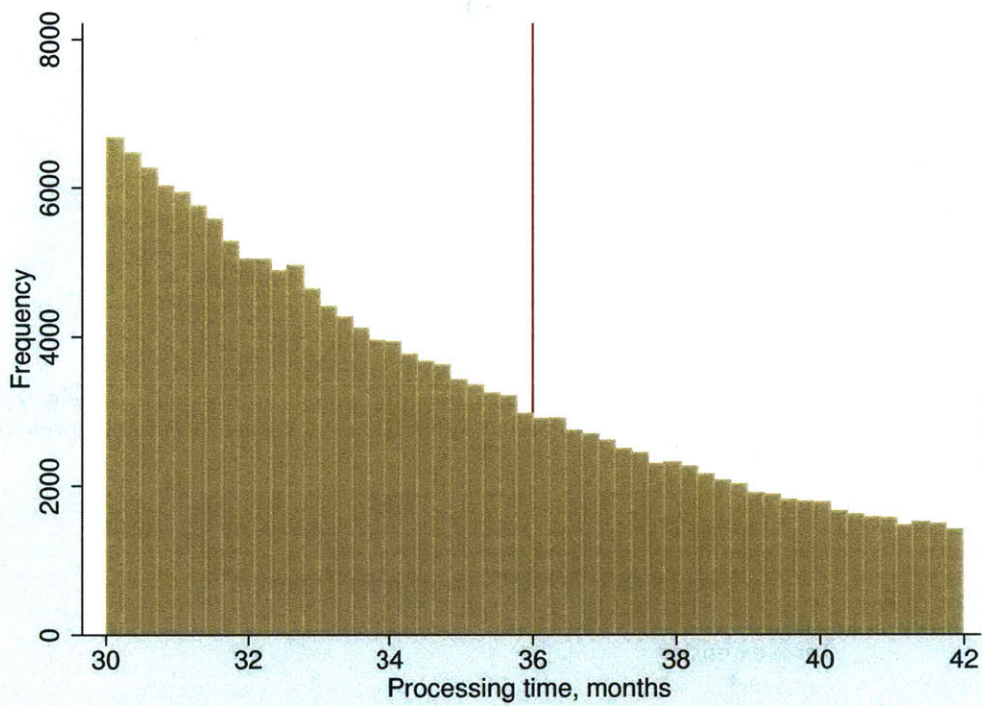
Notes: Plot of weekly patent filing counts in dark blue and patent grant counts in light red at the USPTO. Patent filing counts correspond to patent applications for patents eventually granted before December 16, 2016.

Figure 2-3: TRIPS treatment function




Notes: Plot of TRIPS treatment function in red. The x-axis is patent processing time, defined as grant date minus application date. The y-axis is patent term extension due to TRIPS.

Figure 2-4: Histogram of processing time around 3-year TRIPS cutoff.




Notes: Histogram of patent processing times in months around the three year mark (vertical red line). The sample of patents corresponds to patents affected by TRIPS.

Figure 2-5: Difference in yearly citation counts per extra term day




Notes: This figure plots the values and confidence intervals of the coefficient on treatment by calendar year between 1990 and 2005 for the panel specification, with 1994 taken as baseline year. Patents considered are all patents granted in the 1980s still outstanding on June 8, 1995. Citation counts considered are citations by other firms. The specification includes patent and calendar year times HJT subclass fixed effects, as well as a lapsed-patent dummy, but no controls for the patents' processing time interacted with calendar year. Standard errors are clustered at the patent level, and confidence intervals are set at the 95% level.

Figure 2-6: Difference in yearly citation counts per extra term day, controlling for processing time




Notes: This figure plots the values and confidence intervals of the coefficient on treatment by calendar year between 1990 and 2005 for the panel specification, with 1994 taken as baseline year. Patents considered are all patents granted in the 1980s still outstanding on June 8, 1995. Citation counts considered are citations by other firms. The specification includes patent and calendar year times HJT subclass fixed effects, a lapsed-patent dummy, and controls for the patents' processing time. Standard errors are clustered at the patent level, and confidence intervals are set at the 95% level.

Figure 2-7: Difference in yearly citation counts per extra term year, RKD without controls




Notes: This figure plots the values and confidence intervals of the coefficient on treatment for the regression kink design specification for each calendar year between 1990 and 2005. Each year corresponds to a separate RK specification, without fixed effects nor any other covariate controls. Plotted are the non-parametric estimates of the treatment effect from local quadratic regressions, as well as bias-corrected robust 95% confidence intervals based on Calonico et al. (2014) optimal bandwidth selector with cluster-robust variance estimator at the application year times HJT technology subclass level.

Figure 2-8: Difference in yearly citation counts per extra term year, RKD with controls




Notes: This figure plots the values and confidence intervals of the coefficient on treatment for the regression kink design specification for each calendar year between 1990 and 2005. Each year corresponds to a separate RK specification on the residual of the post-TRIPS citation outcome on a set of fixed effects and covariate controls. Plotted are the non-parametric estimates of the treatment effect from local quadratic regressions, as well as bias-corrected robust 95% confidence intervals based on Calonico et al. (2014) optimal bandwidth selector with cluster-robust variance estimator at the application year times HJT technology subclass level.

Figure 2-9: Difference in yearly citation counts by technology type



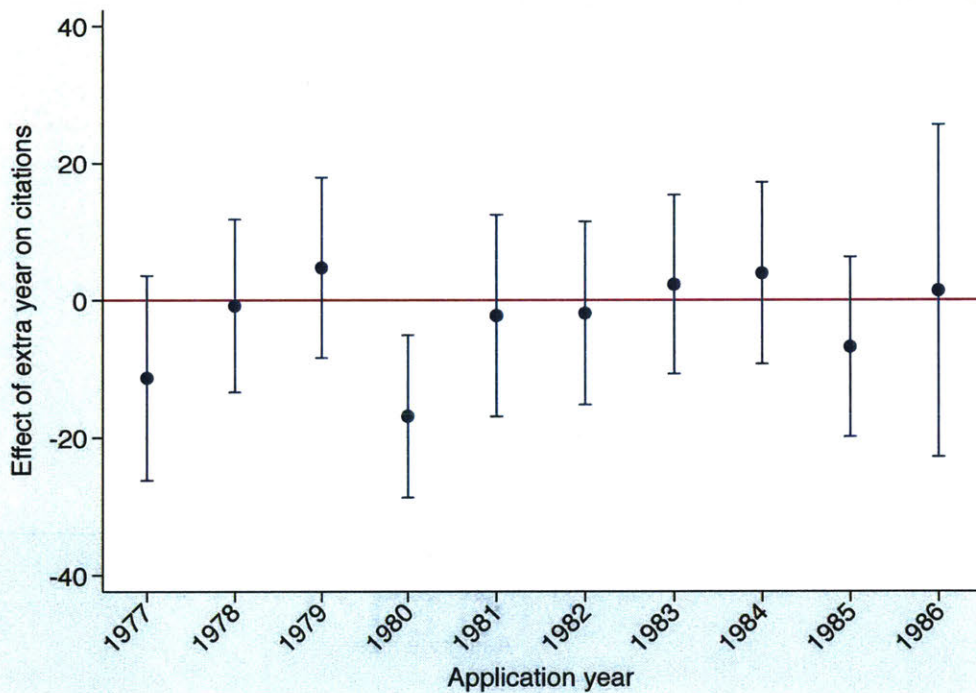
Notes: This figure plots the values and confidence intervals of the coefficient on treatment, separated by technology classes and types, for the cross-sectional specification. The outcome variable is log of post-TRIPS citations by others, scaled up by 100. The specification includes the baseline fixed effects for the cross-sectional specifications, covariate controls and controls for the patents' processing time. Chemical corresponds to HJT class 1, Communications to HJT class 2 *Computers & communications*, Medical to HJT class 3 *Drugs & Medical*, Electrical to HJT class 4 *Electrical & Electronic*, Mechanical to HJT class 5, and Others to HJT class 6. Complex technologies correspond to HJT classes 2 and 4, as well as subclasses 32 and 33; Discrete technologies span the rest of technologies. Standard errors are robust to heteroskedasticity, and confidence intervals are set at the 90% level.

Figure 2-10: Difference in yearly citation counts by grant year



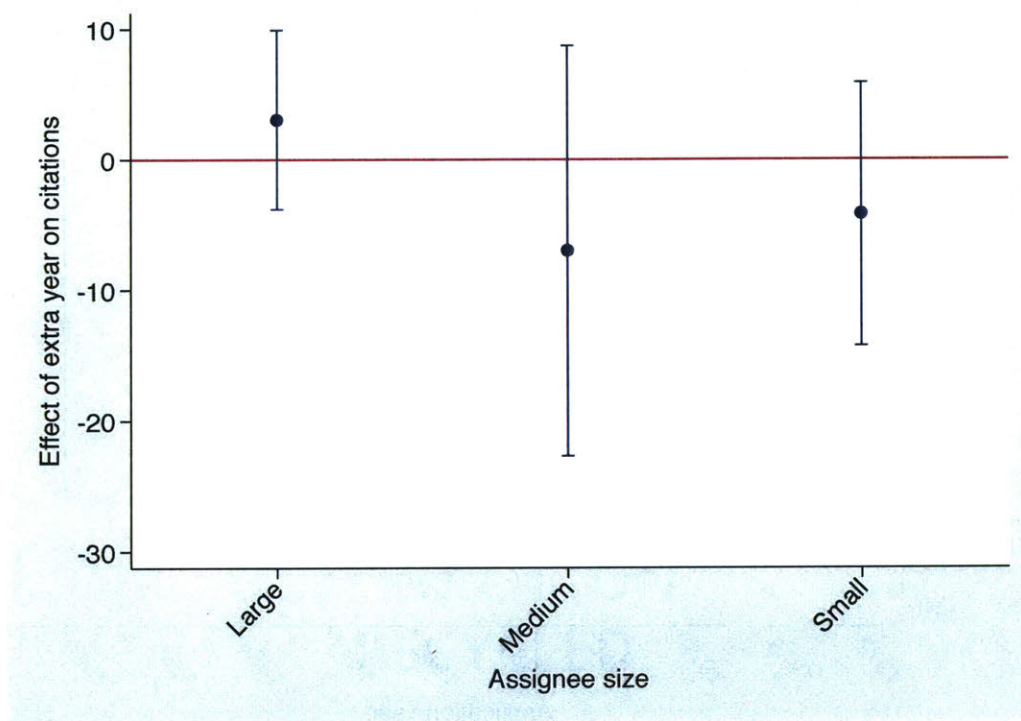
Notes: This figure plots the values and confidence intervals of the coefficient on treatment, separated by grant year of patents, for the cross-sectional specification. The outcome variable is log of post-TRIPS citations by others, scaled up by 100. The specification includes the baseline fixed effects for the cross-sectional specifications, covariate controls and controls for the patents' processing time. Grant years span from 1980 to 1989. Standard errors are robust to heteroskedasticity, and confidence intervals are set at the 90% level.

Figure 2-11: Difference in yearly citation counts by application year



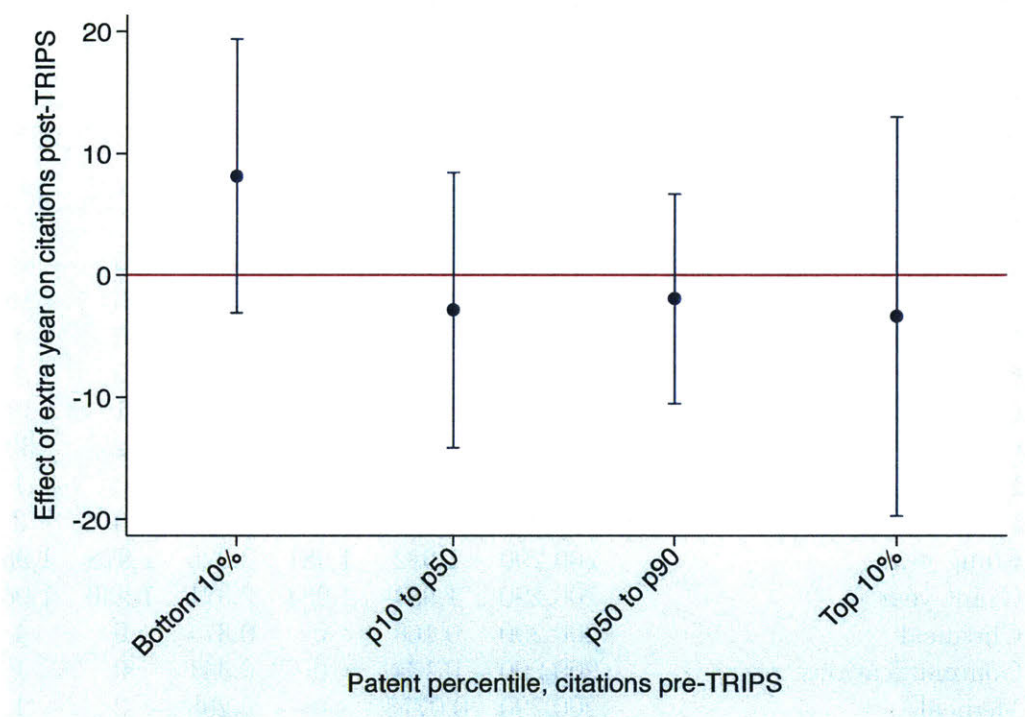
Notes: This figure plots the values and confidence intervals of the coefficient on treatment, separated by filing year of patents, for the cross-sectional specification. The outcome variable is log of post-TRIPS citations by others, scaled up by 100. The specification includes the baseline fixed effects for the cross-sectional specifications, covariate controls and controls for the patents' processing time. Application years considered span between 1977 and 1986 for patents granted between in the 1980s. Standard errors are robust to heteroskedasticity, and confidence intervals are set at the 90% level.

Figure 2-12: Difference in yearly citation counts by assignee size



Notes: This figure plots the values and confidence intervals of the coefficient on treatment, separated by the size of the patent's original assignee, for the cross-sectional specification. The outcome variable is log of post-TRIPS citations by others, scaled up by 100. The specification includes the baseline fixed effects for the cross-sectional specifications, covariate controls and controls for the patents' processing time. Patents considered are those granted between 1985 and 1989 that were still outstanding by June 8, 1995. Large corresponds to assignees filing 15 patents or more between 1990 and 1993, Medium corresponds to between 2 and 14 patents, and Small to 1 or less patents. Standard errors are robust to heteroskedasticity, and confidence intervals are set at the 90% level.

Figure 2-13: Difference in yearly citation counts by pre-TRIPS patent citation count



Notes: This figure plots the values and confidence intervals of the coefficient on treatment, separated by the pre-TRIPS citation count of the patent, for the cross-sectional specification. The outcome variable is log of post-TRIPS citations by others, scaled up by 100. The specification includes the baseline fixed effects for the cross-sectional specifications, covariate controls and controls for the patents' processing time. Patents considered are those granted between 1985 and 1989 that were still outstanding by June 8, 1995. The distribution of pre-TRIPS citations is taken for each HJT class times grant year group, and the patents are classified into the different categories depending on how they fall in their corresponding within-group distribution. Standard errors are robust to heteroskedasticity, and confidence intervals are set at the 90% level.

Table 2.1: Descriptive statistics, variables

VARIABLES	N	mean	p50	sd	p5	p95
Others' citations post-TRIPS	200,290	11.06	4	27.55	0	41
Others' citations pre-TRIPS	200,290	4.779	3	6.144	0	16
Own citations post-TRIPS	200,290	0.774	0	6.165	0	4
Own citations pre-TRIPS	200,290	0.852	0	2.402	0	4
Out-citations	200,290	6.820	6	5.656	1	16
Claims	200,290	11.62	9	10.34	2	30
Foreign	200,290	0.390	0	0.488	0	1
Fees	200,290	1.352	1	1.329	0	3
Filing year	200,290	1,982	1,981	2.925	1,978	1,986
Grant year	200,290	1,984	1,984	2.921	1,980	1,989
Chemical	200,290	0.168	0	0.374	0	1
Communications	200,290	0.134	0	0.341	0	1
Medical	200,290	0.0778	0	0.268	0	1
Electrical	200,290	0.191	0	0.393	0	1
Mechanical	200,290	0.235	0	0.424	0	1
Others	200,290	0.194	0	0.396	0	1

Notes: The statistics are taken over all patents granted between 1980 and 1989 that were still outstanding on June 8, 1995, by paying their required maintenance fees.

Table 2.2: t-tests of covariates for large processing time differences

	Untreated	Treated	Difference	t-stat
Others' citations pre-TRIPS	5.252	4.201	1.051*** (0.132)	7.967
Own citations pre-TRIPS	0.838	0.870	-0.0322 (0.0453)	-0.711
Out-citations	7.784	7.165	0.619*** (0.152)	4.083
Originality	0.768	0.796	-0.0279*** (0.00558)	-5.004
Claims	12.91	12.19	0.716*** (0.225)	3.176
Foreign	0.426	0.438	-0.0120 (0.0108)	-1.108
Fees	2.455	2.358	0.0976*** (0.0167)	5.858
Filing year	1983.9	1985.3	-1.436*** (0.0302)	-47.52
Grant year	1986.9	1987.3	-0.393*** (0.0299)	-13.14

Notes: The table reports means of covariates from two samples, their difference together with standard deviation in column (3), and a t-test of the significance of their difference in column (4). The untreated sample in column (1) corresponds to patents with $procrun \in (-50, 0)$ and treated patents in column (2) with $procrun \in (350, 400)$. These correspond to patents with processing time between 3 years and 3 years + 50 days in column (1) and around two years in column (2). The population of patents considered are granted between 1980 and 1989 and still outstanding on June 8, 1995.

Table 2.3: t-tests of covariates around the kink

	Untreated	Treated	Difference	t-stat
Others' citations pre-TRIPS	5.252	5.287	-0.0348 (0.144)	-0.243
Own citations pre-TRIPS	0.838	0.863	-0.0247 (0.0474)	-0.522
Out-citations	7.784	7.782	0.00213 (0.159)	0.0134
Originality	0.768	0.766	0.00185 (0.00563)	0.328
Claims	12.91	13.19	-0.284 (0.257)	-1.104
Foreign	0.426	0.438	-0.0114 (0.0107)	-1.062
Fees	2.455	2.447	0.00887 (0.0157)	0.564
Filing year	1983.9	1984.0	-0.0889*** (0.0306)	-2.910
Grant year	1986.9	1986.9	0.0309 (0.0303)	1.018

Notes: The table reports means of covariates from two samples, their difference together with standard deviation in column (3), and a t-test of the significance of their difference in column (4). The samples correspond to patents with $procrun \in (0, 50)$ in the treated group in column (2) and patents with $procrun \in (-50, 0)$ in the untreated group in column (1). These correspond to patents with processing time in (3 years, 3 years + 50 days) in column (1) and (3 years - 50 days, 3 years) in column (2). The population of patents considered are granted between 1980 and 1989 and still outstanding on June 8, 1995.

Table 2.4: Panel specifications

	Baseline			Before expiration	Reduced years
	(1)	(2)	(3)	(4)	(5)
Panel A: Others' citations					
Patent extension	0.039 (0.302)	0.075 (0.323)	-0.787 (0.576)	0.510 (0.477)	-0.441 (0.524)
Processing time	0.494* (0.284)	0.537* (0.325)	-0.269 (0.552)	1.207** (0.483)	1.303** (0.527)
Observations	6,709,952	6,709,840	6,707,648	3,718,814	2,096,825
Panel B: Own citations					
Patent extension	-0.742*** (0.113)	0.043 (0.122)	0.137 (0.219)	-0.011 (0.191)	-0.189 (0.218)
Processing time	0.257** (0.106)	0.214* (0.122)	0.301 (0.209)	0.186 (0.193)	0.325 (0.219)
Observations	6,709,952	6,709,840	6,707,648	3,718,814	2,096,825
Fixed effects:					
Patent	✓	✓	✓	✓	✓
Year × tech	✓	✓	✓	✓	✓
Year × tech × grant		✓	✓	✓	✓
Year × tech × app.		✓	✓	✓	✓
Full FE			✓		

Notes: The dependent variable in all columns in Panel A is $100 * \ln(1 + \text{others' citations}_{it})$ for years between 1990 and 2005, and $100 * \ln(1 + \text{own citations}_{it})$ in Panel B. The focal patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. The coefficients shown are interacted with a post-TRIPS dummy. All specifications include the non-interacted terms as well, treatment dummies interacted with post-TRIPS, and a lapsed-patent dummy. Column (1) includes a set of fixed effects encompassing: patent, and HJT subclass times calendar year. Column (2) adds a set of more flexible fixed effects: HJT subclass times grant year times calendar year, HJT subclass times filing year times calendar year. Column (3) adds fixed effects for: HJT subclass times grant year times filing year times calendar year. Column (4) estimates the specification in (2) but only on calendar years when the patent is not lapsed. Column (5) estimates it on years 1990-1993 and 1995. Standard errors in brackets are clustered at the patent level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 2.5: Cross-sectional specifications

	Baseline			Around kink	Years 1987-89	Years 1983-85	Intent- to-treat
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Others' citations							
Term extension	-1.571 (2.194)	-4.265 (8.693)	-0.961 (2.908)	-7.415 (5.819)	1.469 (4.279)	-6.717 (4.080)	-0.111 (2.015)
Processing time	2.962 (2.080)	-5.392 (6.263)	3.579 (2.654)	0.095 (5.136)	6.547 (4.150)	-3.178 (4.342)	4.608*** (1.262)
Observations	200,287	200,287	200,280	75,413	53,989	54,215	259,032
Panel B: Own citations							
Term extension	-1.556 (1.136)	-1.446 (4.549)	-1.035 (1.480)	-0.877 (3.120)	-1.890 (2.190)	-0.293 (2.256)	-1.374 (1.020)
Processing time	-0.637 (1.073)	1.571 (3.394)	-0.149 (1.355)	-1.501 (2.576)	-1.446 (1.985)	0.910 (2.056)	-0.602 (0.661)
Observations	200,287	200,287	200,280	75,413	53,989	54,215	259,032
Covariates	✓	✓	✓	✓	✓	✓	✓
Base FE	✓	✓	✓	✓	✓	✓	✓
Cubic proctime		✓					
Full FE			✓				

Notes: The dependent variable in all columns in Panel A is $100 * \ln(1 + \text{post-TRIPS others' citations}_i)$, and $100 * \ln(1 + \text{post-TRIPS own citations}_i)$ in Panel B. The patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. All specifications include a set of fixed effects for HJT subclass times grant year and HJT subclass times filing year. They also control for originality, log of number of claims, log of out-citations plus one, log of citation counts plus one between 1990-93 as well as up to 1993, and dummies for foreign assignees, as well as for count of maintenance fee payments and for positive treatment (i.e. processing time below 3 years). Column (1) is the baseline specification. Column (2) adds cubic controls in the patents' processing time, only the first term is shown. Column (3) adds a set of more flexible fixed effects to the baseline: HJT subclass times grant year times filing year. Column (4) estimates the baseline specification for a more restricted set of patents around the kink, with processing times between 2.5 and 3.5 years. Column (5) estimates the baseline specification for patents granted in 1987-89. Column (6) estimates the baseline specification for patents granted in 1983-85. Column (7) estimates a 2SLS specification on the previous sample using the possible treatment *had a patent not lapsed* to instrument for the actual TRIPS treatment received. The unreported Kleibergen and Paap (2006) rk Wald F-stat is around $8.5 * 10^3$ in both panels. Standard errors in brackets are clustered at the application year times HJT technology subclass level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 2.6: Alternative cross-sectional specifications

	Baseline	Only 1995	After expiry	1995- expiry	Granted 1985-89
	(1)	(2)	(3)	(4)	(5)
Panel A: Others' citations					
Term extension	-1.571 (2.194)	2.291* (1.201)	-8.956*** (2.117)	7.023*** (1.696)	1.161 (3.210)
Processing time	2.962 (2.080)	4.077*** (1.152)	-5.300*** (1.957)	8.390*** (1.632)	6.960** (3.036)
Observations	200,287	200,287	200,287	200,287	83,981
Panel B: Own citations					
Term extension	-1.556 (1.136)	-0.019 (0.463)	-2.220** (0.866)	0.550 (0.804)	-0.688 (1.671)
Processing time	-0.637 (1.073)	0.411 (0.445)	-1.456* (0.797)	0.716 (0.762)	-0.903 (1.569)
Observations	200,287	200,287	200,287	200,287	83,981
Covariates	✓	✓	✓	✓	✓
Base FE	✓	✓	✓	✓	✓

Notes: All specifications include a set of fixed effects for HJT subclass times grant year and HJT subclass times filing year. They also control for originality, log of number of claims, log of out-citations plus one, log of citation counts plus one between 1990-93 as well as up to 1993, and dummies for foreign assignees, as well as for count of maintenance fee payments and for positive treatment (i.e. processing time below 3 years). The patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. The dependent variables in all columns in Panel A consider only others' citations, while they consider only own citations in Panel B. The dependent variable in column (1) is $100 \cdot \ln(1 + \text{post-TRIPS citations}_i)$. The dependent variable in column (2) is $100 \cdot \ln(1 + 1995 \text{ citations}_i)$. The dependent variable in column (3) is $100 \cdot \ln(1 + \text{post-lapse citations}_i)$. The dependent variable in column (4) is $100 \cdot \ln(1 + 1995\text{-lapse citations}_i)$. Column (5) considers an alternative treatment and processing time definition, where the processing time (and therefore treatment) is normalized by the amount of term left after reform on June 8, 1995. Standard errors in brackets are clustered at the application year times HJT technology subclass level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 2.7: Alternative cross-sectional specifications, citations by years

	1995 -97	1995 -00	1995 -03	1995 -06	1995 -09	1995 -12	1995 -15
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Others' citations							
Term extension	1.872 (1.324)	-0.565 (1.633)	0.088 (1.810)	0.041 (1.897)	0.060 (1.916)	-0.762 (1.979)	-1.153 (2.019)
Processing time	4.668*** (1.052)	3.635*** (1.312)	4.588*** (1.466)	5.246*** (1.525)	5.660*** (1.527)	5.505*** (1.587)	5.355*** (1.610)
Observations	200,287	200,287	200,287	200,287	200,287	200,287	200,287
Panel B: Own citations							
Term extension	-0.363 (0.703)	-0.621 (0.914)	-1.373 (1.038)	-1.341 (1.113)	-1.179 (1.170)	-1.142 (1.215)	-1.344 (1.239)
Processing time	-0.137 (0.554)	-0.364 (0.711)	-1.052 (0.816)	-0.988 (0.872)	-0.756 (0.911)	-0.674 (0.950)	-0.804 (0.971)
Observations	200,287	200,287	200,287	200,287	200,287	200,287	200,287
Covariates	✓	✓	✓	✓	✓	✓	✓
Base FE	✓	✓	✓	✓	✓	✓	✓

Notes: All specifications include a set of fixed effects for HJT subclass times grant year and HJT subclass times filing year. They also control for originality, log of number of claims, log of out-citations plus one, log of citation counts plus one between 1990-93 as well as up to 1993, and dummies for foreign assignees, as well as for count of maintenance fee payments and for positive treatment (i.e. processing time below 3 years). The patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. The dependent variables in all columns in Panel A consider only others' citations, while they consider only own citations in Panel B. The dependent variables are $100 \cdot \ln(1 + \text{citations}_i)$, where citations_i correspond to: 1995-1997 citations in column (1), 1995-2000 citations in column (2), 1995-2003 citations in column (3), 1995-2006 citations in column (4), 1995-2009 citations in column (5), 1995-2012 citations in column (6), 1995-2015 citations in column (7). Standard errors in brackets are clustered at the application year times HJT technology subclass level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 2.8: Extensive margin, assignee counts

	Baseline			Around kink	Granted 1985-89
	(1)	(2)	(3)	(4)	(5)
Term extension	-2.921*	-7.026	-2.935	-5.785	-0.033
	(1.560)	(5.968)	(2.021)	(3.952)	(2.273)
Processing time	0.634	-4.089	0.660	-0.043	3.159
	(1.410)	(4.526)	(1.809)	(3.558)	(2.088)
Observations	200,287	200,287	200,280	75,413	83,981
Covariates	✓	✓	✓	✓	✓
Base FE	✓	✓	✓	✓	✓
Cubic proctime		✓			
Full FE			✓		

Notes: The dependent variable in all columns is $100 \cdot \ln(1 + \text{post-TRIPS citing assignees}_i)$. The patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. All specifications include a set of fixed effects for HJT subclass times grant year and HJT subclass times filing year. They also control for originality, log of number of claims, log of out-citations plus one, log of citation counts plus one between 1990-93 as well as up to 1993, and dummies for foreign assignees, as well as for count of maintenance fee payments and for positive treatment (i.e. processing time below 3 years). Column (1) is the baseline specification. Column (2) adds cubic controls in the patents' processing time, only the first term is shown. Column (3) adds a set of more flexible fixed effects to the baseline: HJT subclass times grant year times filing year. Column (4) estimates the baseline specification for a more restricted set of patents around the kink. Column (5) estimates the baseline specification for patents granted in 1987-89. Column (6) estimates the baseline specification for patents granted in 1983-85. Standard errors in brackets are clustered at the application year times HJT technology subclass level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 2.9: Regression kink discontinuity on covariates

	90-93 cites	Originality	Out-cites	Claims	Foreign	Fees
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Without controls						
Treatment	-34.8	-.0415	-.42	-.108	.227	-.447
Robust p-value	.367	.636	.317	.945	.516	.452
Lower CI	-110	-.443	-1.33	-.85	-.544	-2.21
Upper CI	40.5	.271	.432	.792	1.08	.986
Bandwidth	.812	.647	.69	.806	.583	.726
Observations	58,488	45,324	48,379	58,371	40,133	51,385
Panel B: With controls						
Treatment	-16	-.0914	-.352	-.0538	.215	
Robust p-value	.815	.398	.463	.95	.536	
Lower CI	-80.3	-.521	-1.25	-.737	-.544	
Upper CI	63.2	.207	.569	.692	1.04	
Bandwidth	.804	.611	.643	.874	.593	
Observations	57,841	42,401	44,886	63,990	41,021	

Notes: The table reports the estimated impact of an extra year of patent term on patent covariates. The patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. All estimates are from local quadratic regressions with a triangular kernel based on Calonico et al. (2014) optimal bandwidth selection with cluster-robust variance estimator at the application year times HJT technology subclass level. Reported are the RK point estimate, the robust p-value, the robust uncentered 95% C.I., the optimal bandwidth selection, and the effective number of observations that fall within the bandwidth. Panel A includes no controls, Panel B includes estimates the RK specification on the residual of a regression of the covariate on dummies for: HJT subclass times grant year, HJT subclass times filing year, count of maintenance fees paid for the patent.

Table 2.10: Regression kink discontinuity on outcome variables

	Baseline	95-00 cites	After expiry	1995- expiry	Only 1995
	(1)	(2)	(3)	(4)	(5)
Panel A: Without controls					
Treatment	-98.1	-.852	-81.1	-56.5	-31.2
Robust p-value	.0939	.118	.158	.204	.405
Lower CI	-251	-2.32	-226	-163	-102
Upper CI	19.7	.262	36.8	34.8	41.3
Bandwidth	.737	.666	.736	.844	.731
Observations	51,995	46,213	51,995	61,726	51,871
Panel B: With controls					
Treatment	-46.8	-.419	-51.5	-9.66	-9.26
Robust p-value	.265	.428	.228	.773	.852
Lower CI	-158	-1.66	-177	-107	-72
Upper CI	43.5	.705	42.1	79.7	59.5
Bandwidth	.801	.631	.75	.724	.712
Observations	57,841	43,845	52,999	51,384	49,910

Notes: The table reports the estimated impact of an extra year of patent term on patent citation outcomes. The patents considered were granted between 1980 and 1989 and still outstanding on June 8, 1995. All estimates are from local quadratic regressions with a triangular kernel based on Calonico et al. (2014) optimal bandwidth selection with cluster-robust variance estimator at the application year times HJT technology subclass level. Reported are the RK point estimate, the robust p-value, the robust uncentered 95% C.I., the optimal bandwidth selection, and the effective number of observations that fall within the bandwidth. Panel A includes no controls, Panel B includes estimates the RK specification on the residual of a regression of the outcome variable on controls that include: log of pre-TRIPS citation count plus one, originality, log of out-citations plus one, log of number of claims, and a set of dummies for HJT subclass times grant year, HJT subclass times filing year, maintenance fee counts, and foreign assignee.

Chapter 3

Durable Crises (joint with Nicolas Caramp and Pascual Restrepo)

Abstract

Consumer demand for durable goods is highly pro-cyclical: it falls substantially during recessions and rises sharply during booms. Using U.S. County Business Patterns data between 1988 and 2014, this paper studies how consumer durables amplify business cycle fluctuations. We show that employment in durable manufacturing industries is more cyclical than in other industries, and that this cyclicality is amplified in general equilibrium at the commuting zone level. We provide evidence of three mechanisms that generate amplification. First, employment changes propagate through input-output linkages, which amplify effects on local aggregate employment because industries co-locate. Second, the reduction of employment in durables negatively affects employment in non-tradable sectors, consistent with the existence of demand spillovers. Third, we find that workers do not completely reallocate to other less cyclical tradable industries. Our estimates suggest that consumer durables are responsible for up to 40% of aggregate employment volatility.

3.1 Introduction

The consumption of durable goods is highly cyclical. Relative to other goods, it contracts sharply during recessions and expands in booms (Bils and Klenow, 1998; Bils et al., 2013). In this paper we explore the role that consumer durables play in amplifying business cycles. We ask if the cyclicality of demand for durable goods contributes to aggregate employment reductions during recessions and corresponding raises during booms, or if on the contrary, these sectoral demand shocks are mitigated by reallocation to other industries and do not affect the aggregate employment level.

We first document that employment in industries that produce consumer durables is more cyclical than employment in other industries.¹ Using a measure of life expectancy of consumer goods adapted from Bils and Klenow (1998) and U.S. data from the County Business Patterns (CBP) covering the period from 1988 to 2014, we show that, relative to other manufacturing industries, employment in durable industries declines sharply during downturns but also recovers faster subsequently.² Quantitatively, when the slack in the U.S. labor market rises by 5 percentage points,³ as it did during the recent Great Recession, industry-level employment decreases by 2.265% more per additional year of expected life of the consumer good it produces. For the average durable industry, this translates into an additional 17% decline in employment with respect to industries that produce non-durables. These estimates imply that, when the slack in the U.S. labor market rises by 5 percentage points, employment in industries that produce consumer durables contracts by an additional 700 thousand jobs relative to other industries, or about half a percentage point of the labor force. Our results hold after we control for the secular decline in manufacturing and any potential trend that is specific to industries producing durable goods. The finding that employment in durable industries is particularly pro-cyclical is consistent with the view that business cycles may affect different industries heterogeneously.⁴

We next explore the implications of this volatility on aggregate employment. Our estimates of the cyclicity of employment in durable industries capture the differential effect across industries of the volatility of consumer durable consumption. These effects do not, however, correspond to the equilibrium impact on aggregate employment, which also encompasses indirect channels that could mitigate or amplify the impact of this sectoral shock on employment levels. One could expect the reduction in durable employment during recessions to have a small or no aggregate impact if workers quickly reallocate across sectors, in which case our estimates could simply reflect the reallocation of workers to less cyclical industries during downturns (Loun-

¹In what follows, we use interchangeably the terms durable industries and industries that produce durable consumer goods. These do not include industries producing materials used mainly as intermediary goods, such as primary metals, concrete and cement, or lumber and wood products except furniture.

²This result, also shown recently by Bils et al. (2013), complements the literature showing that durable goods have a more cyclical demand (Bils and Klenow, 1998); affect the volatility of exports (Engel and Wang, 2011); and affect the exposure to risk among firms that produce durables (Gomes et al., 2009).

³Measured as the difference between the unemployment rate and the natural rate of unemployment.

⁴See Abraham and Katz (1986).

gani and Rogerson, 1989). If indeed workers reallocate to less cyclical industries, the volatility of durable good consumption would not affect the aggregate behavior of employment (Baxter, 1996). Even if workers need to spend some time in unemployment to reallocate, the large gross flows of workers across industries imply that this reallocation could be achieved without any significant impact on aggregate employment (Pilossoph, 2012).

To estimate the equilibrium impact of the volatility in durable employment on the labor market we exploit differences in the industry composition of U.S. commuting zones. Because the bulk of the adjustment to labor demand shocks, and especially the reallocation of workers, takes place locally, commuting zones provide an ideal laboratory to investigate the aggregate effects from the decline in the demand for durables. Indeed, the evidence provided by Autor et al. (2013), Notowidigdo (2013), and Yagan (2014) suggests that the extent of workers' migration in response to labor market shocks is modest.⁵ Using CBP data covering all commuting zones in the contiguous U.S., we document that employment is more cyclical in commuting zones that host more durable industries. This finding holds even after we control for the secular decline in manufacturing and any potential trend specific to commuting zones hosting durable industries.

Quantitatively, when the aggregate slack in the economy rises by 5 percentage points, employment in a commuting zone that produces consumer durables that last for one additional year declines by 3.25% more relative to a region that produces no durables. The estimated impact of durables on a commuting zone is larger than what a shift-share projection based on the initial industry estimates would predict, which suggests that rather than mitigating the shock to durables, the equilibrium forces that operate at the commuting zone level amplify the effect of the decline in the demand for durable goods. Although they affect a single sector, the substantial albeit temporary changes in the demand for durables during recessions and booms have aggregate effects at the local labor market level, and impact national employment cyclicity as a result. These novel results are quantitatively significant. A back of the envelope calculation suggests that if overall U.S. employment behaved as it does in areas that do not produce durables, national employment would be 20% less volatile. Figure 3-1

⁵Though this evidence is in the context of more persistent shocks, we find it reasonable to expect even less migration in response to temporary shocks as the ones we study in this paper. Indeed, we analyze migration patterns in section 3.4 and find evidence of only little reallocation between commuting zones.

previews this result and plots the series for the cyclical component of the observed employment rate in the U.S. (in black circles) and a series of the counterfactual employment rate if no region produced durables (in blue hollow diamonds), resulting in decreased business cycle employment volatility.

We identify three mechanisms that explain why the volatility of demand for durables has a significant effect on aggregate employment. First, changes in the demand for consumer durables affect upstream industries that supply intermediate goods to durable goods producers. In line with this propagation through input-output linkages, we document that employment in upstream suppliers of durable industries is also highly cyclical, and so is employment in the commuting zones that host these suppliers. Quantitatively, when the aggregate slack in the economy rises by 5 percentage points, employment in a commuting zone with the average amount of linkages to durable goods declines by an additional 2.5% relative to a region with no linkages. A back of the envelope calculation suggests that if overall employment behaved as it does in areas that do not produce durables nor supply durable industries, U.S. employment would be 40% less volatile; input-output linkages double the contribution of consumer durables to the volatility of employment. The cyclical component for the counterfactual employment rate if the U.S. produced no durables nor supplied durable industries is also shown in Figure 3-1 in red squares.

In addition, we document that industries locate close to their suppliers (Ellison et al., 2010). Thus, input-output linkages amplify the impact on employment in local labor markets that host durable industries, and contribute to explaining why other industries in affected areas do not expand to pick up the slack in the labor market. Quantitatively, the fact that upstream firms co-locate close to producers of consumer durables explains one third of the impact of durable goods on local labor markets.⁶ The importance of amplification through input-output linkages is in line with recent evidence showing that industry shocks affect upstream industries (Acemoglu et al., 2015, 2016; Pierce and Schott, 2016; Carvalho et al., 2014; Barrot and Sauvagnat, 2016), and with the theoretical literature emphasizing how sectoral shocks could have aggregate effects because of input-output linkages (Acemoglu et al., 2012).

Second, we show that employment in non-tradable services is more volatile in areas that host durable industries. This volatility cannot be explained by input-output linkages, and suggests that lower consumption by laid-off workers may affect employment

⁶We take this co-location into account when aggregating results due to both direct effects of durables and upstream linkages.

in non-tradables through demand spillovers.⁷ This finding is in line with the empirical literature emphasizing how local declines in consumption affect employment in the non-tradable sector (Mian and Sufi, 2014), and with the literature emphasizing how demand externalities may amplify shocks when reallocation is imperfect (Beaudry et al., 2014). Quantitatively, the impact on non-tradable employment explains one fifth of the impact of durable industries on local labor market cyclical. However, whether durable cyclical results in negative spillovers on non-tradables at the national level as well will depend on the response of monetary and fiscal policy.⁸

Finally, we find little evidence of reallocation to non-durable tradable industries during crises. Abstracting from the impact of input-output linkages and the negative spillover on non-tradables,⁹ each additional year in the average expected life of consumer goods produced in a local labor market is associated with an extra decline of 1.5% in employment when the slack in the economy rises by 5 percentage points. This is smaller than our industry-level results predicts but still suggests that workers do not fully reallocate to other sectors and industries that are less cyclical. In line with this observation, we find no evidence that non-durable tradable industries that are not affected by input-output linkages,¹⁰ expand more in regions that host durable industries compared to others. These results are at odds with models in which a frictionless or rapid reallocation of workers mitigates the aggregate impact of a sectoral shock to the durable industry (Baxter, 1996; Pilossoph, 2012). Instead, one interpretation that may be consistent with the data is that, due to reallocation costs and the expectation that sectoral conditions may revert, workers do not reallocate but remain “rest unemployed” (Jovanovic, 1987; Hamilton, 1988; Gouge and King, 1997; Alvarez and Shimer, 2011).

Besides the literature already mentioned, our paper relates to the debate on the role of sectoral shocks in generating employment fluctuations. Lilien and Hall (1986) emphasizes that sectoral shocks generate business cycles, while in our case business cycles are amplified because some sectors are more sensitive to the cycle as argued by Abraham and Katz (1986). A literature going back to Schumpeter (1942) emphasizes that firms in declining sectors may be permanently liquidated during recessions, which

⁷For evidence on decreased consumption by unemployed workers, see Ganong and Jaeger (2016).

⁸Even if monetary or fiscal policy did fully offset the aggregate employment effects on non-tradables, our estimates still suggest major distributional impacts in terms of employment between commuting zones depending on their exposure to durables.

⁹That is, considering only direct effects of durable volatility on aggregate local employment.

¹⁰Nor are they affected by local demand spillovers, because of their tradable nature.

implies that permanent sectoral shifts may coincide with the onset of recessions.¹¹ Although manufacturing is on a secular decline in the U.S., we show that our results are robust to controlling in a number of ways for this decline and that our findings are specific to durable goods, rather than all manufacturing. Moreover, employment in durable industries and the commuting zones that host them rebounds in a pro-cyclical manner following a recession. Finally, our findings differ from Chodorow-Reich and Wieland (2016), who emphasize how secular reallocation, understood as the response of the economy to permanent sectoral shocks, may generate unemployment, especially during recessions.

The rest of the paper is structured as follows. Section 3.2 describes our data. Section 3.3 presents our evidence for industry employment and wages, which shows that there are large sectoral shocks that take place during recessions. Section 3.4 shows that these sectoral shocks have aggregate effects in U.S. commuting zones that host durable industries or their suppliers. Section 3.5 presents our investigation of mechanisms that generate amplification. Section 3.6 concludes by discussing the quantitative implications of our exercise and future avenues for research.

3.2 Data sources

We use yearly data from the County Business Patterns (CBP) between 1988 and 2014. CBP is an annual series covering U.S. employment during the week of March 12 and annual payroll data by county and industry. It covers all employment except self-employed individuals, employees of private households, railroad employees, agricultural production employees, and most government employees. In order to maintain a consistent panel of industries over our time period, we use the industry crosswalks in Autor et al. (2013) and aggregate our data to 479 industry codes. We restrict the analysis to the 48 states of the contiguous United States and aggregate the data to 722 commuting zones to study local labor markets. We supplement this data with information on within-U.S. net migration rates for each commuting zone from the Internal Revenue Service's Statistics of Income U.S. Population Migration Data, which records yearly migration flows between counties. In order to control for demographic covariates at the commuting zone-level, we use the 1990 Census. Finally, we use the

¹¹See also Davis and Haltiwanger (1990); Hall (1991); Caballero and Hammour (1994); Aghion and Saint-Paul (1998); Koenders and Rogerson (2005); Berger (2016); Jaimovich and Siu (2014); Restrepo (2015).

long-term NAIRU, observed unemployment, and potential and realized GDP series from the Federal Reserve Bank of St. Louis Economic Data to define two measures of economic slack. The first measure is defined as the difference between the observed national unemployment rate and the long-term natural rate of unemployment, whereas the second one is defined as the difference in log points between the potential and the realized GDP. Both measures are plotted in Figure 3-2, with NBER recessions periods shaded in grey.

We explore different measures for consumer durable exposure, all of which yield qualitatively similar results. In the main text we focus on a measure adapted from Bils and Klenow (1998), which defines for every industry the durability of the consumer goods it produces. If an industry does not produce consumer durables, it is assigned a zero, which allow us to focus on how changes in consumers' demand for durables affect employment. The average durability of consumer goods is 0.35 years (which takes into account that some industries do not produce consumer durables). Among industries that produce consumer durables, the average durability is 7.5 years.¹²

To measure the upstream and downstream exposure to industries that produce durables and investigate possible propagation through supply chain linkages, we use the 1992 input-output table for the U.S. economy from the Bureau of Economic Analysis.¹³ We compute for each of the 497 industries a measure of the share of their sales that are directly or indirectly used in the production of consumer durables. In particular, we use the matrix of cross-industry sales $S = \{s_{ij}\}$ (in shares) from industry j to i to compute its Leontief inverse $L^U = (I - S)^{-1} - I$. The row vector $L_i^U = (l_{i1}^U, l_{i2}^U, \dots, l_{iI}^U)$ indicates the upstream exposure of industry i to shocks in all the industries it directly or indirectly sells its products to. We compute the *upstream propagation* for a non-durable industry as

$$\text{Upstream Propagation}_i = \sum_j l_{ij}^U \cdot \text{Durability}_j.$$

This measure captures the extent of upstream propagation on non-consumer durable industries; it tells us the share of total production that is eventually used by industries to produce consumer durable goods, weighted by their respective durability. We also compute for each of the industries a measure of the share of their inputs

¹²These averages are weighted by the employed population in each industry, to mimic our specifications in section 3.3.

¹³The table is available at www.bea.gov/industry/io_benchmark.htm.

that are consumer durables or are produced using consumer durables. In particular, we use the matrix of cross-industry purchases $P = \{p_{ij}\}$ (in shares) from industry j to i to compute its Leontief inverse $L^D = (I - P)^{-1} - I$. The row vector $L_i^D = (l_{i1}^D, l_{i2}^D, \dots, l_{iI}^D)$ indicates the downstream exposure of industry i to shocks in all the industries it directly or indirectly purchases inputs from. We compute the *downstream propagation* for a non-durable industry as

$$\text{Downstream Propagation}_i = \sum_j l_{ij}^D \cdot \text{Durability}_j.$$

This measure captures the extent of downstream propagation on non-consumer durable industries; it tells us the share of consumer durable goods, weighted by their respective durability, that is needed to produce final goods in each industry. Figure 3-3 maps the geographic location of commuting zones that host durable industries and their upstream suppliers with extensive geographic variation across the US. The means of the 1990 Census covariates at the commuting zone level are shown in column (1) of Table 3.1.

3.3 Evidence from U.S. industries

We begin by exploring whether national-level employment in industries that produce consumer durables is more cyclical than in other industries over the period 1988-2014.¹⁴ To that end, we estimate the industry-level model:

$$\begin{aligned} \ln E_{it} = & \alpha_i + \delta_t + \beta^I \cdot \text{Slack}_t \times \text{Durability industry}_i \\ & + \gamma^I \cdot t \times \text{Durability industry}_i + \theta^I \cdot t \times \text{Manufacture}_i + \varepsilon_{it}, \end{aligned} \quad (3.1)$$

where $\ln E_{it}$ is the log of national employment in industry i in year t , Slack_t is our national-level measure of slack in the economy, $\text{Durability Industry}_i$ measures the durability of consumer goods produced by the industry, with the convention that industries that do not produce consumer goods are assigned a zero. Also, Manufacture_i is a dummy for manufacturing industries, and α_i and δ_t are a full set of industry and year fixed effects, respectively. ε_{it} is the error term, which we assume is independent across industries but may be serially correlated within each industry over time. When

¹⁴This period covers 3 recessions according to the NBER Business Cycle Dating Committee.

estimating equation (3.1) we weight observations by the employment in each industry in 1988 and report standard errors that are robust to heteroskedasticity and serial correlation within industries.¹⁵

The coefficient β^I that multiplies $\text{Slack}_t \times \text{Durability Industry}_i$ captures the additional cyclical of durable industries compared to nondurable ones.¹⁶ In the above model this effect is identified solely from cyclical fluctuations in employment, and does not confound the secular decline in manufacturing or any potential differential trend in durables. These two forces are accounted for by the trends $\gamma^I \cdot t \times \text{Durability Industry}_i$ and $\theta^I \cdot t \times \text{Manufacture}_i$.

Table 3.2 presents estimates of equation (3.1) using the industry-level data from 1988 to 2014 and covering the 479 industries defined in our data. We multiply our estimates by 100 so they can be interpreted in terms of log points. In panel A we use the unemployment measure of slack. In column (1) we present an estimate that excludes differential trends for durable and manufacturing industries. We estimate a statistically significant coefficient for β^I of -1.787, which suggests that durable industries are more cyclical. However, because our measure of slack rises sharply during the Great Recession, this estimate could confound cyclical movements in employment in durable industries with any secular trend affecting manufacturing or durables. To address this concern in column (2) we control for a time trend specific to durable industries. As expected, we find that employment in durable industries is on a statistically significant secular decline of 0.48% fewer jobs per year for every additional year of durability. Our estimate for the excess cyclical of durable industries now falls to a still highly statistically significant -0.453, which shows the importance of accurately controlling for industry trends. This point estimate suggests that, when the slack in the U.S. labor market rises by one percentage point, employment declines by 0.453% more for every additional year of durability among consumer goods produced by a given industry. Compared to non-durable industries, employment in the average durable industry thus falls by 3.4% more for every percentage point increase in the aggregate unemployment slack. Our estimates for β^I remain essentially unchanged

¹⁵In addition, in all of our models we control for a full set of year effects interacted with construction industry dummies, which leads to not considering construction in our industrial analysis. Though housing is an important durable good, we only focus on manufacturing durables as we want to abstract from the housing cycle and the impact of house prices on employment through net worth effects (Mian and Sufi, 2014). Nonetheless, our results are robust to including construction industry in the analysis.

¹⁶Notice that because recessions are measured through measures of positive slack, negative β^I will be associated with higher cyclical.

when we control for trends in the manufacturing sector in column (3) or industry-specific trends in column (5).¹⁷ Finally, in column (4) we explore the volatility of other manufacturing industries, i.e. non-consumer durable manufacturing industries. The coefficients from column (3) change very little, and we find that non-durable manufacturing is much less pro-cyclical than durable manufacturing. If anything, the sign of the statistically insignificant point estimate suggests non-durable manufacturing industries could exhibit less cyclical than the remaining industries.

In panel B we use the measure of slack defined by the output gap. We find similar results as above: when the slack in the U.S. labor market rises by one percentage point, employment declines by a statistically significant 0.278% more for every additional year of durability among consumer goods produced by a given industry. If we take into account Okun's law, which states that an increase in the unemployment rate of 1 percentage point is associated with an increase in the output gap of 2 percentage points, both sets of estimates yield similar quantitative implications. Using series of HP-filtered industry employment and real GDP between 1990 and 2011, Bils et al. (2013) find that the average durable industry is 1.79 times more volatile than GDP. Our estimates are of the same order or somewhat higher, and suggest that a one-point increase in GDP is associated with a 2.1% larger increase in the average durability.

Our results in this section support the view that the demand for durable goods is more cyclical and declines sharply during recessions. Our findings suggest that firms that produce consumer durables respond by reducing employment during downturns and expanding it during booms more than firms in other industries.¹⁸ However, the high cyclical of employment in the durable sector need not affect aggregate employment. As explained in the introduction, our industry-level estimates could reflect reallocation of workers between industries, as other industries that are less cyclical expand (or decline less) during downturns to absorb workers displaced from durable industries. In the rest of the paper we explore whether the decline in the demand for durable goods and the vast employment losses in this industry contribute to the observed cyclical of aggregate employment.

¹⁷The changes in the coefficients are smaller than the rounding level in the tables.

¹⁸In future work, we plan to analyze the intensive versus extensive margin of this adjustment at the establishment level. That is, do firms respond by closing down establishments, or by downsizing them?

3.4 Evidence from U.S. local labor markets

We now analyze the impact of the excess cyclicality of employment in durable industries on the local labor markets that host them. We estimate the following model using data for 722 commuting zones in the contiguous U.S. covering the 1988-2014 period:

$$\ln E_{ct} = \alpha_c + \delta_t + \beta^C \cdot \text{Slack}_t \times \text{Average durability}_{c1988} + \eta^C \cdot \text{Slack}_t \times \text{Manufacture}_{c1988} + \gamma^C \cdot t \times \text{Average durability}_{c1988} + \theta^C \cdot t \times \text{Manufacture}_{c1988} + \varepsilon_{ct}, \quad (3.2)$$

where $\ln E_{ct}$ is the log of the share of employment in commuting zone c in year t normalized by the the population in c at t ,¹⁹ Slack_t is again our measure of slack in the economy, and $\text{Average durability}_{c1988}$ is the average durability of consumer goods produced in the commuting zone in 1988, computed using the observed employment shares in that year and with the convention that industries that do not produce consumer durables are assigned a zero. Also, $\text{Manufacture}_{c1988}$ is the share of employment in manufacturing industries measured in 1988 for commuting zone c , α_c and δ_t are a full set of commuting zone and year fixed effects, respectively, and ε_{ct} is the error term, which we assume may be serially correlated over time for all commuting zones in a given state. We use the durability and manufacturing shares at the beginning of our sample (1988) instead of the contemporaneous values in order to reduce endogeneity concerns of the local productive structure.²⁰ When estimating equation (3.2) we weight observations by the employment in each commuting zone in 1988 and report standard errors that are robust to heteroskedasticity and serial correlation within states.²¹

Just as in the previous section, the coefficient β^C that multiplies $\text{Slack}_t \times \text{Average durability}_{c1988}$ captures the additional cyclicality of employment in areas that host durable industries. By including the two trends in the specification, we ensure that

¹⁹We also estimate all the specifications using non-normalized log of employment in commuting zone c at year t as an outcome variable and find qualitatively the same results. Moreover, later on in this section we investigate migration responses to durable cyclicality and find only economically small impacts.

²⁰The average durability and the manufacturing share at the commuting zone level are highly persistent over time. The correlation between values in 1988 and 2007 are around 0.8.

²¹In addition, in all of our models we control for a full set of year effects interacted with the share of workers employed in construction in 1988. As with our industry analysis, this allows us to abstract from the housing cycle and the impact of house prices on employment through net worth effects (Mian and Sufi, 2014). Our results are robust to foregoing these controls.

the secular decline in manufacturing or any potential trend in durables that could also affect employment in commuting zones is not confounded. The effect of interest is identified solely from cyclical fluctuations in employment.

Our approach exploits differences in the productive structure across commuting zones, in the extent to which they host consumer durable industries. Unlike our previous estimates, which compared relative changes in employment by industry, the impact of durables on the commuting zones that host them takes into account the possibility for reallocation, which could mitigate the aggregate effect of the shock to durables, or amplification mechanisms that could worsen the aggregate effects on employment. To the extent that most of the reallocation and adjustment to labor demand shocks takes place within a commuting zone, these estimates are informative about the equilibrium impact of the excess cyclicality of durables. To illustrate the value of contrasting our these two estimates consider the following scenarios.²² Suppose that workers displaced from durable industries reallocate immediately to other manufacturing jobs in the same commuting zone. While we would still observe relative changes in employment by industry ($\beta^I < 0$), we would not observe any impact on the overall employment level of the commuting zone ($\beta^C = 0$). If instead, workers displaced from durable industries do not reallocate but remain unemployed or out of the labor force, we would observe relative changes in employment by industry that match the impact of the overall employment level of the commuting zone ($\beta^C \approx \beta^I < 0$). Finally, suppose that because of demand externalities or other possible amplification mechanisms, the decline in demand for durables spills over to other industries in the same commuting zone. In this case, we could have a larger impact on the overall employment level of the commuting zone than in the durable industries ($\beta^C < \beta^I < 0$). These examples illustrate that the difference between the industry and commuting-zone estimates, β^C and β^I , reflects the extent to which reallocation, demand externalities and other general equilibrium effects that operate in a commuting zone mitigate or amplify the sectoral shock to durables.

Table 3.3 presents estimates of equation (3.2). As before, we multiply our estimates by 100 so they can be interpreted in terms of log points. In panel A we use the measure of slack defined by the difference between the national unemployment

²²Notice that the average durability of consumer durables is computed using employment shares by industry in each commuting zone and $d \ln (\sum_i E_{ict}) = \sum_i s_{ict} d \ln E_{ict}$, where s_{ict} is the share of employment in industry i within commuting zone c and time t . As a result, the magnitudes of the coefficients β^C and β^I are directly comparable.

rate and the natural unemployment rate (in percentage points). In column (1) we present an estimate that excludes the trends for commuting zones that host durable and manufacturing industries. We estimate a statistically significant coefficient for β^C of -1.977, which suggests that employment in commuting zones that host more durable industries behaves more cyclically than in other regions. However, because our measure of slack rises sharply during the Great Recession, this estimate could confound cyclical movements in employment in durable industries with any secular trend in manufacturing or durables. To address this concern, in column (2) we control for a time trend multiplied by the average durability of each commuting zone. As expected from the fact that employment in durable industries is on a secular decline, we find that employment in areas that host these industries is also on a decline over time.²³ Our estimate for the excess cyclicity of durable industries now falls to a still significant -1.013, which shows the importance of accurately controlling for secular trends. This point estimate suggests that the average commuting zone experiences a decline in employment that is 0.46% larger than if it hosted no durable industries when the U.S. labor market slack increases by one percentage point.

Our estimates for β^C remain largely unchanged²⁴ when we control for trends in the manufacturing sector in column (3), and remain largely unchanged when allowing areas that host manufacturing industries to have different cyclicity in column (4). Our estimates in columns (3) and (4) show that hosting non-durable manufacturing industries does not make a commuting zone more cyclical, and that once we control for the secular decline in manufacturing, commuting zones that host durables are not on a significant further secular decline. These findings suggest that our estimates for β^C do not confound the secular decline of employment in manufacturing or the possibility that this decline may concentrate during downturns.²⁵ Both results reassure us that our estimates for β^C are capturing the specific impact of the excess cyclicity of durables on local labor markets, and not trends that are common to all manufacturing industries.

One concern with our previous estimates is that areas that host durable industries may differ in unobserved dimensions from the rest of the U.S., or from other areas that also specialize in manufacturing but mostly produce nondurable goods. These differences could explain why these areas experience more pronounced recessions and

²³Chodorow-Reich and Wieland (2016) study the impact of these secular reallocations.

²⁴The changes are smaller than the table's rounding.

²⁵See Jaimovich and Siu (2014).

booms. However, we find no significant geographic bunching of durable industries in Figure 3-3, instead documenting extensive dispersion. Moreover, in most of our empirical specifications we also control explicitly for the share of manufacturing employment in the commuting zone. That is, for a given share of local manufacturing, we are exploiting variation in the differential industrial specialization in consumer durables. In order to study whether commuting zones with a larger share of durables differ from others along observable characteristics, we estimate a set of regression specifications with 1990 Census covariates at the commuting zone-level as dependent variable, and the average durability in 1988 as well as, depending on the specification, the manufacturing share of employment in 1988 as independent variables. The coefficients of interest on the average durability for each regression are shown in Table 3.1, as well as sample means of the covariates split by whether the durability of the commuting zone is below or above median in columns (2) and (3). We find that, although Census covariates vary significantly across commuting zones depending on their average durability, commuting zones with a similar manufacturing share but different within-mix of durability only vary significantly in their population size and share of college graduates, with other demographic characteristics not statistically different. They also differ significantly in their exposure to upstream linkages, which we explore in more detail in subsection 3.5.1.

Nonetheless, we allay these concerns further in columns (5) to (7). Although the distribution of durable industries across the contiguous U.S. shown in Figure 3-3 seems not to be concentrated geographically, we control for eight Census division dummies interacted with a full set of year effects in column (5). These dummies guarantee that we identify β^C only by comparing areas that host durables with other areas in the same division, which ensures that our estimates do not confound broad and unobserved regional differences. Our estimates are somewhat reduced, but we still find an economically and statistically meaningful estimate for β^C of -0.747. Besides the division dummies, in column (6) we include a series of covariates measured for each commuting zone using the 1990 Census interacted with a full set of year effects. We control for the log of population, the log of the workforce, the share of people in different age bins, the shares of people with high school and college degrees, and the shares of Blacks and Hispanics. Though differences in these demographic characteristics could make some commuting zones more sensitive to business cycles, we do not find that their inclusion affects our estimates, as we find a coefficient for β^C of -0.649.

Finally, in column (7) we include a full set of commuting zone trends, which control flexibly for unobserved heterogeneity and the possibility that areas that host durables are on a secular decline for reasons that are unrelated to the decline in manufacturing employment. Our estimates in column (7) suggest that, when the slack in the U.S. labor market rises by one percentage points, employment in the average commuting zone declines 0.23% more than if it hosted no durables.²⁶

Panel B presents our findings when we measure the slack in the U.S. economy using the output gap. Our point estimates in column (7) show that when the output gap rises by one percentage point, employment in the average commuting zone declines 0.18% more than if it produced no consumer durables. This is again in line with the results of Panel A, taking Okun's law into account.

Another potential concern with our estimates is that workers may respond to the decline in the demand for durables by moving to other commuting zones. If this were the case, our cross-sectional estimates would confound the (potential) reallocation of workers across commuting zones with a decline in employment. Although the existing evidence suggests that changes in migration are not an important response to local shocks,²⁷ we can test directly if the decline in employment documented above is driven by migration. Table 3.4 has the same structure as Table 3.3 but explores whether the net migration rate (inflow minus outflow) is more cyclical in areas that host durables. Our point estimates are quite small and precisely estimated. Moreover, once we account for differences across commuting zones and trends in columns (6) and (7), we do not find a significant effect of durables on the cyclicity of net migration. Quantitatively, when the slack in the U.S. labor market rises by one percentage point, the yearly net migration rate in the average commuting zone declines by only a statistically insignificant 0.01% more than if it did not host durables.²⁸

To gauge the economic significance of the estimates in this section we compute the counterfactual behavior of U.S. employment if it produced no durable goods.

²⁶The weighted average of the exposure to durables at the commuting zone is about 0.35.

²⁷Bartik (2017) finds large geographic moving costs that inhibit labor market adjustment. Likewise, Autor et al. (2013) and Notowidigdo (2013) find large persistence in local labor market shocks, consistent with low geographic adjustment.

²⁸Besides migration, there is an additional concern when interpreting our estimates of β^C as the equilibrium impact of the decline in the demand for durable goods. Because durables are traded across commuting zones, non-durable industries in other regions may benefit from the low price of durables and expand their employment. This reallocation of production through trade cannot be captured in our data and could lead to our estimates for β^C overstating the negative consequences of the decline in the demand for durables. However, in subsection 3.5.1 we find no evidence of benefits for downstream industries that use durables as inputs.

This counterfactual illustrates the behavior of employment in a scenario in which the demand for consumer durables were not more cyclical than the demand for other goods, or in which all durable consumer goods were imported from other countries. To compute our counterfactual, we multiply our estimate for β^C by the share of employment in durables in each commuting zone and subtract these employment losses or gains from the observed employment. This procedure gives us a series for employment in each commuting zone absent the effect of hosting durable industries:

$$E_{ct}^{nd} = \exp \left(\ln E_{ct} - \widehat{\beta}^C \cdot \text{Slack}_t \times \text{Average Durability}_{c1988} \right). \quad (3.3)$$

The observed and the counterfactual employment series coincide in areas that host no durables or when the aggregate slack in the economy is zero. We aggregate both series to compute their national average. For each average series we use the Holdrick-Prescott filter to compute the log deviations from its trend. Figure 3-1 plots the cyclical components of both series. As is evident from the figure, employment in the U.S. would be less cyclical if business cycles did not involve vast changes in the demand for durables. Quantitatively, the standard deviation of employment is 20% lower in the counterfactual scenario, which suggests that the high cyclicity of durable goods amplifies the impact of aggregate shocks by 20%.

3.5 Mechanisms that amplify the shock to durables

The evidence in the previous sections suggests that, when we look at local labor markets, the impact of the decline in the demand for durable goods is roughly of the same size as our industry estimates ($\beta^I \approx \beta^C < 0$), or even larger. We now explore three mechanisms that can explain why the sectoral shock to durables is not mitigated, and if anything is amplified, at the local labor market level. We first explore whether input-output linkages propagate the initial demand shock on durables across industries. We then analyze whether local demand spillovers impact non-tradable employment at the commuting zone level. Last, we analyze patterns of reallocation of employment from durable industries to other tradable industries.

3.5.1 Input-output linkages

We explore the possibility that the cyclical changes in the demand for durables propagate through input-output linkages. In particular, we expect shocks to the demand for durables to negatively affect upstream industries that supply inputs to durable good producers –what we refer to as upstream propagation. On the other hand, changes in the demand for durables have an ambiguous effect on downstream firms that use durables –what we refer to as downstream propagation. Though downstream industries may benefit from having access to cheaper durable goods –the low demand by consumers implies there are more durables to be used by downstream industries– the shock to durables may also push some upstream firms out of business, thus affecting downstream firms.²⁹

To assess the extent of upstream and downstream propagation at the industry level, we augment equation (3.1) as follows:

$$\begin{aligned} \ln E_{it} = & \alpha_i + \delta_t + \beta^I \cdot \text{Slack}_t \times \text{Durability Industry}_i \\ & + \beta_U^I \cdot \text{Slack}_t \times \text{Upstream Propagation}_i + \beta_D^I \cdot \text{Slack}_t \times \text{Downstream Propagation}_i \\ & + \gamma_i^I \cdot t + \varepsilon_{it}. \end{aligned} \tag{3.4}$$

Here, the terms $\beta_U^I \cdot \text{Slack}_t \times \text{Upstream Propagation}_i$ and $\beta_D^I \cdot \text{Slack}_t \times \text{Downstream Propagation}_i$ capture both potential sources of propagation. We also include industry trends $\gamma_i^I \cdot t$, specific to each of the groups of industries analyzed to isolate the effect of the secular decline in some industries from their cyclical responses.

Table 3.5 presents our industry-level estimates. Panel A uses the unemployment rate to measure slack while Panel B uses the output gap as a measure for slack. In column (1) we estimate the impact of upstream propagation controlling for industry trends. In panel A we estimate a statistically significant coefficient for upstream propagation of $\widehat{\beta}_I^U = -1.338$. This effect is large: our point estimate suggests that, when slack in the U.S. labor market rises by one percentage points, employment in the average non-durable industry declines by an additional 0.7% as a consequence of upstream propagation.³⁰ Meanwhile, employment in the average durable industry

²⁹For example, it could be that existing customer-supplier relationships are more productive or involve customized inputs. Barrot and Sauvagnat (2016) find that idiosyncratic supplier production shocks impose large output losses on their customers, especially when suppliers produce specific inputs.

³⁰The average non-durable industry has an upstream exposure of 0.5.

declines by an additional 3.9%. In column (2) we also estimate the impact of downstream propagation but find no evidence of downstream spillovers. In panel A we estimate a coefficient for downstream propagation of $\hat{\beta}_I^D = -0.249$. Though small and not significant, our imprecise estimates do not allow us to rule out large effects on downstream industries. However, since we are focusing on consumer durables, it is reasonable that there will not be significant downstream effects; few industries use consumer durables as intermediates while many act as suppliers to industries producing consumer durables.

In column (3) we explore if non-durable manufacturing industries are more cyclical once we take into account the upstream propagation of changes in the demand for durables. Our estimates show that, once we control for these sources of propagation, employment in non-durable manufacturing is not significantly less cyclical than employment in non-manufacturing industries. In contrast, employment in durable industries is considerably more cyclical. Column (4) goes one step further and restricts our analysis to manufacturing industries. It shows that once we account for upstream propagation, employment in durables is more cyclical than in other manufacturing industries.

We now explore how input-output linkages affect our commuting zone estimates. First, there is a high spatial correlation in the location of non-durable manufacturing firms and the durable industries they sell to in our data.³¹ Because we cannot identify the individual consumer-supplier relationships between production plants, we use the BEA national input/output tables to obtain average supply relationships between industries. We find that commuting zones with larger average durabilities also host more industry that supply consumer durables. Because of agglomeration gains, it is likely that these supply industries are also more connected to the local durable manufacturing, and implies that, through input-output linkages, the decline in employment can be amplified in commuting zones hosting large shares of durable industries. In addition, the upstream propagation that we document implies that the high cyclicity of durables may also affect commuting zones that do not host durable industries but that do host their suppliers.

To explore both mechanisms empirically at the commuting zone-level, we augment

³¹In line with Ellison et al. (2010).

equation (3.2) as follows:

$$\begin{aligned} \ln E_{ct} = & \alpha_c + \delta_t + \beta^C \cdot \text{Slack}_t \times \text{Average durability}_{c1988} \\ & + \beta_U^C \cdot \text{Slack}_t \times \text{Upstream Propagation}_{c1988} + \text{Trends} + \varepsilon_{ct}, \end{aligned} \quad (3.5)$$

where $\text{Upstream Propagation}_{c1988}$ is the average upstream exposure to durables among industries in commuting zone c measured using employment shares in 1988 to mitigate endogeneity concerns.

We present the results from this exercise in columns (1) and (2) of Table 3.6. In column (1) we report our baseline estimates from column (7) in Table 3.3 for comparison purposes. In column (2) we augment this regression by estimating whether upstream propagation makes employment more cyclical in commuting zones that host upstream suppliers to consumer durables. We find that in a commuting zone with the average amount of upstream linkages to durables (0.5), employment declines by 0.55% more than in a region with no upstream linkages when labor market slack rises by one percentage point. Moreover, our estimate for the impact of hosting durable industries falls from -0.649 to -0.416. This is in line with the fact that part of the effect of durables estimated in column (1) reflects propagation to upstream firms that locate close to durable good producers. Quantitatively, this co-location of suppliers close to their customers explains about a third of the effect of consumer durables on local employment found in section 3.4.

To gauge the economic significance of the estimates in this subsection we compute the counterfactual behavior of overall U.S. employment if it produced no durable goods and absent the upstream propagation. To compute our counterfactuals, we multiply our estimate for β^C by the share of employment in durables in each commuting zone and subtract these employment losses or gains from the observed employment. This procedure gives us a series for employment in each commuting zone absent the effect of hosting durable industries, as in equation (3.3). The observed and the counterfactual employment series coincide in areas that host no durables or when the aggregate slack in the economy is zero. We then compute an additional counterfactual in which we also subtract the role of upstream propagation:

$$\begin{aligned} E_{ct}^{ndu} = & \exp(\ln E_{ct} - \widehat{\beta}^C \cdot \text{Slack}_t \times \text{Average durability}_{c1988} \\ & - \widehat{\beta}_U^C \cdot \text{Slack}_t \times \text{Upstream Propagation}_{c1988}). \end{aligned} \quad (3.6)$$

We aggregate both counterfactual series to compute their national average. Figure 3-4 plots these counterfactual series (normalizing their level to 0 in 2007). As is evident from the figure, employment in the U.S. would be less cyclical if business cycles did not involve vast changes in the demand for durables. For each series we use the Holdrick-Prescott filter to compute log deviations from trend. Quantitatively, the cyclicalities of industries that produce durable goods explains 13% of aggregate employment cyclicalities. This is below our initial estimate because it does not take into account the propagation to suppliers that co-locate close to industries that produce consumer durables. Upstream propagation explains an additional 27% of the cyclicalities of aggregate employment, of which 7% is due to propagation in areas that also host industries that produce consumer durables, and the rest is due to upstream propagation to other regions.

3.5.2 Demand spillovers affecting non-tradables

Another potential source of propagation is through demand spillovers. If unemployed workers consume less,³² the demand for non-tradable goods produced and consumed locally in recessions may decline by more in areas more affected by the cyclicalities of durable industries. If that is the case, non-tradables will not expand in relative terms to pick up the extra slack in the labor market caused by downsizing in durable industries.

To assess the extent of negative spillovers on non-tradables, we estimate equation (3.5) but use the log of commuting zone share of employment in non-tradable services as our dependent variable.³³ We present the results from this exercise in columns (3) and (4) of Table 3.6. In column (3) we report our estimates without controlling for the upstream exposure of non-tradable industries, and in column (4) we include the upstream exposure of retail and service industries to durables. We find that employment in non-tradable services is highly cyclical in commuting zones that host durable industries. In the absence of demand spillovers, and because non-tradables do not include consumer durable goods, we would expect employment in these industries to expand during recessions in commuting zones that host durable industries relative to

³²Ganong and Jaeger (2016) find that spending on non-durable goods and services drops by 6% at the onset of unemployment and continues to fall during the unemployment period. When unemployment insurance is exhausted, spending falls by an additional 11%.

³³Non-tradable services include retail and other services, but exclude professional services, as in Autor and Dorn (2013).

other regions, as displaced workers laid off from durable industries reallocate to the non-tradable sector. That is, absent demand spillovers, employment in non-tradables should be comparatively less cyclical in commuting zones with larger durable industries. We thus attribute our opposite results to local demand spillovers.³⁴ We find a similar effect when we control for the upstream linkages of non-tradables to durables, and find no strong upstream linkages between these two types of industries.

In both columns we find that, when slack in the U.S. labor market rises by one percentage point, non-tradable employment in the average commuting zone declines by 0.234% more than if it did not host any durables. Quantitatively, the negative spillover on non-tradable industries explains about a fifth of the decline in overall local employment associated with durable goods.³⁵

Because of our empirical strategy, this spillover is a differential effect, measured by comparing different commuting zones with different shares of durable industries. Just as the high cyclical of durable employment at the industry level found in section 3.3 could be mitigated through reallocation, spillovers due to local aggregate demand externalities need not be present at the national level. Whether they are still present depends in part on the monetary and fiscal policy adjustments used to stimulate aggregate demand. However, if fiscal stimuli are in part geographically directed, they are likely to target differentially those commuting zones particularly affected by a recession.³⁶ That is, we find evidence of more pro-cyclical demand externalities in areas that are likely to already be benefiting more from counter-cyclical fiscal transfers.

On the other hand, a loosening of monetary policy as a result of worsening economic conditions may neutralize our negative spillover results in the aggregate. However, if nominal interest rates are already close to zero, as in the Great Recession, monetary policy may not have room to adjust. Moreover, if the demand channel is driven by complementarities between durable goods and non-tradable consumption, monetary policy will be ineffective against the increased cyclical of non-tradable consumption. Furthermore, even if the demand spillovers are neutralized in aggre-

³⁴The fall in demand for non-tradables driving these results can be due to local aggregate demand externalities. However, we cannot rule out the possibility that it is driven by income effects due to strong complementarities between durable good consumption and consumption of non-tradables.

³⁵The estimates need to be scaled down by the share of non-tradable employment at the commuting zone to obtain effects on total employment.

³⁶Automatic fiscal stabilizers in the form of unemployment insurance, for example, are likely to flow differentially more to areas with larger drops in employment.

gate, our evidence still suggests large distributional impacts of durable consumer goods across commuting zones. Summarizing, durable industries amplify the cyclicity of U.S. employment due to TFP or demand shocks by between 32% and 40%, depending on whether demand spillovers are present in the aggregate.

3.5.3 Lack of reallocation

Abstracting from the contribution of upstream propagation and demand externalities, we find that for every additional year of durability in the consumer goods produced in a commuting zone, its employment declines by 0.32% more when the slack in the economy rises by one percentage point. Starting from an estimate for β^C of -0.649, we have that one third of the effect is explained by the co-location of upstream suppliers in the same commuting zones that host durables and one fifth is explained by demand externalities. Our residual estimate with these adjustments is now below the comparable industry-level β^I , but still close.

These computations suggest that workers laid-off from durable industries are not reallocating to other tradable industries that are less cyclical and that are not affected by the upstream propagation. To test this idea we estimate equation (3.5) but use the log of employment in non-durable manufacturing industries as our dependent variable. We present the results from this exercise in columns (5) and (6) of Table 3.6. In column (5) we report our estimates without controlling for the upstream exposure of these nondurable industries, and in column (6) we include the upstream exposure of non-durable industries to durables at the commuting zone-level. In both columns, we estimate a positive impact of durable cyclicity on non-durable employment, but the estimated effect is economically small and not statistically significant. Notice that in order to compare the coefficient with the impact on aggregate employment, we need to scale it down by the average employment share of non-durable manufacturing, about 8.5% in 1988. In line with our previous findings, we estimate that one of the factors that keep these industries from expanding when durable industries shrink is their input-output linkages to durables and the upstream propagation these generate. However, even when we control for these linkages, we still find that even non-durable industries that are not affected by demand externalities nor input-output linkages fail to expand significantly when employment in durables, suppliers to durables, and non-tradable services shrinks. Our point estimate in column (6) suggests that unaffected nondurable industries only expand by about a tenth of the overall decline in

employment, and this effect is not statistically significant. The lack of reallocation explains why the decline in the demand for durables has a negative effect on overall local employment comparable to our industry estimates, even after we control for other sources of propagation.

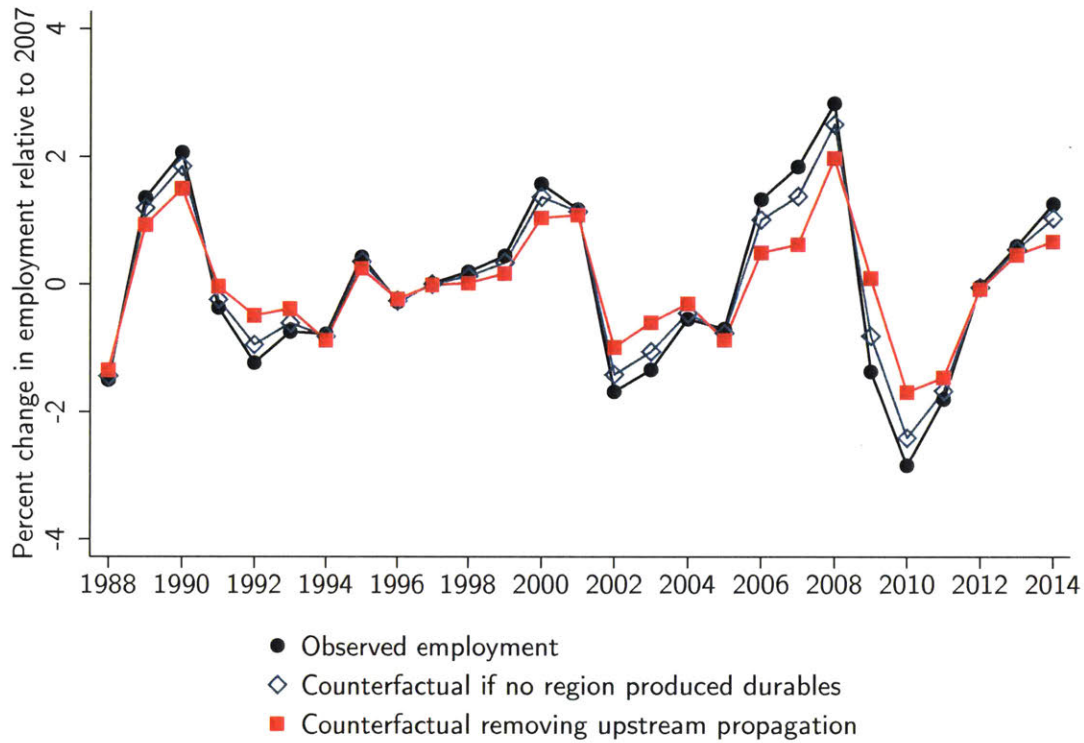
These findings raise the question of why workers are not fully reallocating from highly cyclical industries to less cyclical ones. One possibility is that, anticipating that the shocks to durables and their suppliers are only temporary, workers do not reallocate but remain “rest unemployed” until conditions improve. This is a hypothesis that we are currently investigating using other sources of data.

3.6 Quantitative implications and remarks

Consumer demand for durable goods is highly pro-cyclical. We find that this cyclical-ity has large implications for the volatility of U.S. aggregate employment. Consumer durables, and the propagation mechanisms highlighted above, explain between 32% and 40% of the business cycle volatility of aggregate employment. This effect can be decomposed into: a direct increase in volatility due to the cyclical-ity of durable industries of 10%, a subsequent effect through input-output linkages on suppliers to durable industries of 22%, and a spillover effect through aggregate demand external-ities of 8% that may or may not be present nationally depending on the room for adjustment in monetary policy.


Much work remains to be done. We are currently analyzing the effect of consumer durables on measures of payroll and establishment counts, to decompose the intensive versus extensive margin of adjustment by firms. Moreover, the lack of reallocation of workers to less cyclical tradable sectors is surprising. We find that reallocation forces only mitigate up to 10% of the cyclical-ity of consumer durable employment, and plan to investigate this further.

Figure 3-1: Cyclical component of U.S. employment and its counterfactual behavior if no U.S. region produced durables, nor supplied durable industries.



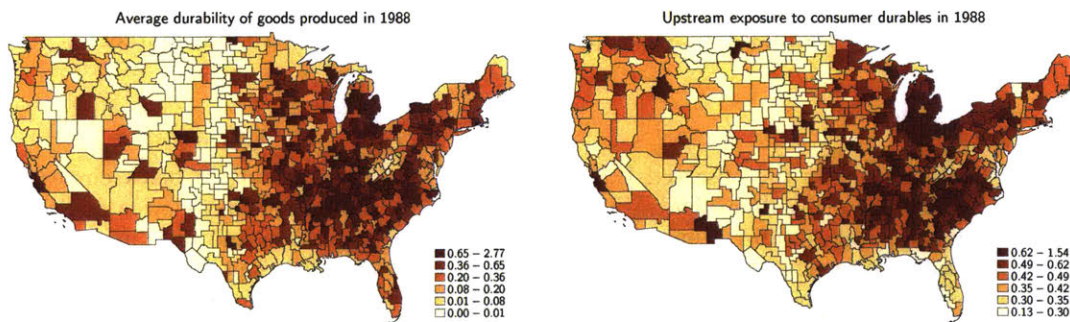
Notes: This figure shows the cyclical component of U.S. non-farm private employment (in black circles), the cyclical component of the counterfactual employment absent industries producing consumer durables (in blue hollow diamonds), and the cyclical component of the counterfactual employment absent industries producing consumer durables and their upstream suppliers (in red squares) between 1988 and 2014. Series are expressed as log deviations from their trends, computed with the Holdrick-Prescott filter. More details of the calculations of the counterfactuals in sections 3.4 and 3.5.1.

Figure 3-2: Unemployment slack and output gap in the U.S. over time.




Notes: This figure plots yearly values between 1988 and 2015 for the unemployment slack in dark blue and output gap in light red. The unemployment slack is defined as the difference between the observed U.S. unemployment rate and the long-term NAIRU, and the output gap is defined as $100 \times$ the difference between log of potential output and log of realized GDP. All data series are taken from the Federal Reserve Bank of St. Louis Economic Data. Also plotted are business cycle peaks (in dashed black) and troughs (in dashed blue) according to the NBER.

Figure 3-3: Average durability of goods and upstream exposure to consumer durables by commuting zone in 1988.



Notes: The map on the left shows the average durability of goods produced in each commuting zone in 1988, with the convention that industries that do not produce consumer durables have a durability of zero. The map on the right shows the upstream exposure of suppliers to consumer durables by commuting zone in 1988. This is calculated using the Leontief inverse of the 1992 BEA input-output table and the durability measure.

Figure 3-4: Employment and its counterfactual behavior if no U.S. region produced durables, there were no upstream propagation to industries in affected regions, or there were no upstream propagation to industries in other regions. All series are expressed in percent deviations from their 2007 level.



Notes: This figure shows the observed U.S. non-farm private employment (in black circles), the counterfactual employment absent industries producing consumer durables (in yellow triangles), the counterfactual employment absent industries producing consumer durables and their upstream suppliers in the same commuting zones (in blue hollow diamonds), and the counterfactual employment absent industries producing consumer durables and all their upstream suppliers (in red squares) between 1988 and 2014. More details of the calculations of the counterfactuals in sections 3.4 and 3.5.1.

Table 3.1: Descriptive statistics of covariates at the commuting zone level.

	Mean	Mean low durability	Mean high durability	Correlation durability	Partial correlation durability
	(1)	(2)	(3)	(4)	(5)
Share < 25	0.387	0.389	0.385	-0.005 (0.005)	-0.003 (0.005)
Share 25-44	0.297	0.297	0.298	-0.002 (0.003)	0.004 (0.003)
Share 45-64	0.189	0.186	0.191	0.008*** (0.003)	0.001 (0.003)
Share college	0.151	0.154	0.147	-0.019*** (0.006)	0.011** (0.004)
Share high school	0.553	0.534	0.572	0.055*** (0.014)	-0.010 (0.011)
Share hispanic	0.0583	0.0839	0.0327	-0.060** (0.024)	-0.007 (0.007)
Share black	0.0730	0.0405	0.105	0.047*** (0.018)	-0.020 (0.020)
Log population	11.48	10.89	12.08	0.803*** (0.251)	0.410* (0.238)
Upstream exposure	0.464	0.368	0.560	0.257*** (0.018)	0.065** (0.027)
Construction share	0.0465	0.0461	0.0469	-0.005** (0.003)	0.001 (0.003)
Commuting zones	722	361	361	722	722

Notes: Column (1) shows the mean of 1990 Census covariates at the commuting zone level, while columns (2) and (3) split the sample between below- and above-median average durability. In column (4), each cell shows the coefficient, and standard error in parentheses clustered at the state level, of a regression involving the 1990 census covariate on the average durability of each commuting zone in 1988. Cells in column (5) are defined as in column (4), but with each specification including a control for the share of manufacturing industry employment in each commuting zone in 1988. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 3.2: Estimates at the industry level of the different response of durable industries to economic fluctuations.

	INDUSTRY ESTIMATES FROM 1988 TO 2014				
	(1)	(2)	(3)	(4)	(5)
<i>Panel A. Slack measured using unemployment.</i>					
Industry that produces durables × Slack in year t	-1.787*** (0.166)	-0.453*** (0.093)	-0.453*** (0.093)	-0.436*** (0.095)	-0.453*** (0.095)
Industry that produces durables × Yearly trend		-0.478*** (0.060)	-0.140** (0.058)	-0.133** (0.057)	
Industry in manufacture × Yearly trend			-3.976*** (0.289)	-4.060*** (0.318)	
Manufacturing industry that produces nondurables × Slack in year t				0.843 (0.752)	
Observations	12906	12906	12906	12906	12906
Number of industries	479	479	479	479	479
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Panel B. Slack measured using output gap.</i>					
Industry that produces durables × Slack in year t	-1.048*** (0.096)	-0.278*** (0.062)	-0.278*** (0.062)	-0.274*** (0.063)	-0.278*** (0.063)
Industry that produces durables × Yearly trend		-0.498*** (0.059)	-0.160*** (0.056)	-0.158*** (0.056)	
Industry in manufacture × Yearly trend			-3.976*** (0.289)	-3.995*** (0.302)	
Manufacturing industry that produces nondurables × Slack in year t				0.200 (0.430)	
Observations	12906	12906	12906	12906	12906
Number of industries	479	479	479	479	470
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Unreported covariates:</i>					
Industry and year effects	✓	✓	✓	✓	✓
Construction × year effects	✓	✓	✓	✓	✓
Industry trends					✓

Notes: Dependent variable is log employment at the industry and year level. All specifications include a full set of industry and year fixed effects, as well as a set of construction dummy-times-year fixed effects. Column (5) includes industry-specific time trends. In panel A slack is measured as the observed U.S. unemployment rate minus the natural unemployment rate, whereas it is defined as the U.S. output gap in panel B. Robust standard errors in brackets are clustered at the industry level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 3.3: Estimates at the commuting zone level of the different response of regions that host durable industries to economic fluctuations.

	COMMUTING ZONE ESTIMATES FROM 1988 TO 2014						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. Slack measured using unemployment.</i>							
Baseline share of durables × Slack in year t	-1.977*** (0.416)	-1.013*** (0.218)	-1.013*** (0.218)	-0.870** (0.339)	-0.747*** (0.185)	-0.649*** (0.144)	-0.649*** (0.147)
Baseline share of durables × Yearly trend		-0.439*** (0.092)	-0.116 (0.121)	-0.131 (0.118)	-0.121 (0.108)	0.063 (0.081)	
Baseline share of manufacture × Slack in year t				-0.777 (1.154)			
Baseline share of manufacture × Yearly trend			-1.765*** (0.578)	-1.683*** (0.589)	-1.696** (0.670)	-2.798*** (0.447)	
Observations	19494	19494	19494	19494	19494	19494	19494
Number of regions	722	722	722	722	722	722	722
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Panel B. Slack measured using output gap.</i>							
Baseline share of durables × Slack in year t	-1.148*** (0.245)	-0.623*** (0.147)	-0.623*** (0.147)	-0.577*** (0.214)	-0.460*** (0.121)	-0.397*** (0.093)	-0.397*** (0.094)
Baseline share of durables × Yearly trend		-0.484*** (0.091)	-0.160 (0.120)	-0.165 (0.120)	-0.154 (0.107)	0.034 (0.079)	
Baseline share of manufacture × Slack in year t				-0.250 (0.646)			
Baseline share of manufacture × Yearly trend			-1.765*** (0.578)	-1.740*** (0.585)	-1.696** (0.670)	-2.798*** (0.447)	
Observations	19494	19494	19494	19494	19494	19494	19494
Number of regions	722	722	722	722	722	722	722
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Unreported covariates:</i>							
Commuting zone and year effects	✓	✓	✓	✓	✓	✓	✓
Share of construction × year effects		✓	✓	✓	✓	✓	✓
Census division × year effects					✓	✓	✓
Demographics × year effects						✓	✓
Commuting zone trends							✓

Notes: Dependent variable is log employment at the commuting zone and year level. All specifications include a full set of commuting zone and year fixed effects, and columns (2) to (7) include construction share-times-year fixed effects. Columns (5) to (7) include fixed effects for the eight Census divisions interacted with year dummies, columns (6) and (7) include commuting zone-level demographic controls interacted with year fixed effects, and column (7) adds controls for commuting zone-specific time trends. In panel A slack is measured as the observed U.S. unemployment rate minus the natural unemployment rate, whereas it is defined as the U.S. output gap in panel B. Robust standard errors in brackets are clustered at the state level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 3.4: Estimates at the commuting zone level of the different response of net migration in regions that host durable industries to economic fluctuations.

	NET MIGRATION RATE AT THE COMMUTING ZONE FROM 1988 TO 2014						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. Slack measured using unemployment.</i>							
Baseline share of durables × Slack in year t	-0.020 (0.040)	-0.082* (0.041)	-0.082* (0.041)	-0.074* (0.040)	-0.090*** (0.032)	-0.033 (0.029)	-0.033 (0.029)
Baseline share of durables × Yearly trend		-0.009 (0.007)	-0.007 (0.008)	-0.008 (0.010)	0.002 (0.007)	-0.002 (0.006)	
Baseline share of manufacture × Slack in year t				-0.042 (0.166)			
Baseline share of manufacture × Yearly trend			-0.010 (0.045)	-0.006 (0.055)	-0.012 (0.043)	0.020 (0.039)	
Observations	18050	18050	18050	18050	18050	18050	18050
Number of regions	722	722	722	722	722	722	722
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Panel B. Slack measured using output gap.</i>							
Baseline share of durables × Slack in year t	-0.013 (0.028)	-0.057* (0.029)	-0.057* (0.029)	-0.051* (0.027)	-0.062*** (0.022)	-0.026 (0.021)	-0.026 (0.022)
Baseline share of durables × Yearly trend		-0.013* (0.007)	-0.011 (0.009)	-0.012 (0.009)	-0.002 (0.007)	-0.003 (0.006)	
Baseline share of manufacture × Slack in year t				-0.031 (0.117)			
Baseline share of manufacture × Yearly trend			-0.010 (0.045)	-0.007 (0.050)	-0.012 (0.043)	0.020 (0.039)	
Observations	18050	18050	18050	18050	18050	18050	18050
Number of regions	722	722	722	722	722	722	722
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Unreported covariates:</i>							
Commuting zone and year effects	✓	✓	✓	✓	✓	✓	✓
Share of construction × year effects		✓	✓	✓	✓	✓	✓
Census division × year effects					✓	✓	✓
Demographics × year effects						✓	✓
Commuting zone trends							✓

Notes: Dependent variable is log of net migration (population inflows minus outflows) at the commuting zone and year level. All specifications include a full set of commuting zone and year fixed effects, and columns (2) to (7) include construction share-times-year fixed effects. Columns (5) to (7) include fixed effects for the eight Census divisions interacted with year dummies, columns (6) and (7) include commuting zone-level demographic controls interacted with year fixed effects, and column (7) adds controls for commuting zone-specific time trends. In panel A slack is measured as the observed U.S. unemployment rate minus the natural unemployment rate, whereas it is defined as the U.S. output gap in panel B. Robust standard errors in brackets are clustered at the state level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 3.5: Estimates at the industry level of the different response of durable industries and their suppliers to economic fluctuations.

	INDUSTRY ESTIMATES FROM 1988 TO 2014			
	(1)	(2)	(3)	(4)
<i>Panel A. Slack measured using unemployment.</i>				
Industry that produces durables × Slack in year t	-0.521*** (0.096)	-0.524*** (0.101)	-0.521*** (0.102)	-0.490*** (0.104)
Upstream propagation of durables × Slack in year t	-1.338*** (0.177)	-1.333*** (0.186)	-1.417*** (0.262)	-1.109*** (0.193)
Downstream propagation of durables × Slack in year t		-0.249 (1.823)	-0.219 (1.820)	
Manufacturing industry that produces nondurables × Slack in year t			0.367 (0.632)	
Observations	12906	12906	12906	10584
Number of industries	479	479	479	392
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014
<i>Panel B. Slack measured using output gap.</i>				
Industry that produces durables × Slack in year t	-0.320*** (0.064)	-0.324*** (0.067)	-0.323*** (0.068)	-0.302*** (0.068)
Upstream propagation of durables × Slack in year t	-0.832*** (0.121)	-0.826*** (0.127)	-0.869*** (0.168)	-0.686*** (0.122)
Downstream propagation of durables × Slack in year t		-0.307 (1.148)	-0.291 (1.151)	
Manufacturing industry that produces nondurables × Slack in year t			0.187 (0.402)	
Observations	12906	12906	12906	10584
Number of industries	479	479	479	392
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014
<i>Unreported covariates and sample:</i>				
Industry and year effects	✓	✓	✓	✓
Construction × year effects	✓	✓	✓	✓
Industry trends	✓	✓	✓	✓
Only manufacturing				✓

Notes: Dependent variable is log employment at the industry and year level. All specifications include a full set of industry and year fixed effects, a set of construction dummies-times-year fixed effects, and industry-specific time trends. Column(4) restricts the analysis to manufacturing industries. In panel A slack is measured as the observed U.S. unemployment rate minus the natural unemployment rate, whereas it is defined as the U.S. output gap in panel B. Robust standard errors in brackets are clustered at the industry level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table 3.6: Estimates at the commuting zone level of the negative spillovers created by the decline in employment in the durable industry on other sectors.

	COMMUTING ZONE ESTIMATES FROM 1988 TO 2014					
	TOTAL EMPLOYMENT		NON-TRADABLE SERVICES		NON-DURABLE MANUFACTURE	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Slack measured using unemployment.</i>						
Baseline share of durables × Slack in year t	-0.649*** (0.147)	-0.416*** (0.131)	-0.519* (0.265)	-0.520* (0.266)	0.365 (0.773)	0.676 (0.766)
Upstream propagation for all industries × Slack in year t		-1.089*** (0.307)				
Upstream propagation for non-durables × Slack in year t						-1.679*** (0.494)
Upstream propagation for retail and services × Slack in year t				-0.489 (2.761)		
Observations	19494	19494	19494	19494	19494	19467
Number of regions	722	722	722	722	720	720
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Panel B. Slack measured using output gap.</i>						
Baseline share of durables × Slack in year t	-0.397*** (0.094)	-0.280*** (0.088)	-0.387** (0.172)	-0.387** (0.173)	0.427 (0.533)	0.619 (0.533)
Upstream propagation for all industries × Slack in year t		-0.547*** (0.202)				
Upstream propagation for non-durables × Slack in year t						-1.033*** (0.356)
Upstream propagation for retail and services × Slack in year t				0.555 (1.781)		
Observations	19494	19494	19494	19494	19494	19467
Number of regions	722	722	722	722	720	720
Years in panel	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014	1988-2014
<i>Unreported covariates:</i>						
Commuting zone and year effects	✓	✓	✓	✓	✓	✓
Share of construction × year effects	✓	✓	✓	✓	✓	✓
Census division × year effects	✓	✓	✓	✓	✓	✓
Demographics × year effects	✓	✓	✓	✓	✓	✓
Commuting zone trends	✓	✓	✓	✓	✓	✓

Notes: Dependent variable is log of total employment at the commuting zone and year level in columns (1) and (2), log of employment in non-tradable services at the commuting zone and year level in columns (3) and (4), and log of employment in non-durable manufacturing industries in columns (5) and (6). All specifications include a full set of commuting zone and year fixed effects, a set of construction share-times-year fixed effects, census division as well as commuting zone demographic controls interacted with year dummies, and commuting zone-specific time trends. In panel A slack is measured as the observed U.S. unemployment rate minus the natural unemployment rate, whereas it is defined as the U.S. output gap in panel B. Robust standard errors in brackets are clustered at the state level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Appendix A

Appendices for Chapter 1

A.1 Variable construction

In this section I discuss the construction of variables and the data used in more detail. I follow the past literature closely in constructing the static spillover variables, in order to keep my estimates comparable.

A.1.1 Static spillovers

In subsection 1.2.4 I explain how to construct the *CitSpill* and *SicSpill* spillover measures. I discuss the construction of these variables in more detail here, and define other static spillover measures introduced by Schnitzer and Watzinger (2015), Jaffe (1986) and Bloom et al. (2013). I start with the basic static spillover measures and then define more sophisticated extensions.

All the static spillover measures share the same underlying construction logic. They correspond to a pool, or weighted sum of other firms' contemporaneous R&D stocks. That is, the spillovers for firm i at time t is defined as $Spill_{it} = \sum_{j \neq i} \omega_{ij} G_{jt}$, where G_{jt} is the R&D stock of firm j at time t and ω_{ij} is a measure of *proximity* between firms i and j representing the likelihood that firm j 's R&D activity spills over unto firm i 's innovative activity. The difference between each spillover measure lies thus in the weighting or proximity matrix used to construct it.

One measure used to quantify static knowledge spillovers is *CitSpill*, in which the proximity between firm i and j is defined according to equation 1.7. The idea for this proximity matrix is that patents citing other patents directly build upon them. As a result, the degree to which firm i 's innovation is inspired by firm j 's will increase

with its share of citations to j 's patents. This follows Azoulay et al. (2015) and Schnitzer and Watzinger (2015). This matrix is however different from Schnitzer and Watzinger (2015)'s citation proximity matrix in that I normalize by the total amount of outcitations rather than by the patent count. In line with the construction of the dynamic knowledge proximity matrix, I believe it is important to use citation shares per patent rather than citation counts in order to keep the innovation production function with constant returns to scale in terms of its cumulative nature.

In order to build *CitSpill*, I use the linked patent-firm data and consider all patents filed between 1984 and 2001. The regression analysis is carried out between 1990 and 2001, and I restrict the analysis with patent data to end in 2001 to avoid attrition concerns. In order to set the starting date, I consider the average time it takes for a patent to be granted after application (2.14 years in the raw patent data) and leave up to 3 years for citations to accrue. I thus start considering patent citations since 1984. The analysis is however robust to modifications in the starting date.

For this citation proximity measure, I also construct an instrument to account for the endogeneity in the citation decision. For the 340 firms in my sample, I calculate the average citation propensity between 1976 and 1984. I then regress the citation proximity between 1984 and 2001 on the 1976-1984 propensity together with citing and cited firm fixed effects. I use the resulting predicted proximity *PredCit* as instrument for the actual proximity *CIT*, and combine it with the tax-predicted R&D to construct an instrument for the static knowledge spillovers *CitSpill*.

I also use the technological proximity in order to build static knowledge spillovers *TechSpill*. It uses the positioning of each firm in the "technology space", and calculates the correlation between firms' position. The position of a given firm in terms of technology is given by a 426x1 vector T indicating the share of patents filed by that firm in each of 426 technology classes between 1970 and 1999. The technological proximity between firms i and j is then defined as

$$TECH_{ij} = \frac{T_i T_j'}{(T_i T_i')^{\frac{1}{2}} (T_j T_j')^{\frac{1}{2}}}. \quad (\text{A.1})$$

This proximity measure was first introduced in Jaffe (1986), and defines a symmetric proximity measure. Companies active in very different technologies will have a proximity close to 0, while firms working in similar fields will see their proximity increase.

In terms of static business stealing spillovers, I use the measure proposed by Bloom et al. (2013). Similar to Jaffe (1986)'s technological proximity, they use the position of a given firm in the product market space to define its proximity to other firms. That position is defined by the 597x1 vector S_i indicating the average share of sales of firm i in each of four digit SIC-code industry. The breakdown of sales by four digit industry is available for firms in the Compustat Business Segments database from 1993 onward, and so the average over 1993-2001 is used. Once the product market position vectors are defined, the proximity between firm i and j is calculated as in equation 1.9 as the uncentered correlation between the vectors. Or in other words, the cosine of the angle between the vectors.

One drawback of the *Tech* and *Sic* distance measures is that they assume that spillovers only occur within narrowly defined technological classes or industries, and they rule out spillovers between classes or industries. In order to relax this assumption, Bloom et al. (2013) introduce a Mahalanobis distance. This distance takes into account how often different technology classes (respectively, industries) coincide within a given firm's technological (product market) vector in order to define distances between classes (industries). These distances are then used to further weight the correlation between technological (product market) vectors. A detailed derivation can be found in Bloom et al. (2013). In order to ensure comparability of my estimates with previous literature, I use the *Tech*, *Sic*, *MahTech* and *MahSic* proximity matrices calculated by Bloom et al. (2013).¹

A final extension to Jaffe (1986)'s technological proximity measure is proposed by Schnitzer and Watzinger (2015) and He (2015). It is similar to the Mahalanobis distance defined above in relaxing the assumption of spillovers occurring only within technology classes. It relaxes this assumption by defining an asymmetric, non-binary distance measure based on the citation flow between classes. That is, the proximity between technological classes A and B is defined as

$$m_{AB} = \frac{\#Citations_{A \rightarrow B}}{\#Outcitations_A}. \quad (\text{A.2})$$

The (426, 426) weighting matrix M is then used the same way as the Mahalanobis distance to construct an augmented proximity *CitTech*. That is, defining a normalized (426, N) matrix $\tilde{T} = [T'_1/(T_1T'_1)^{1/2}, T'_2/(T_2T'_2)^{1/2}, \dots, T'_N/(T_NT'_N)^{1/2}]$, we have

¹I would like to thank the authors for sharing their data.

the (N, N) matrix

$$CitTech = \tilde{T}' M \tilde{T}. \quad (A.3)$$

Figures A-1, A-2 and A-3 plot the pairwise values of three different proximity measures (Sic , Tec , and Cit) for the firms in my sample. They show that there is ample variation in all three spaces to separate different spillover types. They also show that firms are less likely to be classified as close in the citation space. Within directed pairs of firms² with a strictly positive TEC proximity, about 90% have $Cit = 0$. As a result, Cit is a much more restricted measure of the pool of R&D likely to spill over.

A.1.2 Dynamic spillovers

The construction of the dynamic spillover measure $DynSpill$ is discussed in subsection 1.2.3. The main differences between the dynamic measure and the static ones pertain to the increased complexity of the proximity matrix, and the use of R&D intensity instead of R&D stocks.³ The proximity matrix for the dynamic measure is extensively discussed in subsection 1.2.3. Here, I discuss the convenience of combining the proximity weights with R&D intensity to construct spillover metrics.

Static spillover measures represent complementarities or business stealing effects between the R&D efforts of different firms. In the case of static knowledge spillovers for example, the spillovers accrue with the direct inclusion of firm j 's R&D effort in firm i idea production function, as marked in the system of equations 1.1. However, dynamic spillovers of R&D accrue through the initial production of knowledge. That is, past R&D creates ideas that get codified as patents. This codified knowledge then is diffused to subsequent innovators, who use it to produce new ideas. Therefore, in order to measure dynamic R&D spillovers, they have to originate through initial innovation production. That is, larger R&D efforts in the past should result in the production of higher-quality ideas.⁴ Building upon these originally higher-quality ideas then results in increased subsequent idea production (either in terms of quantity or quality).

Thus, in order for R&D to result in dynamic spillovers, it must first have an effect

² Cit is asymmetric, whereas Tec is symmetric.

³R&D intensity is R&D stock normalized by a measure of firm size, usually total assets.

⁴Where quality of ideas is understood as "the magnitude of inventive output associated with them" (Griliches, 1990).

on initial idea production. In Table A.1, I estimate OLS regressions of innovative output on innovation production function inputs such as R&D, at the firm-year level. I find that R&D stock is positively associated to the average citation count per patent granted only when controlling also for the size of the firm. That is, higher R&D intensive firm-years are associated to higher quality of innovative output, whereas the relationship with R&D stock is tenuous when not controlling for firm size. Because of this relation, and dynamic spillovers conceptually accruing through the initial production of higher-quality innovation, I use R&D intensity rather than stock to construct spillover measures.

Moreover, in unreported regressions I also find that other normalized measures of R&D spending, such as R&D stock by patent filed, do not correlate with average patent quality after controlling for firm and year fixed effects. This could be consistent with research-intensive firms working on a number of research lines, with the amount of lines in each firm increasing with firm size. Within each research line, higher R&D effort leads to both more and better innovations. As a result, higher R&D intensity in terms of R&D over assets (firm size) will lead to higher R&D spending within each research line, and thus to higher innovation output. Meanwhile, measures such as R&D over patents will not be indicative of the quantity and quality of innovative output.

A.2 Instruments

In this section, I discuss the instrumental variable identification strategy, with particular emphasis on the construction of the instrumental variables, and the necessary assumptions for the exclusion restriction to hold.

A.2.1 Tax instrument

In order to instrument endogenous R&D spending decisions, I use tax-induced shocks to the supply-side user cost of corporate R&D. This discussion follows Bloom et al. (2013) closely. The Hall-Jorgenson formula for the user cost of R&D induced by corporate income taxes and tax credits is

$$\rho_{it} = \frac{1 - D_{it}}{1 - \tau_{s_{it}}} \left[I_t + \delta - \frac{\Delta p_t}{p_{t-1}} \right], \quad (\text{A.4})$$

where D_{it} is the discounted value of R&D tax credits, $\tau_{s_i t}$ is the corporate income tax rate that includes a federal as well as a state s_i component, I_t is the real interest rate, δ is the depreciation rate of R&D capital, and p_t is the R&D price. Since the terms in the bracket do not vary at the firm level, I focus on the first ratio, or the tax component of the cost of R&D capital $\rho_{it}^T = \frac{1-D_{it}}{1-\tau_{s_i t}}$.

The variation in ρ^T is decomposed in two components. The first one, ρ^F , uses federal variation in corporate income taxes and R&D tax credits. The second, ρ^S , uses state-level variation in tax incentives. The first component is constructed following Hall (1992) and uses time series variation in the corporate income tax as well as in the "Research and Experimentation Tax Credit". This tax credit was introduced in 1981 and, after expiring in 1985, has been extended fifteen times. It provides a tax credit for 20% of qualified R&D expenses above a firm-specific base.⁵ The definition of this base has varied across time. Between 1981 and 1990, the base was the maximum of the previous three years' R&D expenses (with a minimum of 50% of the current R&D level). From 1990 onward, the base was set as the average of the R&D to sales ratio between 1984 and 1988 (with a maximum of 16%) multiplied by current sales. Firms incorporating after 1983, or with less than three years of qualified R&D expenditures and revenue between 1984 and 1988, have a 3% base ratio for the first five years and modified subsequently. On top of the firm-specific R&D bases, the tax credit interacts with firm-specific income taxes. If the credit exceeds taxes, it must be carried forward. With discounting, this reduces its value for firms with small tax bills. The tax credit rates and bases therefore also interact with corporate income tax rates and deductions, leading to firm-specific variation.

The second component uses state-level tax credits, recapture rules, and corporate income tax rates. State-level R&D tax credits exhibit large variation both in the time-series and in the cross-section, with the first tax credit being introduced by Minnesota in 1982 and 28 other states introducing credits by 2001. Tax credit rates vary from 2.5% in Minnesota in 1992 to 20% in Arizona and Hawaii in 2000, and there are multiple changes in credit rates per state (e.g., California changed their rate four times between 1986 and 2000). Firms will be differentially affected by state-level R&D tax incentives depending on the state in which their R&D activity is located, as state-level credits are meant for R&D carried out in a given state, and can be used to offset state corporation taxes. Since state tax liabilities on total firm profits

⁵The rate was 25% between 1981 and 1985, and 0 in 1995 when the tax credit lapsed.

are apportioned using combinations of the distribution of firm sales, employment and property, any firm with R&D labs in a given state is also likely to be liable for state income taxes. Therefore, patenting firms' inventor location appears to provide a good proxy for eligibility for state tax credits. Using the address of each patent's inventor, each patent is allocated to a state. θ_{its} , the exposure of firm i at time t to state s , is then the 10-year moving average of the yearly share of patents filed by firm i in state s . The state component of the tax price of R&D is then $\rho_{it}^S = \sum_s \theta_{its} \rho_{st}^S$.

The endogenous R&D expenses are logged and projected on the two components of the tax instrument for the 340 firms in the final sample between 1980 and 2001, and the results are shown in Table 1.3. Column (1) shows the basic results, column (2) adds year fixed effects, column (3) additionally includes firm fixed effects, and column (4) adds industry-times-year fixed effects. The instruments have considerable power in all specifications, with all the F -statistics above 28. Specification (3) is used to construct a predicted R&D variable which is stocked into RDS_{it}^{Tax} using perpetual inventory methods as in the true R&D case. This predicted R&D stock is then used to instrument for endogenous R&D stocks and to construct the spillover instrumental variables.

The predictive power of the tax-induced supply-side cost of R&D is strongly rooted in the empirical literature. Surveys of the literature such as Hall and Van Reenen (2000), or more recently Becker (2015), find that elasticities of R&D to their tax-induced price are estimated to be broadly around unity. Wilson (2009), from which the state tax data originates, estimates in-state elasticity of aggregate R&D to be between -1.2 and -2.2, whereas the elasticity of in-state R&D to neighboring states' tax-induced prices is even larger in magnitude, reaching 4.4 in the long-run. In more recent work using a regression-discontinuity design in the UK, Dechezleprêtre et al. (2016) find an elasticity of about -2.6. The estimates in Table 1.3 are therefore well within the range found in the prior literature.

One concern is that changes in R&D tax credits might be endogenous, with states responding to falls in R&D levels or economic activity by increasing tax credits. Alleviating this concern are the four following checks. First, I experiment with lagging and leading the tax prices one period. I find that lagged tax prices affect R&D expenses in qualitatively similar ways as current tax prices, but lead tax prices lose predictive power, with the coefficient on federal tax prices becoming insignificant. This is in-

dicative of R&D expenses responding to taxes and not vice-versa. Second, the state tax prices exhibit significant variation both in the time-series and in the cross-section. States have become more generous over time in terms of their tax credits, with Minnesota introducing the first credit in 1982 and 28 other state introducing a credit by 2001. The cross-sectional variation in rates is large relative to their mean and average growth rate, with state tax credits ranging between 2.5% and 20%. Moreover, there is also within-state variation in rates after introduction, with for example California changing their rate three additional times after its introduction in 1986. Third, the level and timing of state tax credits do not seem to be correlated with state-level observables after controlling for state and year fixed effects. Chirinko and Wilson (2008, 2013) find that aggregate variables such as the federal tax credit rate have explanatory power on state corporate tax credits,⁶ but not local economic or political observables. Likewise, Bloom et al. (2013) do not find any predictive power of lagged state-level R&D expenditures or GDP per capita on current tax credits. This may reflect the delays in passing regulation changes through state legislatures, and the fact that the costs of corporate tax credits are largely small so that their adoption and levels does not seem driven by budget concerns.

A.2.2 Network instrument

I also take into account the endogenous network formation, that may be driven by better researchers both producing knowledge more efficiently and identifying and citing higher-quality prior art. As a result, both the dynamic spillover measure and the static citation spillover measure will be endogenous. I instrument the current citation network between firm-year nodes using the past network structure. In particular, I use the *past* 1976-1984 network to predict the *subsequent* 1987-2001 network. Within the 715 firms originally present in the full network for which the business stealing proximity measures can also be constructed, 340 of them are found to be originating citations (and thus at the receiving end of knowledge flows) in both past and subsequent network. In terms of choosing the cut-off years 1984 and 1987, three forces are at play. First, I want to include as many years as possible in the subsequent network in order to carry out the empirical analysis on a panel with a long time dimension, as well as leave some time for patents to accumulate citations. Second, I want to include

⁶The possible influence of aggregate variables or shocks is controlled for in the regressions using the year fixed effects, and the more flexible industry-times-year dummies.

as many years as possible in the past network to obtain a more precise estimate of the citation propensity between firms to predict the subsequent network structure. Third, I want to leave a large enough gap between the pre-period network and the current network, to ensure that the exclusion restriction holds.

The logic for using this past network structure as an instrument is as follows. The structure of the subsequent network, in terms of citations between firms i and j and the gap between years $t - t'$, that is predicted by the past network citation propensity will most likely be a reflexion of an underlying firm-specific culture or strategy in terms of its absorptive capacity of outside knowledge. Because it is predicted by citation patterns up to 25 years in the past, it is likely to reflect constancy in institutional, structural, or cultural aspects affecting the absorptive capacity of firms. For a particular example, Lim (2009)'s analysis of the copper interconnect technology for semiconductor chips exposes clearly that the firms involved in the development and diffusion of this technology had very different and clearly-defined strategies in place to produce and absorb knowledge, involving different ways to manage internal R&D and external links. IBM's *disciplinary* strategy involved carrying out very early stage exploratory R&D in-house by hiring scientists and developing ties with the academic community. Meanwhile, Motorola's *domain specific* strategy was to seek solutions to specific technical problems with focused R&D, including by funding external R&D in specific areas. As a result of these firm-specific strategies, Motorola will have a tendency to build on technology created by IBM, albeit with a lag, and its citation patterns incorporate that information. In order to account for the heterogeneity across firms, that is the firm-specific time-invariant culture or strategy, I use firm fixed effects in all specifications. Conditional on firm fixed effects, I therefore expect the exclusion restriction to hold: the past network instrument will be uncorrelated with the error term.

A.3 Analytical framework

In the first subsection of this appendix, I examine the consequences of a model of competition and innovation production between firms, in order to extract predictions for the empirical results on the R&D spillovers. In the second subsection, I derive the expressions for the private and the social returns to R&D as a function of the parameters of interest in the empirical section. Finally, I derive the implications of

the private and social returns regarding the optimal provision of R&D relative to the decentralized level.

A.3.1 Cumulative innovation production model

I study the implications for a given firm of an innovation production function in which the production of new ideas results from three inputs, R&D effort by the firm itself, R&D effort by technological neighbors and past innovation on which to build. I consider a dynamic game of two periods, each consisting of two stages. In each period, firms decide on the level of their R&D effort in order to produce knowledge in the first stage. In the second stage, with a given knowledge stock, firms compete in the product market and I assume a pure strategy Nash equilibrium exists. Firms decide on a variable x in the second stage, where x can represent prices or quantities, conditional on their knowledge level k . In the first stage, k is produced using own and spillover R&D r , as well as existing foundational innovation stock S . In the second period, the previous two-stage game is repeated with an updated stock of foundational innovation on which to build S' . In order to keep the model simple and tractable, I consider three firms labeled 0, τ and m : firms 0 and τ are technological neighbors, whereas firms 0 and m compete in the product market. Likewise, I only consider dynamic linkages between periods through the stock of foundational innovation S , which accrues with each new quantum of knowledge created. This model is based on Bloom et al. (2013), but deviates in two main respects. First, it adds a dynamic dimension to account for dynamic linkages in the firms' R&D decisions across time. Second, it considers a more complex innovation production function incorporating past innovations as an input. The timing of the model is shown in Figure A-4.

In each of the second stages, and for a given level of knowledge k , the decision variables x only affect profits for that given period. The second stages of the game can thus be solved as in a static game, and we can abstract from the period subscripts for now. Let firm 0's profit function in each second stage be $\pi^0(x^0, x^m, k^0)$, with π^0 concave and increasing in k^0 . The interpretation for this profit function is that firms 0 and m compete in x in the product market, and firm 0 has a knowledge level k^0 that affects its profits positively through e.g. increased productivity (i.e. decreased production costs), or increased market power, or reduced elasticity of residual demand. The derivative with respect to x^0 is ambiguous (depending for example on whether x represents prices or quantities), but

we have $\text{sign}(\pi_1) = -\text{sign}(\pi_2)$, where the subscripts refer to the partial derivatives with respect to each argument. That is, $\text{sign}\left(\frac{\partial \pi}{\partial x^0}\right) = -\text{sign}\left(\frac{\partial \pi}{\partial x^m}\right)$. If the firms are symmetric in their profit function, the best response for firms 0 and m is given by $x^{0*} = \text{argmax}_{x^0} \pi(x^0, x^m, k^0)$ and $x^{m*} = \text{argmax}_{x^m} \pi(x^m, x^0, k^m)$. Solving for the Nash equilibrium yields $x^{0*} = f^0(k^0, k^m)$ and $x^{m*} = f^m(k^m, k^0)$. We can thus write $\Pi^0(k^0, k^m) = \pi(x^{0*}, x^{m*}, k^0)$, with Π^0 increasing in k^0 . If there is no strategic interaction in the product market between 0 and m , then Π^0 will not depend on k^m . I assume that Π^0 is non-increasing in k^m , and concave.⁷

In each of the first stages, let the knowledge production function be $k^0 = \phi(r^0, r^\tau, S)$, where r^0 is the R&D spending of firm 0, r^τ is R&D spending by its technological neighbor τ , and S is the stock of existing knowledge. I assume that the innovation production function ϕ is common to all firms and is non-decreasing and concave in its arguments. Therefore, knowledge spillovers, both static and dynamic, are positive. Let us assume that the existing innovation stock S evolves as new research creates new ideas. Therefore, the innovation stock in period two will be given by $S' = \psi(r^0, r^\tau, S)$, where ψ is increasing and concave in all its arguments. In the second period, firm 0 thus solves

$$\max_{r_0} V_2^0 = \Pi(\phi(r_2^0, r_2^\tau, S'), k_2^m) - r_2^0. \quad (\text{A.5})$$

The FOC that pins down the optimal level of R&D for firm 0 in the second period is thus $\Pi_1 \phi_1 = 1$ at the equilibrium levels. In the first period, the maximization problem for firm 0 is more complex, as it includes both the influence of R&D on immediate profits and on future profits through its effect on S' :

$$\max_{r_0} V_1^0 + \delta V_2^0 = \Pi(\phi(r_1^0, r_1^\tau, S), k_1^m) - r_1^0 + \delta [\Pi(\phi(r_2^0, r_2^\tau, \psi(r_1^0, r_1^\tau, S)), k_2^m) - r_2^0]. \quad (\text{A.6})$$

The FOC with respect to r_1^0 is now $\Pi_1 \phi_1 + \Pi_1 \phi_3 \psi_1 = 1$ at the equilibrium levels. There is now an extra term taking into account the added benefit of carrying out research in period 1 because of it increasing the foundational knowledge base S' available for period 2. In order to analyze how exogenous shifts in the R&D spending by neighbors (τ and m) affect outcomes for firm 0 we draw comparative statics in

⁷This assumption is reasonable unless innovation by m yields such large value creation through market expansion that it actually helps competitors. Bloom et al. (2013) show that this assumption is reasonable empirically.

period two. Regarding firm 0's stock of knowledge and its value, we find that

$$\frac{\partial k_2^0}{\partial r_2^\tau} = \phi_2 \geq 0, \quad (\text{A.7})$$

$$\frac{\partial k_2^0}{\partial r_m} = 0, \quad (\text{A.8})$$

$$\frac{\partial k_2^0}{\partial r_1^\tau} = \phi_3 \psi_2 \geq 0, \quad (\text{A.9})$$

$$\frac{\partial k_2^0}{\partial r_1^0} = \phi_3 \psi_1 \geq 0, \quad (\text{A.10})$$

$$\frac{\partial V_2^{0*}}{\partial r_2^\tau} = \Pi_1 \phi_2 \geq 0, \quad (\text{A.11})$$

$$\frac{\partial V_2^{0*}}{\partial r_2^m} = \Pi_2 \phi_1 \leq 0. \quad (\text{A.12})$$

$$\frac{\partial V_2^{0*}}{\partial r_1^\tau} = \Pi_1 \phi_3 \psi_2 \geq 0, \quad (\text{A.13})$$

$$\frac{\partial V_2^{0*}}{\partial r_1^0} = \Pi_1 \phi_3 \psi_1 \geq 0, \quad (\text{A.14})$$

As shown above, the predictions of the baseline model with respect to knowledge spillovers are qualitatively similar irrespective of the dynamic or static nature of those spillovers. Both spillovers accruing from own past R&D and those from technological neighbors' past R&D should have a positive in both firm knowledge and profit levels. Likewise, the ratio between the effect of own and others' dynamic spillovers will be equal to ψ_1/ψ_2 for both sets of outcomes. The predictions of the model are summarized in Table 1.2, under the assumptions of positive technology spillovers and strategic complementarity between the product market competitors' knowledge stock k_0 and k_m . The only difference in predictions on the effect of spillovers on knowledge stock and market value occurs for the business stealing spillovers, that are supposed to affect market value but not knowledge capital. As I discuss in section 1.4, in my empirical analysis I can unfortunately not cleanly separate effects of spillovers through prices and through "physical quantity" productivity.⁸ I therefore do not attempt to

⁸See Syverson (2011) for an extensive discussion on physical quantity productivity measures and

test this differential prediction in my empirical specifications.

A.3.2 Computing private and social returns without amplification

In this appendix, I show how to derive the equations for the private and the social returns to R&D, abstracting for now from R&D amplification mechanisms.⁹ I define the *marginal social returns* to R&D of firm i at time t (MSR_{it}) as the marginal increase in aggregate output due to a marginal increase in the R&D stock of firm i at time t . Likewise, the *marginal private returns* to R&D of firm i at time t (MPR_{it}) are defined as the marginal increase in firm i 's output due to a marginal increase in its R&D stock at time t . Because of the inherent dynamism of the spillovers involved, I will consider as a measure of output the net present value of all current and future output (that is, at times $t' \geq t$), discounted at rate r . Therefore, we have

$$MPR_{it} = \frac{d \left[\sum_{\tau=0}^{\infty} Y_{it+\tau} \frac{1}{(1+r)^\tau} \right]}{dRDS_{it}}, \quad (\text{A.15})$$

$$MSR_{it} = \frac{d \left[\sum_{\tau=0}^{\infty} Y_{t+\tau}^T \frac{1}{(1+r)^\tau} \right]}{dRDS_{it}}, \quad (\text{A.16})$$

where $Y_t^T = \sum_j Y_{jt}$ is aggregate output at time t and RDS is R&D stock. From my empirical specification, I take the following output equation:

$$\begin{aligned} \ln Y_{it} = & \phi_{rds} \ln RDS_{it} + \phi_{dyn} \ln \sum_{j,t' \leq t} Dyn_{ijt'} RDS_{jt'} \\ & + \phi_{cit} \ln \sum_{j \neq i} Cit_{ij} RDS_{jt} + \phi_{sic} \ln \sum_{j \neq i} Sic_{ij} RDS_{jt} + \phi_5 X_{it}^Y. \end{aligned} \quad (\text{A.17})$$

Let us define the $N \times T$ vectors Y and RDS of all the Y_{jt} and RDS_{jt} terms. In order to have a linear relationship between log-output and log-R&D stock, I take a first-order expansion of the three spillover terms $\ln \sum_{j,t' \leq t} D_{ijt'} RDS_{jt'}$ in terms of $\ln RDS$ around a point $\ln RDS^0$. That is, for a function $f_{it}(RDS) = \ln \sum_{j,t' \leq t} D_{ijt'} RDS_{jt'}$, we have

revenue-based productivity.

⁹These amplification mechanisms would occur through the effect of the spillover terms on the R&D investment decisions of the firms, as in equation 1.16.

$$f_{it}(RDS) = a_{it}^D + \sum_{j,t' \leq t} b_{ijt'}^D \ln RDS_{jt'}, \quad (\text{A.18})$$

where $b_{it}^D = \sum_{j,t' \leq t} \frac{D_{ijt'} RDS_{jt'}^0}{\sum_{k,t'' \leq t} D_{ikt''} RDS_{kt''}^0}$
and $a_{it}^D = \ln \sum_{j,t' \leq t} D_{ijt'} RDS_{jt'}^0 - \sum_{j,t' \leq t} \frac{D_{ijt'} RDS_{jt'}^0}{\sum_{k,t'' \leq t} D_{ikt''} RDS_{kt''}^0} \ln RDS_{jt'}^0$.

Therefore, and denoting $\lambda_{it} = \phi_{dyn} a^{dyn} + \phi_{cit} a^{cit} + \phi_{sic} a^{sic}$, we can write the relationship between output and R&D stock as

$$\begin{aligned} \ln Y_{it} = & \lambda_{it} + \phi_{rds} \ln RDS_{it} + \phi_{dyn} \sum_{j,t' \leq t} b_{ijt'}^{D_{yn}} \ln RDS_{jt'} \\ & + \phi_{cit} \sum_{j \neq i} b_{ij}^{Cit} \ln RDS_{jt} + \phi_{sic} \sum_{j \neq i} b_{ij}^{Sic} \ln RDS_{jt} + \phi_5 X_{it}^Y. \end{aligned} \quad (\text{A.19})$$

In fact, we can even separate the dynamic knowledge spillover term into *intra* and *between* spillovers, as in subsection 1.7.1. In that case, the marginal private returns of R&D are then

$$\text{MPR}_{it} = \frac{d \left[\sum_{\tau=0}^{\infty} Y_{it+\tau} \frac{1}{(1+r)^\tau} \right]}{dRDS_{it}} = \sum_{\tau=0}^{\infty} \frac{1}{(1+r)^\tau} \frac{Y_{it+\tau}}{RDS_{it}} \frac{d \ln Y_{it+\tau}}{d \ln RDS_{it}}, \quad (\text{A.20})$$

where we have $\frac{d \ln Y_{it+\tau}}{d \ln RDS_{it}} = \phi_{rds}$ if $\tau = 0$ and $\frac{d \ln Y_{it+\tau}}{d \ln RDS_{it}} = \phi_{intra} b_{it+\tau it}^{intra}$ for $\tau > 0$. Let us define $\delta'_{it} = \sum_{\tau=0}^{\infty} \frac{1}{(1+r)^\tau} b_{it+\tau it}^{intra}$. If we assume that all firm-years have the same citation patterns with each other and their past selves,¹⁰ then $\delta'_{it} = \delta'$ does not depend on the firm nor time considered. Moreover, if we assume that all firms are also symmetric and equal in terms of their baseline output Y and R&D stock levels RDS , we then have that $b_{it+\tau it}^{intra}$ is just the share of citations within firm i and between times $t + \tau$ and t , over the share of intra-citations (within i) from $it + \tau$. The term δ' is therefore the discounted, weighted by the within-firm citation shares for each time gap τ , of the unitary constant. We then have

$$\text{MPR}_{it} = \frac{Y}{RDS} (\phi_{rds} + \delta' \phi_{intra}). \quad (\text{A.21})$$

¹⁰That is, $Intra_{it+\tau it} = Intra_\tau$.

As for the marginal social returns, the process is similar but now considering the output of all the firms involved

$$\text{MSR}_{it} = \frac{d \left[\sum_{\tau=0}^{\infty} \sum_j Y_{jt+\tau} \frac{1}{(1+r)^\tau} \right]}{dRDS_{it}} = \sum_{\tau=0}^{\infty} \frac{1}{(1+r)^\tau} \sum_j \frac{Y_{jt+\tau}}{RDS_{it}} \frac{d \ln Y_{jt+\tau}}{d \ln RDS_{it}}. \quad (\text{A.22})$$

If we assume again that all proximities between firms are the same,¹¹ and that the firms are of the same size in terms of baseline output and R&D stock levels, we have

$$\text{MSR}_{it} = \frac{Y}{RDS} (\phi_{rds} + \phi_{cit} + \phi_{sic} + \delta \phi_{dyn}), \quad (\text{A.23})$$

where δ is similar to δ' above, but using all the dynamic (intra and between) citation patterns Dyn .

In terms of elasticities, we have that the percentage increase in firm i 's own output relative to a percentage increase in own R&D stock is $\phi_{rds} + \delta' \phi_{intra}$. Assuming symmetric firms, the elasticity of aggregate output to firm i 's R&D is $\frac{1}{N} (\phi_{rds} + \phi_{cit} + \phi_{sic} + \delta \phi_{dyn})$. Finally, if we are instead interested in how output changes with a marginal change in R&D expenditures, rather than stock, we have to take into account the durable impact of investment on R&D stock. Therefore, if we assume that R&D depreciates at a 15% rate and using a discount rate of 6%, we should multiply the previous expressions by about 5.

A.3.3 Optimal provision of R&D

Once we have estimates for the marginal private and social returns to R&D, the question we move to is what the implications are with respect to the optimal provision of corporate R&D investment in the economy. For this, let us assume that firms can invest in R&D r at a constant marginal cost c , which represents the user cost of R&D investment.¹² R&D increases profits in the product market $\pi(r)$, through for example reductions in the costs of production. In this case, the firm solves for $\text{Max}_r \pi(r) - cr$. As discussed in appendix A.2.1 and section 1.7.2, we find that in general the relationship between corporate R&D expenditures r and the tax-induced

¹¹That is, $Sic_{ij} = Sic$, $Cit_{ij} = Cit$, and $Dyn_{itjv} = Dyn_{v-t}$.

¹²Remember that the user cost of R&D investment is affected among others by the corporate income taxes and R&D tax credits, but also by the real interest rate and the depreciation rate of R&D capital.

user cost of capital c is unitary iso-elastic, that is $\ln r = \ln a - \ln c$. This relationship can be rationalized in the framework above with a profit function $\pi(r) = a \ln r + b$.

If we abstract for now from consumer surplus, the difference between the socially optimal level of R&D investment by firm i and the private optimum is that the social optimum also takes into account the effect of r_i on all the other firms' profits. That is, the social optimum solves $\sum_j \frac{\partial \pi_j}{\partial r_i}(r^S) = c$ while the private optimum solves $\frac{\partial \pi_i}{\partial r_i}(r^*) = c$. If we assume that the relationship between π_i and all the R&D investments r_j is given by $\pi_i(r) = \sum_j a_{ij} \ln r_j + b_i$, inspired by the relationship in the previous paragraph, and we also assume that the marginal increases in profits are proportional to the marginal increases in output,¹³ then we would have that

$$\frac{\frac{\partial \pi_i}{\partial r_i}(r)}{\sum_j \frac{\partial \pi_j}{\partial r_i}(r)} = \frac{\text{MPR}_i}{\text{MSR}_i} = \frac{a_{ii}}{\sum_j a_{ji}}. \quad (\text{A.24})$$

However, at the private and social optimal levels r^* and r^S we have $\sum_j \frac{\partial \pi_j}{\partial r_i}(r^S) = \frac{\partial \pi_i}{\partial r_i}(r^*) = c$.¹⁴ That is, $\frac{\sum_j a_{ji}}{r_i^S} = \frac{a_{ii}}{r_i^*}$. Therefore, the ratio between the socially optimal level of R&D and the privately optimal $\frac{r_i^S}{r_i^*}$ will be equal to the ratio between the MSR and the MPR.

This analysis does not fully take into account that the social planner also cares about consumer surplus, rather than only about profits. Accounting for consumer surplus would increase the under-provision of R&D in the decentralized economy. Let us assume that the marginal increases in consumer surplus across product markets are proportional to the increases in output, and that we have an estimate of the ratio γ in a given market between the privately optimal level of R&D (only accounting for profits) and the socially optimal level, accounting for total surplus. In that case, the ratio between the socially optimal level of R&D accounting for total surplus and spillovers to the private equilibrium level will be equal to $\gamma \frac{\text{MSR}}{\text{MPR}}$, proportional to the ratio between MSR and MPR.

¹³This assumption would hold, for example, if the mark-ups are constant and equal for all firms.

¹⁴This assumes that the user cost of capital c is the same for the social and the private returns calculation. This is likely to be the case for the decentralized economy without distortionary taxes, that is if c is not affected by taxes. Since the tax treatment does affect the user cost c , the observed level of R&D is likely to be different than the decentralized optimum.

A.4 Using alternative proximity measures

In section 1.4 I analyze the effect of including both dynamic and static knowledge spillovers, as well as business stealing, in the empirical analysis. I run here an extended amount of specifications to account for the various possibilities in constructing static knowledge and business stealing spillover measures. Results are shown in Tables A.2 and A.3, for the productivity and market value equations respectively. For the productivity equation in Table A.2, the coefficients on dynamic spillovers are remarkably robust across all specifications. Columns (1), (2) and (7) use *CitSpill* as a measure of static knowledge spillovers and *SicSpill* for static business stealing spillovers. These regressions are the same as those reported in Table 1.5, which are discussed in subsection 1.4.2.

I also use *TechSpill* and *MahTechSpill* as measures of static knowledge spillovers.¹⁵ These measures yield large and significant coefficients in the first OLS specification, which are robust to including three digit industry-times-year dummies. The 2SLS specification, which takes into account the endogeneity in the R&D decisions due for example to correlated shocks, yields insignificant coefficients for these two measures however. Among the firms in my sample, it thus seems that the Jaffe (1986) and extended Mahalanobis technology measures may not be appropriate to measure knowledge spillovers. As discussed in subsection A.1.1, and particularly compared to *CitSpill*, technological space proximities may be too lax in their definition of which R&D is likely to spill over.

In terms of business stealing, the evidence is not strong. Although all of the point estimates in the table are negative, only specification (7) leads to a statistically significant coefficient. In the market value equation in Table A.3, the coefficients on dynamic spillovers are again robust and statistically significant across specifications. As for the static spillover coefficients, only *CitSpill* together with *SicSpill* seem to pick up statistically significant spillovers. Nonetheless, it is important to remark that in general I cannot reject the significant results in Bloom et al. (2013) for the static spillovers in my sample.

There are a number of reasons for the departures between the results in Bloom et al. (2013) and those discussed here. First, I only analyze about half of their firms in my sample –generally larger and longer-lived ones– as I restrict my analysis to firms

¹⁵Columns (5), (6) and (9) include *MahSicSpill* rather than *SicSpill*.

also observed in the pre-period network. I also consider only years starting in 1990 rather than 1985 as they do. These two factors drive most of the change in the loss of significance in business stealing effects. Second, I use different definitions of the market-to-book value than they do. I base myself on most of the finance literature in doing so¹⁶ and I believe my measures to be more relevant for the task at hand. In order to ensure that my market value results are robust and not driven solely by the definition of the market-to-book ratio, I carry out robustness exercises using an alternative definition. In particular, I also follow Barrot et al. (2016), Rauh and Sufi (2012), Rhodes-Kropf et al. (2005), and Warusawitharana and Whited (2016) among others, in defining market value as $AT + (PRCC_F * CSHO) - CEQ - TXDB$ and dividing it by total assets AT to obtain the market-to-book ratio. Dynamic spillovers are robust across all three definitions of the market-to-book ratio. Third, I allow for more flexible correlation in the error terms, both cross-sectionally and serially. I cluster two-way at the firm and at the year level in all my specifications. Fourth, I include a flexible set of two digit SIC-code industry-times-year dummies to account for common or correlated shocks. Fifth, there a number of small variable construction details that vary between both our papers. Although each of these decisions is not likely to be significant, it can lead to larger effects in the end. Examples include the fact that I winsorize all my data at the 1st and 99th percentile to ensure that results are not driven by outliers; or that I separate $\ln(\text{R\&D stock}/AT)$ into $\ln(\text{R\&D stock})$ and $\ln(AT)$ in the value equations in order to avoid biasing estimators due to dividing the RHS and LHS by the same variable.

A.5 Additional results

I show here more detailed results on the heterogeneity of dynamic spillovers within and across industrial sectors and technology classes.

A.5.1 Dynamic spillovers within and between industries

In section 1.5, I examine the heterogeneity in dynamic spillovers across industries. Here I show how spillovers vary depending on whether they accrue within or between industries. That is, whether the emitting and the receiving firm are in the same

¹⁶See Davis, Fama, and French (2000), or Kenneth French's website for details on variable construction.

industry or not. 58% of weighted edges in the patent citation network underpinning the dynamic spillover measure occur between firms belonging to the same 2-digit SIC code industry, so dynamic spillovers occur majoritarily within industries.

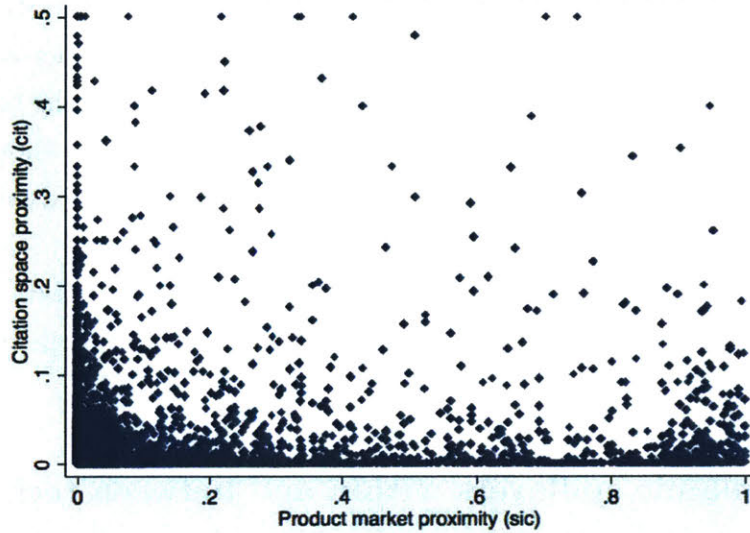
I separate dynamic spillover measures depending on whether they accrue within or between industries, and estimate OLS and 2SLS specifications of the baseline productivity equation including both measures of dynamic spillovers. I show the resulting estimates in the left panels of Figure A-5. The elasticities on dynamic spillovers are positive and generally statistically significant regardless of whether they accrue within or between industries. The point estimates are larger within than between, consistent with them constituting a larger share of citations, but the difference is not statistically significant.

A.5.2 Dynamic spillovers within and between technologies

I then examine a similar heterogeneity across technology types, depending on whether the originating patent and the receiving patent belong to the same Hall et al. (2001) technology class. A large majority, 77%, of weighted citations in the patent citation network accrue within the same technology class.

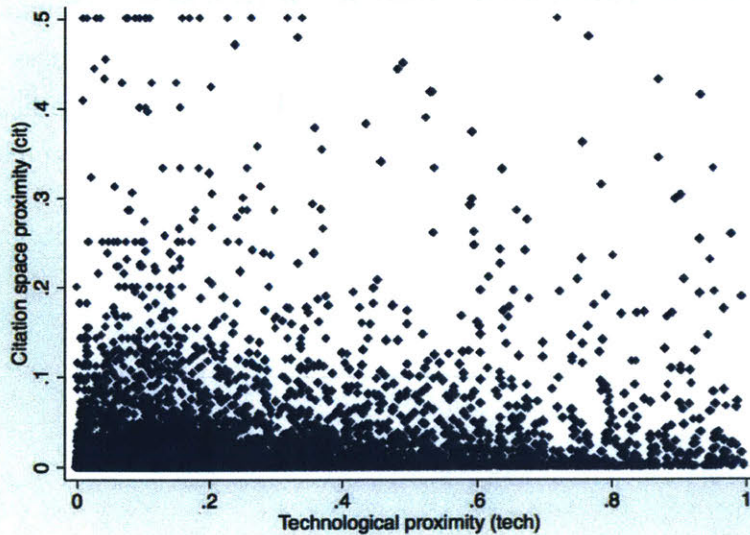
I separate dynamic spillover measures depending on whether they accrue within or between technology class, and again estimate the OLS and 2SLS specifications of the baseline productivity equation including both measures of dynamic spillovers. The resulting estimates are shown in Figure A-5, in the right panels. The elasticity estimates on dynamic spillovers are positive and statistically significant for within-class spillovers in both specifications, but not for spillovers between classes.

Figure A-1: Correlation between CIT and SIC proximity measures



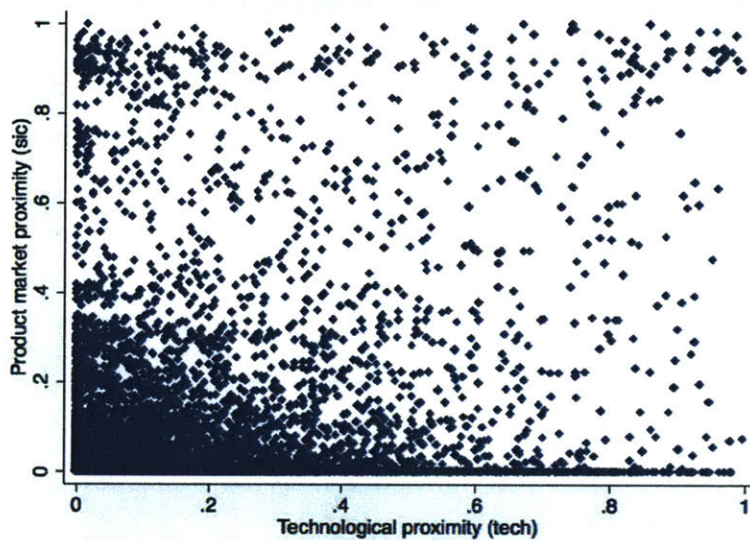
Notes: This figure plots the pairwise values of the proximity in citation space *CIT* and proximity in product market space *SIC* for all pairs of firms in my sample. In order to make the figure clearer, *CIT* values are topped at 0.5.

Figure A-2: Correlation between CIT and TEC proximity measures



Notes: This figure plots the pairwise values of the proximity in citation space *CIT* and proximity in technological space *TEC* for all pairs of firms in my sample. In order to make the figure clearer, *CIT* values are topped at 0.5.

Figure A-3: Correlation between SIC and TEC proximity measures



Notes: This figure plots the pairwise values of the proximity in technological space *TEC* and proximity in product market space *SIC* for all pairs of firms in my sample.

Figure A-4: Timing of analytical model

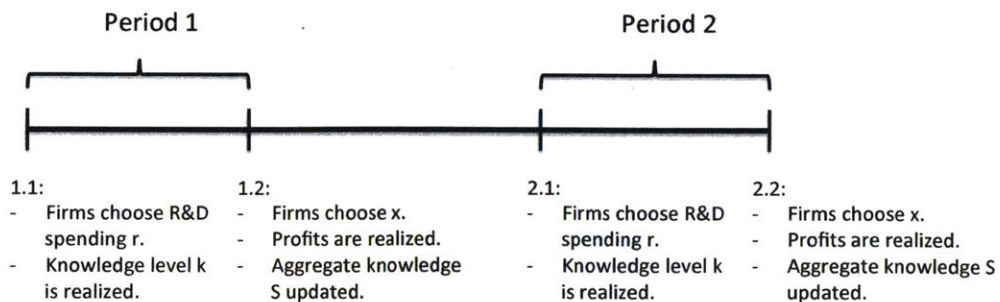
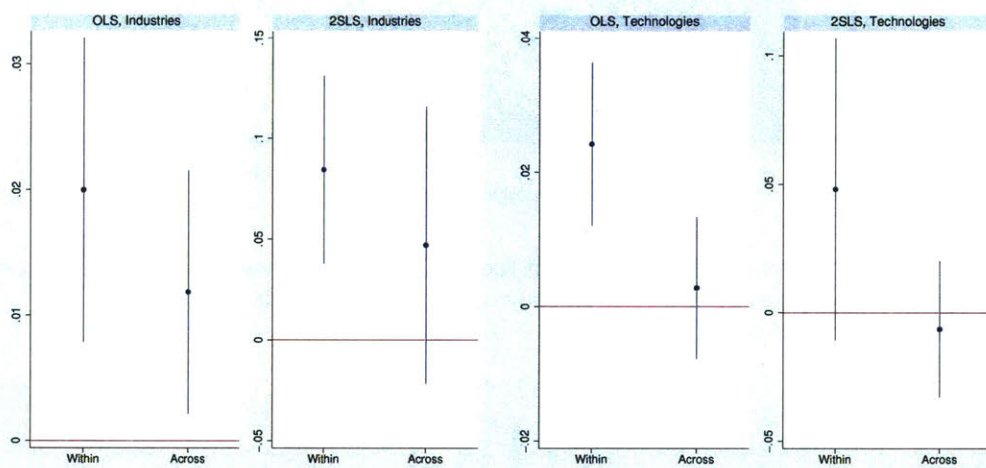


Figure A-5: Dynamic spillovers between and within industries and technology categories



Notes: This figure plots the values and confidence intervals of the coefficient on dynamic spillovers, separated by whether they accrue across or within industries of the originating and receiving firm, and Hall et al. (2001) technology classes of the citing and cited patent, for the productivity equation, and OLS and 2SLS specifications. The regressions run are the same as in the baseline regressions in Table 1.5, albeit incorporating both types of dynamic spillovers. Standard errors are clustered two-way at the year and firm level, and confidence intervals are set at the 90% level.

Table A.1: Effect of R&D intensity on average patent citations

	log patent count		Average citation count			Average Standardized citation count	
	OLS (1)	OLS (2)	OLS (3)	OLS (4)	2SLS (5)	OLS (6)	2SLS (7)
log (R&D stock)	0.356*** (0.047)	0.009 (0.018)	0.037* (0.020)				
log (employees)			-0.054** (0.022)				
log (R&D stock/Assets)				0.036** (0.015)	0.232* (0.136)	0.061*** (0.022)	0.349* (0.194)
F-statistic					32.34549		32.34549
Observations	10757	10757	10401	10757	10757	10757	10757
R-squared	.87	.6	.6	.6	.59	.53	.51
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry x year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable is log of patent count in column (1), log of average citation count per patent in columns (2) to (5), and average of the standardized citation count per patent, where citations are standardized to mean zero and unitary variance within each technological subcategory and application year. Regressions include firm and 3-digit SIC-code industry-times-year fixed effects. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table A.2: Alternative static spillover measures, sales

	OLS						IV		
	Citation		Jaffe tech.		Mah. tech.		Cit.	Jaffe	Mah.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
ln(R&D stock)	0.036 (0.027)	0.029 (0.026)	0.035 (0.028)	0.029 (0.027)	0.036 (0.027)	0.027 (0.026)	0.092* (0.049)	0.129*** (0.050)	0.132*** (0.049)
Static spill.	0.193** (0.076)	0.122** (0.062)	0.324* (0.170)	0.343* (0.200)	0.610** (0.246)	0.693** (0.287)	0.280** (0.130)	0.035 (0.269)	0.016 (0.392)
Business steal.	-0.040 (0.032)	-0.026 (0.030)	-0.018 (0.031)	-0.024 (0.030)	-0.069 (0.076)	-0.073 (0.083)	-0.087** (0.038)	-0.032 (0.025)	-0.058 (0.164)
Dynamic spill.	0.023*** (0.008)	0.025*** (0.009)	0.022*** (0.008)	0.025*** (0.009)	0.022*** (0.008)	0.025*** (0.009)	0.107** (0.047)	0.098** (0.047)	0.097** (0.047)
First stage F-test							13.68	13.642	11.719
Observations	3631	3631	3631	3631	3631	3631	3631	3631	3631
Firm and year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Industry x year FE		✓		✓		✓			

Notes: Dependent variable is lead ln(Sales). Regressions include the log of an industry-specific price deflator; industry-wide log-sales and lagged log-sales; log counts of patents filed; dummies for no R&D, for no dynamic spillover and for no patents filed, as well as a full set of firm and year FEs. Columns (2), (4) and (6) also include 2-digit SIC-code industry-times-year FEs. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Table A.3: Alternative static spillover measures, MTB

	OLS						IV		
	Citation		Jaffe tech.		Mah. tech.		Cit.	Jaffe	Mah.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
ln(R&D stock)	0.024 (0.050)	0.012 (0.049)	0.026 (0.050)	0.017 (0.049)	0.027 (0.051)	0.011 (0.049)	0.144 (0.142)	0.189 (0.130)	0.177 (0.133)
Static spill.	0.019 (0.162)	-0.044 (0.198)	-0.084 (0.517)	-0.283 (0.490)	-0.189 (0.644)	-0.100 (0.638)	0.277** (0.133)	-0.357 (0.956)	0.419 (1.235)
Business steal.	0.116 (0.103)	0.055 (0.087)	0.127 (0.094)	0.065 (0.072)	0.280 (0.268)	0.115 (0.267)	-0.112* (0.067)	-0.042 (0.067)	-0.421 (0.470)
Dynamic spill.	0.050*** (0.019)	0.051** (0.022)	0.050*** (0.019)	0.051** (0.022)	0.050*** (0.019)	0.051** (0.022)	0.171* (0.089)	0.160* (0.089)	0.156* (0.092)
First stage F-test							11.275	11.239	9.9082
Observations	3561	3561	3561	3561	3561	3561	3561	3561	3561
Firm and year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Ind. x year FE		✓		✓		✓			

Notes: Dependent variable is ln(MTB). Regressions include a sixth-order polynomial in ln(R&D intensity), only the first term is shown for brevity; industry-wide log-sales and lagged log-sales; log counts of patents filed; dummies for no R&D, for no dynamic spillover and for no patents filed, as well as a full set of firm and year FEs. Columns (2), (4) and (6) also include 2-digit SIC-code industry-times-year FEs. Standard errors in brackets are clustered two-way at the year and firm level. ***, ** and * denote statistical significance at the 1%, 5% and 10% respectively.

Bibliography

- K. G. Abraham and L. F. Katz. Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances? *Journal of Political Economy*, 94(3):507, 1986.
- D. S. Abrams. Did TRIPS Spur Innovation? An Analysis of Patent Duration and Incentives to Innovate. *University of Pennsylvania Law Review*, 157:1613–1647, 2009. ISSN 0041-9907.
- D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016, 2012.
- D. Acemoglu, D. Autor, D. Dorn, G. H. Hanson, and B. Price. Import Competition and the Great US Employment Sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198, 2015.
- D. Acemoglu, U. Akcigit, and W. Kerr. Innovation Network. *NBER Working Paper 22783*, 2016.
- P. Aghion and P. Howitt. A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–351, 1992.
- P. Aghion and G. Saint-Paul. Interaction Between Productivity Growth and Economic Fluctuations. *Macroeconomic Dynamics*, (2):322–344, 1998.
- J. Alcacer and M. Gittelman. Patent Citations as a Measure of Knowledge Flows : the Influence of Examiner Citations. *The Review of Economics and Statistics*, 88(4):774–779, 2006.
- F. Alvarez and R. Shimer. Search and Rest Unemployment. *Econometrica*, 79(1):75–122, 2011.
- B. N. Anand and T. Khanna. The structure of licensing contracts. *Journal of Industrial Economics*, 48(1):103–135, 2000. ISSN 0022-1821. doi: 10.1111/1467-6451.00114.
- A. Arora, A. Fosfuri, and A. Gambardella. *Markets for Technology : The Economics of Innovation and Corporate Strategy*. MIT Press, 2001.
- A. Arora, W. M. Cohen, and J. P. Walsh. The acquisition and commercialization of invention in American manufacturing: Incidence and impact. *Research Policy*, 45(6):1113–1128, 2016.

- K. Arrow. Economic welfare and the allocation of resources for invention. In R. R. Nelson, editor, *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pages 609–626. Princeton University Press, Princeton, NJ, 1962.
- D. H. Autor and D. Dorn. The growth of low-skill service jobs and the polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597, 2013. ISSN 00028282. doi: 10.1257/aer.103.5.1553.
- D. H. Autor, D. Dorn, and G. H. Hanson. The China Syndrome: Local labor market impacts of import Competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.
- P. Azoulay, J. S. G. Zivin, D. Li, and B. N. Sampat. Public R&D Investments and Private-sector Patenting: Evidence from NIH Funding Rules. *NBER Working Paper*, 2015.
- J.-N. Barrot and J. Sauvagnat. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *Quarterly Journal of Economics*, Forthcomin, 2016.
- J.-N. Barrot, E. Loualiche, and J. Sauvagnat. The Globalization Risk Premium. *Mimeo*, 2016.
- E. J. Bartelsman and W. Gray. The NBER Manufacturing Productivity Database, 1996.
- A. W. Bartik. Worker Adjustment to Changes in Labor Demand : Evidence from Longitudinal Census Data. *Mimeo*, 2017.
- M. Baxter. Are Consumer Durables Important for Business Cycles? *The Review of Economics and Statistics*, 78(1):147–155, 1996.
- P. Beaudry, D. Galizia, and F. Portier. Reconciling Hayek’s and Keynes’ views of recessions. *NBER Working Paper #20101*, 2014.
- B. Becker. Public R&D Policies and Private R&D Investment: A Survey of the Empirical Evidence. *Journal of Economic Surveys*, 29(5):917–942, 2015.
- S. Belenzon. Cumulative Innovation and Market Value : Evidence From Patent Citations. *The Economic Journal*, 122(March):265–285, 2012.
- D. Berger. Countercyclical Restructuring and Jobless Recoveries. *Mimeo*, 2016.
- J. Bessen. Holdup and licensing of cumulative innovations with private information. *Economics Letters*, 82(3):321–326, 2004. ISSN 01651765. doi: 10.1016/j.econlet.2003.08.004.
- J. Bessen. Estimates of patent rents from firm market value. *Research Policy*, 38(10):1604–1616, 2009.

- B. Biasi and P. Moser. Effects of Copyrights on Science - Evidence from the WWII Book Republication Program. 2017.
- M. Bils and P. Klenow. Using Consumer Theory to Test Competing Business Cycle Models. *Journal of Political Economy*, 106(2), 1998.
- M. Bils, P. Klenow, and B. Malin. Testing for Keynesian labor demand. In *NBER Macroeconomics Annual*, volume 27, pages 311–349. 2013. ISBN 9780226052779. doi: 10.1086/669182.
- N. Bloom, R. Griffith, and J. V. Reenen. Do R & D Tax Credits Work? Evidence from a Panel of Countries 1979-1997. *Journal of Public Economics*, 85:1–31, 2002.
- N. Bloom, M. Schankerman, and J. Van Reenen. Identifying Technology Spillovers and Product Market Rivalry. *Econometrica*, 81(4):1347–1393, 2013.
- J. Bound, C. Cummins, Z. Griliches, B. H. Hall, and A. B. Jaffe. Who Does R & D and Who Patents ? In Z. Griliches, editor, *R&D, Patents, and Productivity*, chapter Chapter 2, pages 21 – 54. University of Chicago Press, 1984.
- E. Budish, B. N. Roin, and H. Williams. Do firms underinvest in long-term research? Evidence from cancer clinical trials. *American Economic Review*, 105(7):2044–2085, 2015. ISSN 00028282. doi: 10.1257/aer.20131176.
- R. J. Caballero and M. L. Hammour. On the Timing and Efficiency of Creative Destruction. *NBER Working Paper #4768*, 1994.
- S. Calonico, M. D. Cattaneo, and R. Titiunik. Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326, 2014. ISSN 00129682. doi: 10.3982/ECTA11757.
- V. M. Carvalho, M. Nirei, and Y. U. Saito. Supply Chain Disruptions : Evidence from Great East Japan Earthquake. *Mimeo*, 2014.
- A. C. Chang. Tax Policy Endogeneity : Evidence from R & D Tax Credits. *Mimeo*, 2014.
- K. W. Chauvin and M. Hirschey. Advertising, R&D expenditures and the market value of the firm. *Financial Management*, 22(4):128–140, 1993.
- R. S. Chirinko and D. J. Wilson. State investment tax incentives: A zero-sum game? *Journal of Public Economics*, 92(12):2362–2384, 2008. ISSN 00472727. doi: 10.1016/j.jpubeco.2008.07.005.
- R. S. Chirinko and D. J. Wilson. Tax Competition among U.S. States: Racing to the Bottom or Riding on a Seesaw? 2013.
- G. Chodorow-Reich and J. Wieland. Secular Labor Reallocation and Business Cycles. *NBER Working Paper #21864*, 2016.

- A. Coad and R. Rao. Innovation and firm growth in high-tech sectors: A quantile regression approach. *Research Policy*, 37(4):633–648, 2008. ISSN 00487333. doi: 10.1016/j.respol.2008.01.003.
- W. M. Cohen and D. A. Levinthal. Absorptive Capacity : A New Perspective on Learning and Innovation. *Administrative Science Quarterly*, 35(1):128–152, 1990.
- W. M. Cohen, R. R. Nelson, and J. P. Walsh. Protecting their Intellectual Assets: Appropriability Conditions and Why US Manufacturing Firms Patent (or not). *NBER Working Paper*, 2000.
- D. Colino. Knowledge Diffusion and the Patent Citation Network. 2016.
- S. Correia. Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix. 2015.
- S. Correia. A Feasible Estimator for Linear Models with Multi-Way Fixed Effects. 2016.
- J. L. Davis, E. F. Fama, and K. R. French. Characteristics, covariances, and average returns: 1929 to 1997. *The Journal of Finance*, 55(1):389–406, 2000.
- S. J. Davis and J. Haltiwanger. *Microeconomic Evidence and Macroeconomic Implications**, volume 5. 1990.
- A. Dechezleprêtre, E. Einiö, R. Martin, K.-T. Nguyen, and J. Van Reenen. Do Tax Incentives for Research Increase Firm Innovation? An RD Design for R&D. 2016.
- G. Ellison, E. L. Glaeser, and W. R. Kerr. What causes industry agglomeration? Evidence from coagglomeration patterns. *American Economic Review*, 100(3):1195–1213, 2010.
- C. Engel and J. Wang. International trade in durable goods: Understanding volatility, cyclicity, and elasticities. *Journal of International Economics*, 83(1):37–52, 2011.
- M. Fletcher. Dole wins Helms pledge to temper attacks on Clinton (The Times), nov 1994.
- J. L. Furman and S. Stern. Climbing atop the Shoulders of Giants : The Impact of Institutions on Cumulative Research. *American Economic Review*, 101(5):1933–1963, 2011.
- J. L. Furman, M. Kyle, I. Cockburn, and R. Henderson. Public & Private Spillovers, Location and the Productivity of Pharmaceutical Research. *NBER Working Paper*, 2006.
- A. Galasso and M. Schankerman. Patents and Cumulative Innovation: Causal Evidence from the Courts. *The Quarterly Journal of Economics*, pages 317–369, 2015. doi: 10.1093/qje/qju029.Advance.

- P. Ganong and S. Jaeger. A Permutation Test for the Regression Kink Design. 2016.
- R. Gilbert and C. Shapiro. Optimal Patent Length and Breadth. *The RAND Journal of Economics*, 21(1):106–112, 1990.
- X. Giroud and H. M. Mueller. Does corporate governance matter in competitive industries? *Journal of Financial Economics*, 95(3):312–331, 2010.
- J. F. Gomes, L. Kogan, and M. Yogo. Durability of Output and Expected Stock Returns. *Journal of Political Economy*, 117(5):941–986, 2009.
- J. González-uribe. Venture Capital and the Diffusion of Knowledge. *Mimeo*, 2012.
- R. Gouge and I. King. A Competitive Theory Employment Dynamics. *Review of Economic Studies*, 64:1–22, 1997.
- J. R. Green and S. Scotchmer. On the Division of Profit in Sequential Innovation. *The RAND Journal of Economics*, 26(1):20–33, 1995.
- Z. Griliches. Issues in assessing the contribution of research and development to productivity growth. *The Bell Journal of Economics*, 10(1):92–116, 1979.
- Z. Griliches. Market Value, R&D, and Patents. *Economics Letters*, 7:183–187, 1981.
- Z. Griliches. Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature*, 28(4):1661–1707, 1990.
- Z. Griliches. Introduction to "Output Measurement in the Service Sectors". In Z. Griliches, editor, *Output Measurement in the Service Sectors*, pages 1–22. University of Chicago Press, 1992.
- Z. Griliches, B. H. Hall, and A. Pakes. R&D, Patents, and Market Value Revisited: Is There A Second (Technological Opportunity) Factor? *Economics of Innovation and New Technology*, 1(3):183–201, 1991.
- B. H. Hall. R&D Tax Policy during the Eighties: Success or Failure? 1992.
- B. H. Hall and J. Van Reenen. How Effective are Fiscal Incentives for R&D? A New Review of the Evidence. *Research Policy*, 29:449–469, 2000.
- B. H. Hall, A. B. Jaffe, and M. Trajtenberg. The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools. *NBER working paper*, 2001.
- B. H. Hall, A. Jaffe, and M. Trajtenberg. Market Value and Patent Citations. *The RAND Journal of Economics*, 36(1):16–38, 2005.
- R. E. Hall. Recessions as Reorganizations. In *NBER Macroeconomics Annual*. 1991.
- J. D. Hamilton. A Neoclassical Model of unemployment and the Business Cycle. *Journal of Political Economy*, 96(3):593–617, 1988.

- A. X. He. Technology Spillovers through the Innovation Network. *Mimeo*, 2015.
- M. Hirukawa and M. Ueda. Venture capital and industrial "innovation". *Mimeo*, 2013.
- S. Hsieh. Profit Windfalls and Investments in Innovation : Evidence from Patent Term Extensions. *2017*.
- A. B. Jaffe. Technological Opportunity and Spillovers of R&D : Evidence from Firms' Patents , Profits and Market Value. *NBER Working Paper*, 1986.
- A. B. Jaffe, M. Trajtenberg, and R. Henderson. Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics*, 108(3):577–598, 1993.
- N. Jaimovich and H. E. Siu. The Trend is the Cycle: Job Polarization and Jobless Recoveries. *Mimeo*, 2014.
- C. I. Jones. R&D-Based Models of Economic Growth, 1995.
- C. I. Jones and J. C. Williams. Too much of a good thing? The economics of investment in R&D. *Journal of Economic Growth*, 85(March):65–85, 2000.
- B. Jovanovic. Work, Rest, and Search: Unemployment, Turnover, and the Cycle. *Journal of Labor Economics*, 5(2):131–148, 1987.
- E. W. Kitch. The Nature and Function of the Patent System. *The Journal of Law and Economics*, 20(2):265–290, 1977. ISSN 02729490. doi: 10.1525/sp.2007.54.1.23.
- F. Kleibergen and R. Paap. Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1):97–126, 2006. ISSN 03044076. doi: 10.1016/j.jeconom.2005.02.011.
- P. Klemperer. How broad should the scope of patent protection be? *Rand Journal of Economics*, 21(1), 1990.
- T. J. Klette and S. Kortum. Innovating Firms and Aggregate Innovation. *Journal of Political Economy*, 112(5):986–1018, 2004.
- K. Koenders and R. Rogerson. Organizational dynamics over the business cycle: A view on jobless recoveries. *Federal Reserve Bank of St. Louis Review*, 87(4):555–579, 2005.
- S. Kortum and J. Lerner. Assessing the contribution of venture capital to innovation. *RAND Journal of Economics*, 31(4):674–692, 2000. ISSN 07416261. doi: 10.2307/2696354.
- M. a. Lemley and B. Sampat. Examiner Characteristics and Patent Office Outcomes. *Review of Economics and Statistics*, 94(3):817–827, 2012. ISSN 0034-6535. doi: 10.1162/REST_a_00194.

- J. Lemus and G. Marshall. When the Clock Starts Ticking : Measuring Strategic Responses to TRIPS ' s Patent Term Change. 2017.
- R. C. Levin, A. K. Klevorick, R. R. Nelson, S. G. Winter, R. Gilbert, and Z. Griliches. Appropriating the Returns from Industrial Research and Development. *Brookings Papers on Economic Activity*, 3(Special Issue on Microeconomics):783–831, 1987.
- X. Li, M. Macgarvie, and P. Moser. Dead Poets' Property - How Does Copyright Influence Price? 2016.
- D. M. Lilien and R. E. Hall. Chapter 17 Cyclical fluctuations in the labor market. *Handbook of Labor Economics*, 2(C):1001–1035, 1986.
- K. Lim. The many faces of absorptive capacity: Spillovers of copper interconnect technology for semiconductor chips. *Industrial and Corporate Change*, 18(6):1249–1284, 2009.
- P. Loungani and R. Rogerson. Cyclical fluctuations and sectoral reallocation. Evidence from the PSID. *Journal of Monetary Economics*, 23(2):259–273, 1989.
- S. Lychagin, J. Pinkse, M. E. Slade, and J. V. Reenen. Spillovers in Space: Does Geography Matter? *Journal of Industrial Economics*, 64(2):295–335, 2016. ISSN 14676451. doi: 10.1111/joie.12103.
- E. Manresa. Estimating the Structure of Social Interactions Using Panel Data. *Mimeo*, 2016.
- A. Mehta, M. Rysman, and T. Simcoe. Identifying the Age Profile of Patent Citations: New Estimates of Knowledge Diffusion. *Journal of Applied Econometrics*, 25:36–37, 2010.
- F. Mezzanotti. Roadblock to Innovation: The Role of Patent Litigation in Corporate R&D. *Mimeo*, 2015.
- A. Mian and A. Sufi. What Explains the 2007-2009 Drop in Employment? *Econometrica*, 82(6):2197–2223, 2014.
- P. Moser and A. Voena. Compulsory Licensing-Evidence from the Trading with the Enemy Act. *American Economic Review*, 102(1):396–427, 2012. ISSN 0002-8282. doi: 10.1257/aer.102.1.396.
- F. Murray, P. Aghion, M. Dewatripont, J. Kolev, and S. Stern. Of Mice and Academics : Examining the Effect of Openness on Innovation. *American Economic Journal: Economic Policy*, 8(1):212–252, 2016.
- A. Nagaraj. Does Copyright Affect Reuse? Evidence from Google Books and Wikipedia. 2016.
- R. R. Nelson. The Simple Economics of Basic Scientific Research. *Journal of Political Economy*, 67(3):297–306, 1959.

- W. D. Nordhaus. *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. MIT Press, 1969.
- M. Notowidigdo. The Incidence of Local Labor Demand Shocks. *Working Paper*, (November 2009), 2013.
- A. Pakes and M. Schankerman. The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources. *R & D, Patents, and Productivity*, pages 73–88, 1984. doi: 10.3386/w0346.
- J. R. Pierce and P. K. Schott. The Surprisingly Swift Decline of U.S. Manufacturing Employment. *American Economic Review*, 106(7):1632–1662, 2016.
- L. Pilossoph. Sectoral Shocks and Move Unemployment. 2012.
- J. D. Rauh and A. Sufi. Explaining corporate capital structure: Product markets, leases, and asset similarity. *Review of Finance*, 16(1):115–155, 2012.
- P. Restrepo. Skill Mismatch and Structural Unemployment. *Mimeo*, 2015.
- M. Rhodes-Kropf, D. T. Robinson, and S. Viswanathan. Valuation waves and merger activity: The empirical evidence. *Journal of Financial Economics*, 77(3):561–603, 2005.
- P. M. Romer. Endogenous Technological Change. *Journal of Political Economy*, 98(5):S71–S102, 1990.
- B. Sampat and H. L. Williams. How do patents affect follow-on innovation? Evidence from the human genome. *NBER Working Paper*, 2015.
- M. Schnitzer and M. Watzinger. Measuring Spillovers of Venture Capital. *Mimeo*, 2015.
- S. Scotchmer. Standing on the Shoulders of Giants: Cumulative Research and the Patent Law. *Journal of Economic Perspectives*, 5(1):29–41, 1991.
- J. H. Stock and M. Yogo. Testing for Weak Instruments in Linear IV Regression. In D. Andrews and J. H. Stock, editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press, Cambridge, 2005. ISBN 780521844413. doi: citeulike-article-id:1459542.
- N. U. Sukhatme and J. N. L. Cramer. How Much Do Patent Applicants Care About Patent Term ? Cross-Industry Differences in Term Sensitivity. 2014.
- C. Syverson. What Determines Productivity? *Journal of Economic Literature*, 49(2):326–365, 2011. ISSN 0022-0515. doi: 10.1257/jel.49.2.326.
- C. Van Horn. Effects of GATT and NAFTA on PTO Practice. *Journal of Patent and Trademark Office Society*, March, 1995.

- M. Warusawitharana and T. M. Whited. Equity Market Misvaluation, Financing, and Investment. *Review of Financial Studies*, 29(3), 2016.
- H. L. Williams. Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy*, 121(1):1–27, 2013.
- D. J. Wilson. Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits. *Review of Economics and Statistics*, 91(May):431–436, 2009.
- D. Yagan. Moving to Opportunity? Migratory Insurance Over the Great Recession. 2014.
- R. H. Ziedonis. Don't Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms. *Management Science*, 50(6):804–820, 2004. ISSN 0025-1909. doi: 10.1287/mnsc.1040.0208.