

Essays on Auctions, Contests, and Games

by

Vivek Bhattacharya

B.S. Economics and Physics, Duke University (2012)

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

© Vivek Bhattacharya, 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.

Author . **Signature redacted**
Department of Economics
June 1, 2017

Certified by. **Signature redacted**
Glenn Ellison
Gregory K. Palm Professor of Economics
Thesis Supervisor

Certified by **Signature redacted**
Nikhil Agarwal
Castle Krob Career Development Assistant Professor of Economics
Thesis Supervisor

Accepted by **Signature redacted**
Ricardo Caballero
Ford International Professor of Economics
Chairman, Departmental Committee on Graduate Studies



Essays on Auctions, Contests, and Games

by

Vivek Bhattacharya

Submitted to the Department of Economics
on June 1, 2017, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis consists of three chapters broadly in industrial organization, with a focus on contests and auctions, and game theory. Chapter 1 develops a new model of multistage R&D procurement contests, in which firms conduct research over a number of stages to develop an innovative product and then supply it to a procurer. I show that the primitives of this model—the cost of research, the distributions of project values and delivery costs, and the share of the profits captured by the firms—are nonparametrically identified given data on R&D expenditures and procurement contract amounts. I then develop a tractable estimation procedure and apply it to data from the Small Business Innovation Research program in the Department of Defense. I find that within a particular contests, there is low variation in the values of the proposed projects, which are drawn early in the process, but considerably larger variation in the delivery costs, which are drawn later. The DOD provides high-powered incentives, sharing about 75% of the surplus with the firms. I then suggest simple design changes to improve social surplus but find that many of these socially beneficial design changes would in fact reduce DOD profits.

Chapter 2, which is joint with James Roberts and Andrew Sweeting, studies the benefits of regulating entry into procurement auctions, relative to standard auctions in which bidders are allowed to enter and bid freely. Specifically, we study the relationship between auction outcomes and the precision of information bidders have about their costs before entering the bidding stage of the contest. We show that the relative performance of a standard auction with free entry and an “entry rights auction,” which restricts participation in the bidding phase, depends nonmonotonically on the information precision. We finally estimate the model on a dataset of auctions for bridge-building contracts let by the Oklahoma and Texas Departments of Transportation. Entry is estimated to be moderately selective, and the counterfactual implication is that an entry rights auction would significantly increase social efficiency and reduce procurement costs.

Chapter 3, which is joint with Lucas Manuelli and Ludwig Straub, proposes a model of “signal distortion” in a game with imperfect public monitoring. We construct a framework in which each player has the chance to distort the true public signal, and each player is uncertain about the distortion technologies available to his opponent. Continuation payoffs are dependent on the distorted signal. Our main result is that when players evaluate strategies according to their worst-case guarantees—i.e., are ambiguity-averse over certain distributions in the environment—players behave as if the continuation payoffs that incentivize them in the stage game are perfectly aligned with their opponents’. We then provide two examples showing counterintuitive implications of this result: (i) signal structures that allow players to identify deviators can be harmful in enforcing a strategy profile, and (ii) the presence of signal distortion can help sustain cooperation when it is impossible in standard settings. We then extend our equilibrium concept to a repeated game, show that it is a natural generalization of strongly symmetric equilibria, and then prove an anti-folk theorem that payoffs are in general bounded away from efficiency.

Thesis Supervisor: Glenn Ellison
Title: Gregory K. Palm Professor of Economics

Thesis Supervisor: Nikhil Agarwal
Title: Castle Krob Career Development Assistant Professor of Economics

Acknowledgments

I had always intended to write a short acknowledgements section. But after five years of grad school, I realize that doing so would give me too much credit for this dissertation.

I imagine it's difficult to find an advisor as versatile as Glenn Ellison. He was always willing to discuss so many different parts of the research process—extremely preliminary ideas, presentation techniques, details of a proof, interview questions, empirical specifications, responding to referee comments, and anything else. I was constantly impressed by the breadth of his knowledge: his comments on ideas about repeated games were just as insightful as his ones on empirical contests. Coupled with the fact that his door was literally always open, even as department head, Glenn was always my go-to person for any question. But at the end of the day, I liked his directness the most: his most common bit of advice for me was “Just write papers.”

One of the luckiest points in grad school was when Nikhil Agarwal emailed me during my second year to ask whether I wanted to work as an RA for him. I was unsure at the time, since I was still contemplating macro as a field, but I am very glad I agreed. Nikhil has an amazing balance of focusing on the details of a paper—exactly how to implement the estimation procedure, details of the formal identification argument, specifics about the descriptive analysis of the dataset—while simultaneously asking whether the big picture question is interesting; he encouraged me to think more broadly than just my empirical setting. Exceptionally generous with his time, he was always willing to meet on short notice and a phone call away during the market. I bothered him with all sorts of unusually specific questions, both academic and nonacademic. More broadly, he created a very cohesive IO community among the grad students, and the structural reading group he organized was an important part of my experience at MIT. Having Nikhil as a mentor has provided me a role model to try to emulate in the next stage of my career.

While Mike Whinston is formally my third advisor, he was just as integral a part of my development as Glenn and Nikhil. I never had the chance to take a class from him, but much of how I think about IO has been shaped by our incredibly long meetings. Mike thinks through each part of an argument (theoretical or empirical) extremely carefully and does not move on to the next step until he has analyzed the current one from multiple angles. Having the chance to witness his thought process has been a wonderful experience; he has a somewhat distinctive manner of thinking about the connection between theory and empirics that remains faithful to both, and I hope this has rubbed off on me to some extent. He also went through my job market paper with a fine-toothed comb, once again thinking through essentially every sentence, and his comments improved it greatly. Finally, he never shied away from giving brutally honest feedback about my paper, my presentation, and how I come across in academic settings. I always appreciated that.

Many other faculty members at MIT have been instrumental in helping me grow as an economist. Paulo Somaini was an essential source of support early on at my time at MIT; I learned much of what I know about structural econometrics from his course and from serving as an RA from him, and he was always generous with his time even as he moved to Stanford. Sara Fisher Ellison had helpful comments about every one of my talks and helped me think beyond the particular setting in each paper. While I wish I had had more opportunities to interact with Nancy Rose, the training I received in her antitrust course had always inspired me to stay attuned to the actual institutions I was studying. Juuso Toikka is one of the friendliest, sharpest, and most approachable people I have ever met, and he was always open to discussing theory ideas even when it became clear that I was not on the path to becoming a theorist. My friends and I would often joke that Juuso is who we want to be when we grow up, and I'm hoping I can follow through on that ideal at Northwestern. Finally, I enjoyed Iván Werning's classes and reading groups quite a lot when I wanted to become a macroeconomist, although it is probably best for everyone that I did not.

I would definitely not be in this position had it not been for Jimmy Roberts and Andrew Sweeting, my undergraduate advisors at Duke. They introduced me to research in economics, indulging my questions and patiently dealing with all the mistakes I made. Working for them was what convinced me to go to grad school in economics. Even after I left Duke, they have been ready sources of all sorts of advice. Outside my immediate family, it is hard to point to people who have had quite as much an impact on my life as they have had. I feel fortunate that we have remained in touch for so long, and I look forward to interacting with them for many years to come. Indeed, Duke is responsible for putting me in touch with many fantastic role models: Matt Rognlie has been a constant source of support and advice since we were undergrads at Duke (his emails in response to my questions are often as long as the chapters in this dissertation), and it's rare find such a wonderful peer mentor. I feel lucky to call him a good friend, and I'm excited to be following him to a different city once more.

MIT is a great place to be a student not just because of the faculty but also because of the friendliness and dedication of the staff. Cherisse Haakonsen, Emily Gallagher, and especially Kim Scantlebury have been friendly faces throughout my time here. Cherisse and Kim helped me get through the job market (not just logistically). More importantly, MIT wouldn't have been the same without my hallway conversations with them, which I enjoyed very much. Gary King kept us all on track, and I'll miss his emails.

The day-to-day experience in grad school is of course shaped by classmates rather than faculty, and I have been incredibly fortunate in that dimension as well. I have interacted with many wonderful IO students—Gastón Illanes (who was instrumental in helping me navigate the year before the market), Sarah Moshary, Manisha Padi, David Colino, Hongkai Zhang, Mengxi Wu, Dan Waldinger, and many others—across cohorts at MIT. I also had a wonderful group of classmates outside IO who gave feedback and suggestions at every step of the research process: Ben Roth, David Choi, Ludwig Straub, Sebastián Fanelli, Greg Howard, Jonathan Libgober, Ernest Liu, Arianna Ornaghi, Dan Green, Yan Ji, Scott Nelson, Bryan Perry, Peter Hull, and the list goes on. Perhaps more importantly, many of these classmates have become exceptionally close friends. As evidenced by many late-night IMs, early morning basketball games, and frequent board game sessions (yes, we're super cool), it would have been extremely difficult to navigate the ups and downs of grad school without this group. I should devote more than just a paragraph to them, but I wouldn't know where to stop.

Finally, none of this would be possible without the constant support of my parents and my brother. They were with me both when things looked bleak, and I needed someone to talk to late at night, and when the future looked brighter, and I wanted someone to share my excitement with. They have all sacrificed a lot to help me succeed, and I'm not sure I have ever vocalized how much I appreciate everything they do. I don't know how much of the dissertation they will actually read, but I hope they read this.

Contents

1	An Empirical Model of R&D Procurement Contests	15
1.1	Introduction	16
1.1.1	Related Literature	18
1.2	Empirical Setting and Data	21
1.2.1	Overview of the Navy SBIR Program	21
1.2.2	Data Sources	23
1.2.3	Descriptive Statistics	25
1.3	Model	31
1.3.1	Model Timing and Primitives	31
1.3.2	How Are Research Efforts Determined?	32
1.4	Identification and Estimation	37
1.4.1	Identification	37
1.4.2	Empirical Model	41
1.4.3	Estimation Procedure	43
1.5	Structural Estimates	48
1.6	Social Inefficiency in R&D Contests	51
1.7	The Effect of Early- and Late-Stage Competition	54
1.7.1	Changing N_1 and \bar{N}_2	55
1.7.2	Decomposing the Effect of Competition	56
1.8	The Effect of the Bargaining Parameter	58
1.9	Decoupling Research from Delivery: Prizes and IP Sharing	60
1.9.1	A Model of IP Sharing	61
1.9.2	The Costs of Prizes and IP Sharing	63
1.10	DOD Profits Under Alternate Contest Designs	65
1.11	Conclusion	67
2	Regulating Bidder Participation in Auctions	71
2.1	Introduction	71
2.2	A Model of Endogenous Bidder Participation in Auctions	76
2.2.1	Selective Free Entry into First-Price Auctions	77
2.2.2	Entry Rights Auctions	78

2.2.3	Efficiency of Bidder Entry	81
2.3	Data	84
2.4	Estimation	85
2.4.1	Model Setup and Specification	86
2.4.2	Importance Sampling	87
2.4.3	Identification	89
2.5	Results	93
2.5.1	Parameter Estimates	94
2.5.2	Regulating Bidder Entry with an Entry Rights Auction	95
2.6	Conclusion	102
3	Imperfect Public Monitoring with a Fear of Signal Distortion	105
3.1	Introduction	105
3.2	A One-Shot Game	108
3.2.1	Setup	108
3.2.2	Distortion Equilibrium	111
3.3	Properties of Distortion Equilibria	113
3.3.1	Linearity	113
3.3.2	Existence and Multiplicity	115
3.3.3	Discussion	117
3.4	Examples	119
3.4.1	Identifying Deviators Can Be <i>Harmful</i>	119
3.4.2	Cooperating Can Be <i>Easier</i>	121
3.5	An Infinitely Repeated Game	123
3.6	APS-Style Characterization of Perceived Payoffs	126
3.6.1	Recursive Characterization of Perceived Payoffs	126
3.6.2	Totally Linear PPEs	127
3.6.3	Connecting Payoffs of PPEDs and Totally Linear PPEs	129
3.6.4	An Anti-Folk Theorem	130
3.6.5	Examples of Q_P	130
3.7	Conclusion	132
A	Appendix to Chapter 1	135
A.1	Further Details on DOD Profits Under Changes to the Contest Design	135
A.1.1	Changing the Number of Competitors	135
A.1.2	Changing η : The Procurer's Laffer Curve	136
A.1.3	Mandating IP Sharing	139
A.2	Additional Empirical Results	141
A.2.1	Details on the Descriptive Statistics	141
A.2.2	Structural Estimates without Assumption O	144

A.3	Extensions to the Identification Result	147
A.3.1	Bounding η Using Assumption M	147
A.3.2	Asymmetric Firms	148
A.3.3	Unobserved Heterogeneity	148
A.3.4	External Benefits of Research	151
A.4	Incentive Compatibility in the Model of IP Sharing	151
A.5	Omitted Proofs	152
A.5.1	Proof of Propositions 1.2 and 1.3	153
A.5.2	Proof of Proposition 1.5	154
A.6	Data Appendix	155
A.6.1	SBIR Data from the Office of Naval Research	155
A.6.2	Federal Procurement Data System	155
A.6.3	SBIR Solicitations	156
A.7	Computational Methods	157
A.7.1	Computing and Optimizing the Likelihood Function	157
A.7.2	Inverting the Characteristic Function	158
A.7.3	Solving for the Equilibrium Effort Function	159
B	Appendix to Chapter 2	161
B.1	Numerical Methods to Solve the First-Price Auction	161
B.2	Details of Sample Construction	161
B.3	Monte Carlo Studies	162
C	Appendix to Chapter 3	165
C.1	Formally Adding Ambiguity Aversion	165
C.1.1	Sequential Equilibrium	165
C.1.2	Adding Ambiguity Aversion	166
C.2	Extension to N Players	168
C.3	Omitted Proofs	172
C.3.1	Proof of Theorem 3.2	172
C.3.2	Proofs in Section 3.6.1	173
C.3.3	Proof of Theorem 3.4	174
C.3.4	Proof of Theorem 3.5	175
C.3.5	Proof of Theorem 3.6	176

List of Figures

1-1	Distribution of contract amounts in Phases II and III	27
1-2	Measures of model fit	52
1-3	Social surplus as a function of η	59
1-4	Social surplus with prizes and IP sharing	64
2-1	Relative performance of an entry rights auction and a first-price auction with free entry	83
2-2	Simulated outcomes for various values of K and α	90
2-3	Sensitivity analysis following Gentzkow and Shapiro (2013)	92
2-4	Comparing the winner's cost and the total procurement costs in an FPAFE an an ERA	99
2-5	Bid functions for various auction designs	101
3-1	An extensive form representation of the first-stage game	110
3-2	Illustration of consistent triples (w, \tilde{w}, D)	116
3-3	Using a proxy in the prisoner's dilemma	122
3-4	Examples of the limit sets of payoffs in the prisoner's dilemma for various signal structures	132
A-1	Procurer's Laffer curve	139
A-2	DOD profits as a function of prizes and IP sharing	140
A-3	Nonparametric regressions of success rates on (relative) funding	144

List of Tables

1.1	Representative topics generated by the LDA algorithm	24
1.2	Summary statistics for contests from 2000 to 2012	26
1.3	Distribution of the number of competitors in each phase	26
1.4	Descriptive regressions of outcomes on measures of competition	28
1.5	Structural estimates	49
1.6	Improvement in social surplus when correcting for inefficiencies in individual phases	53
1.7	Effect of changing the number of competitors	55
1.8	Decomposing the effect of competition on social surplus	57
1.9	Socially optimal levels of η	59
1.10	Optimal levels of prizes	64
1.11	Social and DOD outcomes for various contest designs	66
2.1	Summary statistics of the sample of auctions	84
2.2	Importance sampling density parameters ($\tilde{\Gamma}$) for estimation	89
2.3	Parameter estimates	94
2.4	Relative social and procurement costs of an FPAFE and an ERA at the estimated parameters	97
2.5	Breakdown of the expected values of major quantities in a FPAFE and an ERA	104
A.1	Effect of competition on DOD outcomes	136
A.2	Decomposing the effect of competition on DOD outcomes	137
A.3	Optimal values of η from the DOD's perspective	138
A.4	DOD-optimal prizes with and without IP sharing	140
A.5	Robustness analysis of descriptive regressions of contest-level outcomes on measures of competition	142
A.6	Robustness analysis of descriptive regressions of individual-level outcomes on measures of competition	143
A.7	Robustness analysis of descriptive regressions of award amounts on measures of competition	145
A.8	Structural estimates, conditional on various values of η	146
B.1	Monte Carlo experiments for the importance sampling estimator	164

Chapter 1

An Empirical Model of R&D Procurement Contests

This chapter develops a new empirical model of a multistage R&D contest for a procurement contract and uses it to study the design of the Small Business Innovation Research program in the Department of Defense. Firms' incentives to innovate depend on the cost of research, the intensity of competition, and the rewards from securing the procurement contract. The cost of research, the distributions of project values and delivery costs, and the fraction of the surplus shared by the procurer are nonparametrically identified and can be tractably estimated using data on the procurement contract amount and the firms' R&D expenditures. Estimates suggest that there is fairly low variation in the values of projects developed by different firms and that most of the variation in the procurement contract is attributable to differences in delivery costs, which are drawn later in the research process. Further, the DOD currently provides high-powered incentives, sharing approximately three-quarters of the surplus from the innovation with the supplier. Increasing the number of competitors in later stages of the contest, lowering the share of the surplus firms receive in procurement, and mandating that firms share intellectual property would all increase total social surplus. However, because the DOD pays for research expenditures but only partially internalizes the gains from improved innovations, many socially beneficial design changes would actually reduce its profits from the contest. Together, these results suggest that at the estimated parameters, the DOD may have an incentive to skew the design of the contest significantly away from the socially optimal one.

I would like to thank Nikhil Agarwal, Glenn Ellison, and Mike Whinston for invaluable advice and support throughout this process. This chapter benefited from discussions with Daron Acemoglu, Attila Ambrus, David Choi, Allan Collard-Wexler, Sara Ellison, Sebastián Fanelli, Greg Howard, Yan Ji, Ernest Liu, Manisha Padi, Parag Pathak, Harry Di Pei, Ben Roth, Paulo Somaini, Ludwig Straub, Juuso Toikka, Dan Waldinger, and especially Gaston Illanes and Jimmy Roberts. I would also like to thank John Williams and G. Nagesh Rao at the Small Business Administration and Robert Smith at the Office of Naval Research for helpful conversations related to the institutions and data sources, as well as Gokul Subramanian for introducing me to this empirical setting. Finally, I am extremely grateful to Jonathan Libgober, whose advice and comments were instrumental in the early stages of this project. This material is based on work supported by the NSF Graduate Research Fellowship under Grant No. 1122374 as well by the John Krob Castle 1963 Fellowship and the George P. and Obie B. Shultz Fund. All errors are mine.

1.1 Introduction

Motivated by settings as diverse as the patent system, grand challenges that offer prizes for particular breakthroughs, and grants for academic research, economists have long studied mechanisms to incentivize research and development and procure innovative products and services.¹ An especially large player in the market for both funding R&D and procuring products that require R&D is the government: of the approximately \$450 billion obligated in federal contracts in FY2014, about 10% was for contracts for R&D, which accounted for about 9% of all R&D expenditures worldwide. The Department of Defense in particular spent \$28 billion on R&D, more than all other government agencies combined (Schwartz, Ginsberg, and Sargent, 2015). Firms with R&D contracts are often incentivized through mechanisms that resemble contests: multiple firms conduct research on similar products, and the procuring agency contracts with one of the firms for delivery or purchases the rights to use the plans in production. Yet, despite both the sizable body of theoretical work on R&D contests and their importance in many real-world settings, there has been little empirical analysis to understand the underlying primitives of these contests. In this chapter, I develop a structural model of R&D contests, provide a methodology for identification and estimation of the model parameters, and study the effect of both competition and contest design on procurement outcomes in the context of contests run by the U.S. Department of Defense.

In the design of procurement contests, a central question allows the degree of competition to allow. In a standard procurement setting, adding competition is unambiguously beneficial for social surplus and the procurer’s profit; Bulow and Klemperer (1996) provides an especially strong result that adding competitors is even preferable to setting an optimal reserve. In settings with R&D, however, an additional consideration comes into play. Although introducing an additional competitor does increase the chance of a successful innovation and can also directly reduce the price the procurer pays for the innovation, each competitor in the contest may reduce their research effort, anticipating a lower expected reward. Taylor (1995) considers this tradeoff in a stylized model of contests with a fixed prize, and Fullerton and McAfee (1999) and Che and Gale (2003) propose related models with the starker result that restricting the number of competitors to exactly two is optimal. Therefore, unlike in standard procurement, the impact of competition on outcomes of R&D contests is an empirical question.

In this chapter, I study multistage research contests in which successful research is awarded with procurement contracts. In an empirical setting, I investigate three main mechanisms to control competition in these contests. First, I study the “extensive margin” of competition by investigating the optimal number of early-stage and late-stage competitors that the procurer—in this case, the DOD—should admit to the contest. In doing so, I decompose the effect of competition into the direct effect of adding competitors and the indirect *incentive effect* of allowing these competitors to change their research effort. I find that the social planner would like to admit a large number of competitors into both phases of the contest whereas the DOD prefers to restrict competition

¹See Maurer and Scotchmer (2004) and Cabral, Cozzi, Denicoló, Spagnolo, and Zanza (2006) for reviews of the literature on R&D procurement. Williams (2012) discusses innovation contests in particular.

severely; the incentive effect is usually beneficial for social surplus and DOD profits, but the DOD does not capture much of the direct effect of adding late-stage competitors. Second, I consider the “intensive” margin of competition, which is modulated by the portion of the generated surplus the procurer allows the firms to capture in the final contracting stage. The procurer trades off incentives for surplus generation with the proportion of the surplus it captures, and it thus faces a natural “Laffer” curve. I show evidence that the current design is on the efficient side of the Laffer curve. However, giving the firms slightly less of the surplus can improve social surplus by discouraging excessive R&D. Finally, I consider changes in the prize structure, first to partially decouple the early-stage incentives for research from the final procurement contract and then to study the benefits of sharing intermediate research breakthroughs. Decoupling research and delivery is always socially beneficial at the estimated parameters but may reduce DOD profits. These counterfactuals suggest that at the estimated parameters, the social planner and the DOD have starkly conflicting incentives for optimal design.

I study these design counterfactuals by developing a model of multistage R&D contests that captures the salient features of my empirical setting: the Small Business Innovation Research (SBIR) program in the Department of Defense. The DOD spends over \$1 billion a year on R&D contracts through this program and almost \$500 million on delivery contracts generated from research funded by this program. It solicits research on technologies related to all major defense acquisition programs. While it is thus an important program in its own right, this program also provides a controlled setting to study multistage R&D contests. In the SBIR program, a set of firms conducts preliminary work to develop initial plans for a specific product. I model this “research” phase as one in which firms exert effort to generate a successful innovative idea and learn its value to the DOD. In the second phase of the SBIR program, the most promising firms then receive contracts to make these plans commercially viable. I model this phase as a “development” phase in which firms choose how much effort to exert based on the value of their particular project, and they receive a draw of a delivery cost from some distribution based on this effort. A firm is successful at developing the project if the draw of the delivery cost is lower than the value the project provides to the DOD. The DOD contracts with at most one of these successful firms for delivery. In my model, this contract amount is set via a natural extension of Nash bargaining. This timeline—a multistage innovation process followed by commercialization or contracting—is representative of many settings of R&D procurement.

I then show that the underlying parameters of the model—the distribution of values and costs, the stochastic map from research effort to the cost draws, and the bargaining parameter—are identified from data on the amount spent on research as well as the delivery contract amounts. I provide a constructive identification proof to make the argument transparent, and the key conditions are relatively weak: because firms with higher-value projects have more of an incentive to exert effort, and because the DOD would presumably never purchase a project whose delivery cost exceeds its value, all parameters but the bargaining parameter are nonparametrically identified. The condition that the research effort is set optimally then identifies the bargaining parameter. This nonparametric identification is robust, and I show that it can be extended to many generalizations

of the model I consider in this chapter.

An added benefit of the constructive identification proof is that it leads to a natural estimation procedure. I propose a multistep estimator that has two main benefits. The practical benefit is that the estimation procedure avoids having to explicitly solve a model of R&D contests for many different parameter values, which can be computationally burdensome. The conceptual benefit is that the estimation procedure allows the researcher to be agnostic about the actual process that determines the effort schedule (i.e., the map from values to research efforts) and thus allows the procedure to apply to a variety of settings instead of being specific to this model. If the researcher has external knowledge of parameters of the contracting process—the bargaining parameter, in this chapter—then this first step is sufficient to estimate all the primitives of the process without imposing any structure on the effort schedule. In later steps of the estimation procedure, I leverage specifics of this empirical setting—in particular, that the firm chooses the research efforts optimally—to estimate the bargaining parameter by matching certain moments in the data. Furthermore, the procedure is designed to control for both observed as well as unobserved heterogeneity that affects both values and costs, borrowing techniques from the literature on auctions (Li and Vuong, 1998; Krasnokutskaya, 2011).

Estimates from the model indicate that the DOD values successful projects at an average of \$11–\$15 million, and the DOD tends to invite more competitors to contests that it finds more valuable. The within-contest variation in values is fairly small: a competitor with a project at the 97.5th percentile of the value distribution has a value that is only about 12% larger than the value at the 2.5th percentile. Most of the final variation in contract amounts comes from variation in delivery costs drawn in the development phase. Finally, the estimates suggest that firms capture about three-quarters of the surplus generated by the program. I then discuss and quantify the inefficiencies inherent in R&D contests. I show that research in the later phase is underprovided due to a holdup effect. Research in the early stages, however, is overprovided due to a combination of a business-stealing effect and a reimbursement effect that stems from the DOD’s practice of refunding later-stage research costs. The identification argument I provide allows me to clearly comment on the patterns in the data that lead to these estimated parameters.

These estimates also allow me to investigate the nature and magnitude of the inefficiencies in this contest. I find that effort is underprovided in the second stage of the contest due to the holdup effect; the cost of this holdup is fairly low, however, and removing it improves social efficiency by 5–10%. Effort is overprovided in the early stage of the contest due to the potential for business stealing and the reimbursement effect, and social efficiency can be improved by as much as 22%. These social inefficiencies are informative by themselves, but they also feed into the analysis of the costs and benefits of alternate contest designs, as discussed above.

1.1.1 Related Literature

The conceptual framework for this chapter is based on the theoretical literature on R&D contests, which stresses the tension between the direct effect of adding another competitor—both in terms

of the added chance of success as well as the increase in total research costs—with the indirect *incentive effect* on the research efforts. The salient conclusion of this literature is the importance of restricting entry. Taylor (1995) considers both “research” contests (in which the best competitor at the end of the contest wins a prize) and “innovation” contests (in which the first competitor to achieve a desired level of innovation wins the prize) and notes that restricting entry in these contests to finitely many competitors can be desirable to counteract the incentives to reduce effort. In a related model, Fullerton and McAfee (1999) extend this insight by considering agents with heterogeneous research costs and studying a mechanism that auctions entry into tournaments. They arrive at the strong conclusion that it is always optimal to restrict entry to two competitors. Che and Gale (2003) focus on the unverifiability of innovation and allow contestants to bid for the prize *before* an innovation stage in order to guarantee economic profits to incentivize innovation; once again, however, they find that it is optimal to restrict entry to just two competitors at the innovation stage.²

There are a number of differences between my empirical setting and the baseline models of R&D contests studied by these papers. First, the incentives in my setting come from a procurement contract instead of a fixed prize. In this sense, Che and Gale (2003) consider the closest incentive scheme to my setting. More recently, Che, Iossa, and Rey (2016) study optimal design in a much more related setup in which innovations have heterogeneous values and costs and the mechanism by which the delivery contract is awarded is chosen to incentivize preliminary research. Second, my setting is explicitly a multistage process in which breakthroughs (or draws of values and costs) happen sequentially and a successful innovation requires successes in both stages. In this sense, progress on values influences the effort exerted on minimizing costs, and the setting is related to the literature on R&D races.³ Furthermore, a more recent theoretical literature has studied incentives in settings in which innovations require multiple breakthroughs themselves.⁴ To my knowledge, this chapter is the first to use the foundations of the literature on R&D contests to build and estimate a structural model of R&D procurement—a setting with costly effort and multistage progress.

The nontrivial interaction between competition and innovation has been of interest beyond the setting of R&D contests. Economists since Schumpeter (1939) and Arrow (1962) have discussed whether firms with large market power also have more incentives to innovate. Not all the effects highlighted by the seminal papers in this literature as well as by the “quality ladder” models inspired by it⁵ are applicable to my setting of R&D procurement.⁶ Nevertheless, this literature is

²See also Fullerton, Linster, McKee, and Slate (2002) for further discussion of a setting where contest winners are rewarded by first-price auctions.

³For instance, “leaders” and “laggards” have differential incentives to conduct research. Papers in this literature include Harris and Vickers (1987) and Choi (1991).

⁴Toxvaerd (2006) studies delays in projects where multiple breakthroughs are required. Green and Taylor (2016) study incentive provision in a principal-agent model where two successes are required and the first success is unobserved. In somewhat related work, Biais, Mariotti, Rochet, and Villeneuve (2010) look at the “opposite” situation in which an agent with limited liability exerts effort to avoid large losses (instead of to generate breakthroughs).

⁵See, for instance, Aghion, Harris, and Vickers (1997), which is the basis for the empirical analysis of Aghion, Bloom, Blundell, Griffith, and Howitt (2005).

⁶The “replacement effect” from Arrow (1962) relies on a monopolist being unable to spread the cost of R&D over a large output. Gilbert and Newbery (1982) suggest an “efficiency effect” that depends on a monopolist being able

relevant to this chapter because much of the empirical work on the relation between innovation and competition involves cross-firm studies inspired by these theoretical models. Blundell, Griffith, and van Reenen (1999) document a positive relationship between innovation and market share. Aghion, Bloom, Blundell, Griffith, and Howitt (2005) show evidence of an inverted-U relationship between citation-weighted patents and the competition a firm faces. Acemoglu and Linn (2004) uncover a robust relationship between the market size—a proxy for the incentives to innovation—and R&D expenditures in the pharmaceutical industry.

A more recent empirical literature studies competition in the relatively new domain of online “ideation” contests, in particular for computer code and logo design. While this market is much smaller and does not lead to procurement contracts in the same sense as the DOD SBIR program, it does provide a controlled setting to study these effects as well as the opportunity for experimentation. Boudreau, Lacetera, and Lakhani (2011) use quasi-experimental variation in competition to show that effort reduction is the dominant force in contests with low uncertainty, and Boudreau, Lakhani, and Menietti (2016) show differential responses to competition by skill level. Gross (2016a) shows evidence of an inverted-U response of “originality” to competition.⁷ While my fundamental question is similar to the ones in these papers, the setting I consider is rather different. First, I focus on multistage contests. Second, projects in my setting differ on multiple dimensions (values and costs). Finally, the prize structure is different, as competitors are rewarded based on the surplus they generate rather than by a fixed prize based on their rank in a tournament.

This chapter contributes to a small academic literature on the SBIR program itself. Lerner (2000) and Howell (2016) document the long-term effects of the SBIR program and show that awardees have increased growth, higher revenues, and more patents than comparable firms that were not awarded grants. Lerner (2000) proposes that the differential growth is due to signaling firm quality, while Howell (2016) suggests it is due to funding early-stage prototyping. Wallsten (2000) uses data on internal financing to conclude that SBIR funds crowd out private investment in R&D dollar-for-dollar. Unlike these papers, I study the effects of competition within the SBIR program itself, and I also focus on an agency that uses this program as part of its procurement process rather than as a potential substitute for private R&D or venture capital funding.

Finally, I study the SBIR program in the Department of Defense rather than in other settings, and this chapter thus is related—albeit loosely—to the literature on defense procurement. A benefit of the DOD SBIR program relative to other instances of DOD procurement is that projects are much smaller in scope and the goals are well-specified, and asymmetric information about values and costs is arguably much less of an issue than in the procurement of major weapons systems from prime suppliers. However, the SBIR program retains many of the salient features of defense procurement as described in Rogerson (1994, 1995) and Lichtenberg (1995). Defense procurement

to preempt entry by engaging in R&D. More recently, Holmes, Levine, and Schmitz (2012) show that monopolists innovate less if they must spend time adopting new technologies.

⁷Other related papers include Gross (2016b), who studies the impact of performance feedback on outcomes in ideation contests, and Kireyev (2016), whose focus is more on the prize structure. While these concerns are in principle applicable to R&D for defense goods and services as well, they are more controlled in the settings of online ideation contests that these authors study.

involves contracting for both R&D and delivery, and the DOD often considers prototypes from multiple competitors before narrowing the competition for the delivery contract.⁸ Furthermore, procurement contracts are structured so that firms earn economic profits, thus providing them incentives for investment in early stages of the process (Rogerson, 1989). Finally, innovation and delivery can be decoupled, and the DOD may choose to contract with two separate firms for the two parts of the process. The counterfactuals I study in this chapter speak to all three of these methods for controlling incentives.

1.2 Empirical Setting and Data

1.2.1 Overview of the Navy SBIR Program

The Small Business Innovation Research program is a federal program designed to encourage small businesses to engage in R&D. The ultimate goal of this program is to provide these firms seed funding to commercialize early-stage research projects—either on the private market or, as will be primarily the case for this project, to the government. Any federal agency with an extramural R&D budget of more than \$100 million must allocate at least 2.8% of it competitively through this program to small businesses. This requirement encompasses eleven federal agencies, including the Department of Defense. I focus on the DOD—and in particular the Navy—because, unlike other federal agencies, it almost always solicits research on technologies that it wishes to acquire. Over 80% of the topics solicited by the Navy are developed by Program Execution Offices (PEOs) to meet specific needs of their acquisition programs.⁹ Furthermore, the market for technologies produced through the SBIR program is more limited than with other agencies—essentially to the DOD itself or to prime contractors through DOD contracts. Finally, the Navy keeps careful track of implementation and delivery contracts that result directly from R&D funded by SBIR, providing a way to track a technology from concept to acquisition. Note that while I focus on the Navy in this chapter, many of the institutional details described in this section are applicable to the entire DOD.

The DOD posts solicitations for specific research projects two to three times a year, with about 800 solicitations per year—between 150–250 each for the three main components of the DOD (i.e., the Army, the Navy, and the Air Force). These solicitations are publicly available and include a description of the required technology, often including relatively detailed technical requirements; goals for Phases I, II, and III; a discussion of possible commercialization potential; and references to both scientific literature and specific DOD liaisons for more information. The Navy in particular connects almost all solicitations to not just systems commands (e.g., the Naval Air Systems Command, or the Space & Naval Warfare Command) but also specific acquisition programs (e.g., the Joint Strike Fighter Program, or the Virginia Class Submarine Program). The

⁸The tradeoff between “early-stage” and “late-stage” competition is related to work on dual sourcing. Anton and Yao (1989, 1992) study theoretical models of dual sourcing and split award contests, and Lyon (2006) provides some empirical analysis of whether this can reduce costs.

⁹See <http://www.navybir.com/natconf14f/presentations/3-09-Navy-Comm-Williams.pdf>.

solicited products are fairly specific to military applications and often are smaller components of major weapons systems.¹⁰

Firms interested in competing for a Phase I contract must submit a 20-page technical proposal discussing a potential approach to meeting the goals of the solicitation as well as a detailed cost volume discussing how the firm will use the Phase I funding provided by the DOD. Upon evaluating these proposals, the DOD awards Phase I contracts to a number of the firms; this number is a function of the R&D budget of the particular component and command in the DOD letting the project as well as potentially project-specific characteristics. According to the Navy SBIR Program Overview, Phase I

is a feasibility study to determine the scientific or technical merit of an idea or technology that may provide a solution to the Department of the Navy's need or requirement.¹¹

Phase I often involves preliminary prototyping, benchtop testing, computer simulations, and other forms of low-cost preliminary research. The specific award amount for Phase I differs slightly across DOD components, but the Navy currently awards approximately \$80,000 for the base contract along with the potential option of \$70,000 (which the firm is usually only allowed to exercise if it is selected to participate in Phase II). In practice, there is very little variation across competitors and projects in the Phase I award amount. Approximately six months after the contract is awarded, the firms submit a Phase I Final Report detailing their findings, a Phase II proposal that includes plans to implement or manufacture the product designed in Phase I, and a detailed cost proposal for Phase II research. The DOD evaluates the proposals primarily on technical merit and essentially excludes any consideration of the proposed cost of Phase II research;¹² in the case of the Navy, the PEO itself is in charge of making Phase I and Phase II selections.¹³ The targeted number of Phase II contestants is about 40% of the number of Phase I awards, although the DOD reserves the right to award Phase II contracts to fewer firms.¹⁴ In fact, in 17% of the projects in my dataset, the DOD chooses not to let the contest continue into Phase II.

To assess the commercial viability of the idea generated in Phase I, Phase II awardees are awarded larger contracts to conduct more intensive research to build and test prototypes. Typical contract amounts are on the order of \$500,000 to \$1.5 million. They vary considerably both across projects and across competitors within a specific project: the Navy solicitation guidelines

¹⁰Examples of recent solicitations (in 2015) include one for a "Compact Auxiliary Power System for Amphibious Combat Vehicles" and one for "Navy Air Cushion Vehicles (ACVs) Lift Fan Impeller Optimization." The former is let by the Program Manager for Advanced Amphibious Assault and the latter is for the Ship-to-Shore Connector Acquisition Program.

¹¹See <http://www.navysbir.com/overview.htm>.

¹²Section 8 of the DOD SBIR solicitation guidelines (<http://www.acq.osd.mil/osbp/sbir/solicitations/sbir20162/preface162.pdf>) notes that the primary dimension on which proposals are evaluated is "the soundness, technical merit, and innovation of the proposed approach and its incremental progress toward topic or subtopic solution" and that this criterion is "significantly more important than cost or price."

¹³See <http://www.navysbir.com/natconf14f/presentations/3-09-Navy-Comm-Williams.pdf>.

¹⁴The Phase II desk reference (http://www.acq.osd.mil/osbp/sbir/sb/resources/deskreference/12_phas2.shtml) includes the following quote: "The DoD Components anticipate that at least 40 percent of its Phase I awards will result in Phase II projects. This is merely an advisory estimate and the government reserves the right and discretion not to award to any or to award less than this percentage of Phase II projects."

note that Phase II is structured in a way “that allows for increased funding levels based on the project’s transition potential.” Throughout the phase the firm remains in contact with the DOD and submits interim progress reports, submitting a final progress at the conclusion of the Phase II contract, which usually lasts approximately two years.

Unlike many other federal agencies, the DOD SBIR process includes a formal “Phase III,” which is the final goal of most firms involved in these contests. Phase III is essentially a delivery phase in which the firm either implements or produces the technology developed in Phases I and II for the DOD or for prime contractors through a DOD contract. Phase III does not use funds that are set aside specifically for SBIR but is instead funded by the specific acquisition program in charge of the contest. Very few contests—just 9% of my dataset—actually result in a Phase III contract. Finally, while SBIR requirements do not stipulate that only one firm can be awarded a Phase III contract, this is almost always the case in practice.¹⁵ I interpret this as a sign that the technologies developed by multiple Phase II competitors are for the most part sufficiently substitutable that the DOD has value for at most one: this provides the fundamental source of competition in each contest.

1.2.2 Data Sources

I first collect information about the set of all SBIR contracts awarded by the Navy from the Navy SBIR Program Office via www.navysbirsearch.com. This data includes firm information, including name, location, and firm size at the time of the award; the topic number associated with the contract, so that each contract can be mapped back to a particular solicitation and contest; the systems command (“SYSCOM”) of the Navy in charge of the contract; whether the contract is for Phase I, II, or III; and dates for the contract. It also includes the title of the proposal from the firm, keywords associated with the proposal, and an approximately 200-word abstract of the project as well as a two- to three-sentence description of the potential benefit of the project to the Navy.

The data from www.navysbirsearch.com also contains contract numbers for each award. I use these contract numbers to map the SBIR contracts to the Federal Procurement Data System (via www.usaspending.gov) and extract information for each contract. In particular, the FPDS contains information about all options exercised as well as all modifications for each contract, which allows me to compute the total amount awarded to the firm through the contract.

Finally, I collect the full text of all Navy solicitations (from 2000 onward) from the DOD SBIR Solicitation Website¹⁶ and match them to the contracts acquired from the Navy SBIR Program Office. Each solicitation is a one- to two-page document containing a title for the solicitation, broad technology areas associated with the topic, and the acquisition program in charge of the topic. The solicitation also includes a large amount of free text describing the project, including a one-paragraph objective, a one-page description of the problem and technical requirements, and

¹⁵A number of the exceptions in the dataset can be explained by idiosyncratic reasons.

¹⁶See <http://www.acq.osd.mil/osbp/sbir/solicitations/index.shtml>.

Title	Keywords
modeling	modeling, simulation, analysis, software, prediction
aircraft	aircraft, control, unmanned vehicles, flight, operations
data	data, network, software, architecture, security
power	power, energy, heat, thermal, cooling
acoustics	acoustics, sonar, underwater, submarine, anti-submarine warfare
radio	radio, communications, rf, signal, interference, frequency
materials	composite, corrosion, coating, materials, structures
optics	optics, laser, fiber, infrared, wavelength
ballistics	armor, gun, shock, fire, blast
engines	engine, turbine, aircraft control, engines, propulsion
battery	fuel, battery, water, energy storage, cell

Table 1.1: Representative topics generated by the LDA algorithm in MALLET. MALLET only returns a representative list of words corresponding to each topic; the topic name is arbitrarily determined by me for presentation.

guidelines for the goals for Phases I, II, and III. It also includes keywords and references to academic, military, and general-audience publications related to the topic.

The free-flowing text from these datasets—abstracts of winning proposals and information from the solicitation—allows me to construct detailed project-level covariates to control for the topic of the contest via an unsupervised machine learning algorithm. Extracting information and generating regressors from unstructured text is a promising frontier in industrial organization that researchers have only begun to explore, and this program provides a setting that is especially conducive to such analysis.¹⁷ I use a Latent Dirichlet Allocation algorithm for topic modeling implemented in the software package MALLET by McCallum (2002). This algorithm infers topics in the dataset as collections of words that appear together frequently and then classifies documents in the dataset as mixtures of topics.¹⁸ Using such a topic generation algorithm allows for finer distinctions between projects than simply using the broad categories listed by the Navy, which can often encompass a rather wide range of projects. For reference, Table 1.1 lists some representative topics in the dataset (when generating 20 topics), along with common words associated with each topic. Details regarding the algorithm used to construct these topics are provided in Appendix A.6.

¹⁷See Bajari, Nekipelov, Ryan, and Yang (2015a,b) for an application to demand estimation as well as a discussion of the potential uses of unstructured text analysis in economics. One of the few other papers that makes use of text analysis is Gentzkow and Shapiro (2010), which categorizes the bias of newspapers by identifying phrases that are differentially associated with Democrats and Republicans. Hansen, McMahon, and Prat (2014) use an LDA algorithm, like the one I use in this chapter, to estimate the effect of central bank transparency on outcomes.

¹⁸Briefly, this algorithm takes as input a set of documents, each of which it treats as a sequence of words, as well as a fixed number of latent topics. It places Dirichlet priors on the distribution of topics for each document as well as on the distribution of words for each topic. The data generating process the model specifies is roughly one in which each document is a mixture of topics and each topic is a mixture of words: a document can be generated by recursively selecting a topic from this mixture (multinomially) and then selecting a word from this topic (again multinomially). Since the Dirichlet distribution is the conjugate prior for the multinomial, this model lends itself to a computationally attractive sampling procedure to generate topics as well as assign documents to mixtures of these topics. Further details can be found in Blei, Ng, and Jordan (2003).

I restrict the sample to all contests solicited between 2000 and 2012.¹⁹ Before 2000, the Navy was not especially careful about classifying follow-on delivery projects from Phase II contracts as Phase III. Restricting to projects solicited before 2012 allows for enough time to ensure that I can identify which contests culminate in Phase III contracts. I discuss further details of the data cleaning, sample selection, topic generation, and the process of matching the datasets from the three sources together in Appendix A.6.

1.2.3 Descriptive Statistics

In this section, I first provide summary statistics about the dataset. I then report a set of descriptive correlations between success rates and funding amounts that motivate the structural model that I develop in Section 1.3.

Table 1.2 presents basic summary statistics for the number of competitors, contract amounts, and contest covariates.²⁰ Most contests do not involve large numbers of competitors: there are on average 2.5 competitors in Phase I of the contest and 1.1 in Phase II (including the contests with zero Phase II competitors). The average Phase II contract is about \$800,000, and the average Phase III contract is about \$8.8 million; these distributions have large standard deviations as well. The number of contests is relatively balanced throughout the time period: almost exactly 50% of the contests are let no later than 2006. The Naval Air Systems Command and Naval Sea Systems Command solicit about three-fifths of the contests. Finally, Table 1.2 lists the proportions of the six most common topics (as generated by MALLETT). No single topic dominates the contests, as the means for the topic proportions are not much larger than $1/19 \approx 0.052$, which we would expect if documents were randomly assigned to topics.²¹ However, each solicitation is not assigned to a large number of distinct topics: the median value for each of the topics in the dataset is extremely small, suggesting that the topic generation algorithm does discriminate between topics.

Table 1.3 shows the full distribution of the number of competitors in each Phase. About 75% of the contests in the dataset have 2 or 3 Phase I competitors, and less than 4% have more than 4. The transition from Phase I to Phase II is usually not the constraining factor in whether the contest succeeds: over 80% of contests proceed to Phase II, but about 75% of contests that enter Phase II have only one competitor.²² However, very few contests—fewer than 9% of the ones in my dataset, or about 11% of the ones that enter Phase II—lead to a Phase III contract.

Figure 1-1 shows histograms for Phase II and III contract amounts. Panels (a) and (b) show that there is a good deal of variation in these contract amounts. While there is a salient peak

¹⁹I also include Small Business Technology Transfer (STTR) contracts, which are structured in the same way but are reserved for small businesses that collaborate with a nonprofit research institution.

²⁰Table 1.2 restricts the sample to all contests let between 2000 and 2012, dropping 17 contests with more than 1 Phase III awardee. In later parts of this section, I further restrict the sample to more closely match the one used in structural estimation by only considering contests that have no more than 4 Phase I competitors. The numbers in Table 1.3 indicate that this further restriction does not drop much of the data at all and the sample used in structural estimation is representative of the entire population of contests in this time period.

²¹As described in Appendix A.6, I generate 20 topics and drop one that I deem to be too generic.

²²Note that this number is a result of both the success rate of individuals in Phase I as well as the constraint on how many competitors are allowed to enter Phase II.

	<i>N</i>	Mean	Median	SD
Number of Competitors				
Phase I	2875	2.51	2	1.09
Phase II	2875	1.09	1	0.74
Phase III	2875	0.087	0	0.283
Contract Amount (Millions)				
Phase II	3143	0.803	0.749	0.453
Phase III	252	8.77	2.93	13.23
Fiscal Year \leq 2006	2875	0.505	1	0.500
Systems Command				
NAVAIR	2875	0.327	0	0.469
NAVSEA	2875	0.272	0	0.445
Topics				
Information/Data	2875	0.080	0.00150	0.182
Materials/Composites	2875	0.074	0.00015	0.188
Algorithms/Sensing	2875	0.071	0.00045	0.163
Aircraft	2875	0.064	0.00230	0.147
Manufacturing	2875	0.063	0.00017	0.167
Power/Energy	2875	0.061	0.00018	0.162

Table 1.2: Summary statistics for the dataset of all solicitations posted between 2000 and 2012, dropping ones in which multiple Phase III contracts were awarded.

	0	1	2	3	4	\geq 5
# Phase I Comp	–	12.9%	41.8%	32.8%	8.9%	3.6%
# Phase II Comp	16.9%	61.1%	19.0%	2.3%	0.6%	0.2%
# Phase III Comp	91.3%	8.8%				

Table 1.3: Distribution of the number of competitors in each phase. As in Table 1.2, I restrict to solicitations posted between 2000 and 2012 and only consider ones in which at most one Phase III contract was awarded.




Figure 1-1: Distribution of (a) Phase II award amounts and (b) Phase III award amounts. The histogram in (a) includes a datapoint for each contract and can thus include multiple contracts for a particular contest. Panel (c) shows the percent difference between the highest and lowest Phase II award amounts within contests, restricting to contests with at least two Phase II competitors.

around \$750,000 for Phase II contracts, the standard amount for Phase II SBIR contracts in other agencies that is sometimes used as a baseline by the Navy, most contracts are for other amounts. Phase II contracts can be as small as \$200,000 and as large as \$2 million. Phase III contract amounts also have a long right tail and can exceed \$25 million. This variation is plausibly due to two sources of heterogeneity. First, certain projects will plausibly generate more surplus for both the Navy and the firms involved, and thus these contests likely receive more funding. Second, different firms likely have ideas that the Navy values differently—even *within* a contest. Panel (c) restricts the sample to contests with at least two Phase II competitors and plots a histogram of the percent difference between the contract amounts for the firms with the largest and smallest contracts. Because this comparison controls perfectly for contest-level heterogeneity, I will interpret large differences in contract amounts as suggestive of variation in the value of the projects of each competitor.²³ Indeed, the histogram shows that differences in amounts can be very large: the best-funded competitor often receives more than 50% more funding than the worst-funded competitor, and the difference is not unlikely to even exceed 100%.

How does the number of competitors affect the probability that the contest transitions into the subsequent stage? Adding a competitor increases the number of draws from the pot and, ignoring any endogenous responses to effort, should increase the probability of at least one competitor succeeding. However, there may be a nontrivial equilibrium response in competitive effort: firms may reduce research effort in response to an increase in competition, because they anticipate a lower probability of capturing the return to effort, or they may increase their effort on the margin in response to the competitive pressure. The former outcome is more likely in a setting with less differentiation across firms; the transition from Phase I to Phase II can be approximated by such a model. The latter outcome can happen if there is some heterogeneity across firms, as may be the case as firms in Phase II compete to enter Phase III. As such, in both cases, the *net* effect is in

²³This interpretation is consistent with the DOD’s claim that it gives more funding to projects that have increased transition potential. Furthermore, it is consistent with the evidence I will present that these projects are indeed more likely to lead to Phase III contracts. On the other hand, an alternate interpretation that attributes this variation solely to heterogeneity in research cost would not immediately be able to explain this correlation.

	Contest Success		Individual Success		Log(Amount)	
	Phase I	Phase II	Phase I	Phase II	Phase II	Phase III
	(1)	(2)	(3)	(4)	(5)	(6)
# Phase I Comp	0.066*** (0.009)	-0.018** (0.008)	-0.128*** (0.008)	-0.023*** (0.008)	0.016 (0.012)	0.234* (0.110)
# Phase II Comp		0.076*** (0.016)		0.028*** (0.010)	-0.002 (0.016)	-0.429** (0.176)
Log([Avg] Phase II Amt)		0.157*** (0.018)		0.250*** (0.031)		0.330** (0.195)
Fiscal Year FE	✓	✓	✓	✓	✓	✓
SYSCOM FE	✓	✓	✓	✓	✓	✓
Topics	✓	✓	✓	✓	✓	✓
R^2	0.083	0.128			0.133	0.422
N	2773	2292	2773	2292	2292	151

Table 1.4: Regressions of a dummy of whether the contest enters Phase II (columns (1) and (3)) or Phase III (columns (2) and (4)) on the number of competitors in Phases I and II, controlling for year fixed effects, SYSCOM fixed effects, and topic covariates. I restrict the sample to contests with no more than 4 Phase I competitors. Columns (3) and (4) restrict to the set of contests that enter Phase II. Columns (5) and (6) regress the log of the contract amount in Phases II and III on observables, controlling for the same covariates. Note that Log([Avg] Phase II Amt) refers to the log of the within-contest average of Phase II contract amounts in columns (2) and (6) and the log of the individual firm’s Phase II amount in column (4).

principle ambiguous.

Table 1.4 reports OLS regressions of contest-level “success” rates from Phase I to Phase II and from Phase II to Phase III. I run linear probability models of the contest transitioning to a particular phase, controlling for contest-level heterogeneity using year fixed effects, SYSCOM fixed effects, and the topics information generated from the text descriptions via MALLETT. For all regressions in this section, I restrict the sample to contests with no more than 4 competitors in Phase I to have a sample that is as close as possible to the one used in the structural estimation in Section 1.4.3. Column (1) indicates that increasing the number of competitors in Phase I by 1 is associated with an average increase in the probability of at least one firm advancing to Phase II by 6.6 percentage points—compared to a mean of 83%. Column (2) reports similar regressions for the transition from Phase II to Phase III, controlling for both competition in Phase I and Phase II as well as average funding per firm in Phase II.²⁴ Adding a Phase II competitor increases the probability of transitioning to Phase III by 7.6 percentage points, which is an especially large number compared to the mean success rate of 10.5%. Somewhat counterintuitively, contests with one additional competitor in Phase I tend to have a *lower* rate of transitioning from Phase II to Phase III by 1.8 percentage points on average—a number that is small but significant at the 5%

²⁴I do not control for Phase I funding in any of these regressions because, unlike for Phase II funding, there is almost no variation in Phase I funding.

level. This correlation is admittedly at odds with the idea that more Phase I competitors are associated with stronger competitors entering Phase II, although other results (discussed below) do suggest that this effect is reasonable. If anything, this correlation suggests that we should be aware of the endogeneity concern that contests with different numbers of Phase I competitors could be systematically different from each other.²⁵ I will allow for this possibility when estimating the structural model.

Columns (3) and (4) of Table 1.4 investigate the probability that an *individual* competitor generates successful research. Because individual successes are not observed,²⁶ I use explicit models of censoring to estimate the probability $p(X_{ij})$ that a contestant i succeeds in contest j as a function of contest-level covariates and individual-level funding. For the transition from Phase I to Phase II, I estimate a censored binomial model in which for each contest j , the unobserved number of successes N_{Sj} is such that $N_{Sj} \sim \text{Binomial}(N_1, p(X_j))$, but the observed quantity is

$$N_{2j} = \begin{cases} N_{Sj} & \text{if } N_{Sj} \leq \bar{N}_{2j} \\ \bar{N}_{2j} & \text{if } N_{Sj} > \bar{N}_{2j} \end{cases}$$

I estimate this model via MLE, letting $p(X_j)$ be a linear function of N_1 (or having fixed effects for all values of N_1) and controlling for the same contest-level covariates, and I report $p(\cdot)$ in columns (1) and (2). I do not directly observe the limit on Phase II competition in the data, so I leverage the 40% rule that I also use in Step 5 of the structural estimation. Since the DOD aims to let at most 40% of the competitors in Phase I into Phase II, I assume that $\bar{N}_2 = 1$ if N_1 is 1 or 2, and $\bar{N}_2 = 2$ if N_2 is 3 or 4. If N_2 exceeds the candidate value of \bar{N}_2 , I say that $\bar{N}_2 = N_1$. Column (3) shows that adding one competitor to Phase I is associated with a decrease in the probability of an individual competitor generating a successful innovation by 12.8 percentage points.

I model the transition from Phase II to Phase III as the following: a contestant i generates a successful innovation in contest j with probability $p(X_j; t_{ij})$, where t_{ij} is the Phase II research funding; if no one succeeds, then the project does not enter Phase III, but if multiple contestants succeed, one contestant is awarded the Phase III contract uniformly at random.²⁷ Column (4) indicates that contestants in contests with one additional Phase II competitor have a *higher* probability of success, by about 2.8 percentage points. Once again, the individual success rate is lower for contests with more Phase I competitors; while this may be due to stronger competition dissuading research effort, it may also be an indication of differences across contests not controlled by these models.

²⁵In principle, this correlation could be explained by stronger competition leading to lower incentives to spend money on research, which in turn leads to a lower success rate. This explanation is, however, at odds with the results in the final two columns of Table 1.4, which show that contests with more Phase I competitors have slightly more funding in Phase II and lead to significantly larger Phase III funding amounts. Appendix A.2.1 models the dependence on N_1 more flexibly, and the source of the negative coefficient on N_1 is primarily contests with $N_1 = 4$.

²⁶That is, while I do observe how many firms entered Phase II, it could be that more firms generated innovations that could have merited Phase II grants.

²⁷Once again, a “success” is a project that would be worthy of a Phase III contract; however, at most one firm is offered a Phase III contract. In Section 1.3, I will develop an explicit model for how the DOD decides between multiple “successful” firms, but I use the uniform-at-random assumption for the descriptive analysis.

What affects Phase II contract amounts? Column (5) of Table 1.4 shows regressions of the average Phase II funding per firm within-contest on the number of competitors. Contests with one more Phase I competitor have on average 1.6% more funding, an amount that is both small and imprecisely estimated. Overall, adding more Phase II competitors has no impact on average funding, although Appendix A.2.1 notes a large drop when moving from contests with 3 to contests with 4 Phase II competitors. The institutional details provided in Section 1.2.1 suggest that firms with more promising research projects are given more funding. Moreover, increased funding probably directly leads to a higher rate of success. Accordingly, we would expect that funding correlates positively with success in Phase III. Indeed, Columns (2) and (4) of Table 1.4 show that on average, increasing the average funding by 10% is associated with an increase in the contest-level success rate of 1.6 percentage points, and an increase in the individual-level funding by the same proportion increases individual success by 2.5 percentage points. Moreover, Appendix A.2.1 shows evidence that even *within contest*, firms with larger Phase II contract amounts are more likely to enter Phase III.

Finally, column (6) of Table 1.4 regresses the Phase III contract amount against Phase II award amounts and measures of competition. Because the Phase III contract is for delivery, one would expect that it increases not only with delivery costs but also with the value the product brings to the DOD: as long as the firm has some bargaining power in the procurement process, it should be able to extract some surplus from the DOD. Moreover, we would expect a competitive effect to lower the Phase III award amount: if there are multiple Phase II competitors, the DOD can capture a larger portion of the surplus by threatening to go to a second-best competitor who may have also produced a successful innovation. The predictions related to Phase III contract amounts are therefore threefold: (1) a larger number of Phase I competitors would possibly indicate that firms with more valuable projects survive into later rounds and thus would lead to larger Phase III contracts, (2) having more Phase II competitors would give the DOD more chances for a lower draw of the cost of delivery—and also more bargaining power by leveraging competition—and thus lead to lower Phase III contracts, and (3) more Phase II funding is associated with both higher-value projects and better draws of cost (via more research) and thus lead to lower Phase III award amounts. The coefficients of OLS regressions agree with these predictions, although having a small fraction of contests with successful Phase III contracts leads to power issues.²⁸ Adding one Phase I competitor is associated with an increase in the Phase III contract amount by about 26%, a large and marginally significant amount. Adding a Phase II competitor is associated with a reduction in the Phase III contract amount by about 35%, which is also large but imprecisely estimated. Finally, a 10% increase in average Phase II funding is associated with a 3.3% increase in the Phase III contract amount.

I use these descriptive correlations as motivation for developing and estimating a structural model with features that are consistent with these correlations. In the model I present in the next

²⁸Furthermore, I restrict to contests where the Phase III contract amount is at least \$1 million to avoid data points where the Phase III contract is unnaturally small. Results are qualitatively robust to using the entire dataset, as discussed in Appendix A.2.1.

section, firms will learn the values of their projects from the end of Phase I, and the strongest firms move on to Phase II. Firms with more valuable projects are awarded larger Phase II research contracts, which makes them more likely to develop technologies with lower delivery costs. Finally, the DOD engages in a form of Nash bargaining that allows it to leverage competition between the successful Phase II competitors in the procurement market. Because a drawback of the descriptive analysis is that it makes it difficult to separately disentangle values and costs, I will leverage the structural model to back out these parameters from the observables. The structural model will also give me an explicit way to control for potential differences across contests with different numbers of Phase I competitors, a concern that became clear through the descriptive analysis.

1.3 Model

In this section, I present a model of a multistage R&D contest that captures the salient features of the DOD SBIR program. In Section 1.3.1, I present the primitives and the timing of the model, detailing how research efforts translate to values and costs, and how values and costs determine how the contracts are awarded in the various phases. Section 1.3.2 then discusses two assumptions for how research efforts are determined. I present both a weak assumption, which may be widely applicable to settings beyond the one considered in this chapter, and a stronger one that is consistent with institutional details of the DOD SBIR program. Presenting these two assumptions separately lets me highlight the identifying power of each of them in Section 1.4.1 and develop a natural estimation procedure in Section 1.4.3. In particular, the weak assumption is sufficient to identify a large subset of the primitives discussed in Section 1.3.1.

1.3.1 Model Timing and Primitives

Each contest in the SBIR program consists of three phases. The primitives of each contest are the number of contestants in Phase I (N_1), the maximum number that will be allowed to enter Phase II (\bar{N}_2), the distributions from which firms draw values (V), the cost functions ($\psi(\cdot)$ and $H(\cdot; \cdot)$), and the firm’s bargaining parameter in the acquisition phase (η). I discuss each phase—and which primitives are relevant for it—in sequence.

Phase I. Phase I is a prototyping phase in which firms exert effort to determine both the feasibility and the potential value of the innovation. Note that while I will refer to this potential value as a “value” throughout the chapter, it is important to conceptualize this quantity as the value to the DOD. The DOD invites N_1 firms to participate in Phase I, and firms are ex-ante identical. If firm i spends the monetary amount $\psi(p_i)$ (with $\psi(\cdot) > 0$, $\psi'(\cdot) > 0$, and $\psi''(\cdot) > 0$) on its Phase I project, then it generates a successful innovation with probability p_i . The events that two different firms succeed at developing the same innovation are mutually independent.

The N_S firms that succeed each independently draw a value $v_i \sim V$ with cdf F . At most \bar{N}_2 of the N_S firms that succeed are allowed to proceed to Phase II. That is, if $N_S \leq \bar{N}_2$, then all firms

that succeed enter Phase II. If $N_S > \bar{N}_2$, then the \bar{N}_2 firms with the highest draws of v are the ones that proceed to Phase II. Note that a contest can fail in Phase I if none of the participants succeed.

Phase II. The goal of Phase II is to develop a commercially viable production plan; that is, firms conduct research to reduce the delivery cost (e.g., manufacturing cost for physical products or implementation cost for software) of their innovation. In Phase II, each firm spends some amount t , which could depend on all the other parameters of the contest. (I suppress this dependence for the sake of brevity.) Exerting effort t results in a draw of the delivery cost c from a distribution $C(t)$ with cdf $H(\cdot; t)$ and density $h(\cdot; t)$. This distribution is first-order stochastically decreasing in the effort t so that more effort corresponds to drawing lower delivery costs. Note that a project fails in Phase II if all participants draw costs that exceed their values. How t is determined will be discussed in Section 1.3.2.

Phase III. This final phase is a delivery phase, in which the procurer contracts with at most one of the firms to deliver the product. The procurer sees the realization (v_i, c_i) for all firms in Phase II and selects a winner based on the following procedure. The procurer approaches the firm with the highest *surplus* (value of $v - c$), as long as it is positive, and Nash bargains as if its outside option is to go to the firm with the second-highest surplus and extract all its surplus. Thus, a firm wins if it has the highest value of $v - c$. The winner gets a profit of η times the excess surplus he generates, which amounts to a transfer of $c + \eta(v - c - s)$, where s is the second-highest value of $v - c$ (and is 0 if all other competitors have $c > v$).²⁹

1.3.2 How Are Research Efforts Determined?

In this section, I present two possible assumptions for how Phase I and Phase II efforts are determined in a particular empirical setting. The first (Assumption M) is especially general and simply states that the map from values to Phase II research efforts is monotone (conditional on the other primitives in the model). The second assumption (Assumption O) is that the firm is the one choosing the optimal amount of research, in a manner consistent with the model outlined in Section 1.3.1. I then show that this second assumption implies the first in many cases and discuss how this stronger assumption is consistent with the institutions of the SBIR program. By separating these two assumptions, I can be clear in Section 1.4.1 about which aspects of the structure imposed in Assumption O are used to identify which parameters. Furthermore, because Assumption M

²⁹Note that in this empirical setting, it is overwhelmingly the case that only one competitor is successful in Phase II. About 75% of contests that enter Phase II have only one firm; even when there is more than one firm entering Phase II, the low success rate suggests it is highly unlikely that multiple firms develop successful innovations. Thus, the precise extension of Nash bargaining to multiple parties is not especially relevant empirically. One could consider alternate models, such as Shaked and Sutton (1984) and Bolton and Whinston (1993), or a bargaining procedure in which the DOD negotiates with the highest-*value* party instead of the highest-*surplus* party first. Many of these models still respect monotonicity, but they do change incentives in the model described in Section 1.3.2 (although, once again, by a small amount in this setting).

is more general, stating it separately can help provide guidance on which other settings—beyond R&D contests—are appropriate for the methodology developed in this chapter.

Throughout I assume that research efforts depend only on one’s own value v and not on opponents’ values, and I thus discuss an effort function $\hat{t}(v)$.³⁰ I begin with the more general assumption.

Assumption M. *The research effort $\hat{t}(v)$ is an increasing function of the firm’s value v . This map may depend on all of the primitives of the contest, as well as on the realization of N_2 .*

First, note that Assumption M(monotonicity) places absolutely no restrictions on how Phase I efforts \hat{p} are set. Second, the restriction that is placed on how Phase II efforts are determined is that higher-value firms exert more effort and that effort only depends on one’s own value. This assumption is relatively weak and may be broadly applicable outside the specific institutional setting considered in this chapter. For instance, in certain contests, small firms may be given a research award that is an institutionally specified function of a quality score (the “value”), and they may exhaust the award on research for the project.³¹ Outside the specific context of R&D contests, one could imagine that higher-quality startups, which are capital-constrained, also attract more external funding and thus spend more money developing their research projects. Finally, Assumption M may be applicable when the firms themselves choose how much to invest in the R&D project. I discuss this case further below.

For the rest of this chapter, I impose an additional assumption that seems appropriate in the particular empirical setting of the DOD SBIR program: the contract amounts for Phases I and II coincide with the research efforts that the firm would choose itself, meaning that the DOD contract amounts are the *firm-optimal* ones. Stating the firm’s problem to define these optimal amounts involves specifying information sets, beliefs, and objectives at each phase.

Phase I. In Phase I, firms are aware of the number of Phase I competitors N_1 as well as the limit \bar{N}_2 on the number of Phase II competitors. Firms also know the primitives of the contest, such as F , η , $\psi(\cdot)$, and $H(\cdot; \cdot)$. At the time of exerting effort, each firm has no further information.

Phase II. In Phase II, each firm is given a lump sum award by the DOD, denoted $t^{DOD}(v)$.³² It then decides on effort to reduce its delivery costs. In doing so, it knows its own value v_i and is informed of the number N_2 firms that entered Phase II. However, they are informed of neither the

³⁰One institutional justification is that firms know their own values at the start of Phase II but do not know their opponents’.

³¹This could be the case when monitoring is especially strong and the monitoring agency can check whether each dollar is spent on the project itself. Alternatively, one can imagine that this is likely to be the case when firms are especially small, i.e., smaller than the typical firm that participates in the DOD SBIR program. Such firms may have no other ongoing R&D projects, and as long as the award cannot literally be pocketed and used as profit, they would exhaust the award on research.

³²This award captures the Phase II contract. Assume that this contract can depend on all primitives of the contest as well as the realization of N_2 . Because this contract is purely a function of primitives and value v , and because the DOD is informed of the firms’ values, this contract is simply a lump-sum transfer and does not affect incentives to exert research effort at this stage. Note that this transfer does affect research incentives in Phase I. In the empirical setting, I make the assumption that $t^{DOD}(v) = \hat{t}(v)$ (Assumption O), which corresponds to the assumption that the DOD fully refunds the *firm-optimal* level of research costs.

number of successes N_S nor about the values of their opponents' projects. They form beliefs (with cdf $F(\cdot; v_i, N_2, p)$) of their opponents' values, where p is their belief of the Phase I effort of each of their competitors,³³ and, based on these beliefs as well as their own values, they exert effort t_i to get the cost draws $c_i \sim H(\cdot; t_i)$.

To compute beliefs, note that a firm's own value can give information about the values of his opponents *only if there is selection in entry into Phase II*. That is, if $N_2 < \bar{N}_2$ or $N_2 = N_1$, then it is common knowledge that every firm that succeeded was granted entry into Phase II. Thus, all firms know that the values of their opponents are drawn from V . The case $1 < N_2 = \bar{N}_2 < N_1$ is complicated by the fact that there is both selection into Phase II as well as competition between firms. Furthermore, beliefs of the values of two different opponents are not independent. If one's own value is v , the probability that the other $\bar{N}_2 - 1$ players have values \mathbf{v}_{-i} is

$$f_v(\mathbf{v}_{-i}; v, \bar{N}_2, p) \propto \sum_{N_S=\bar{N}_2}^{N_1} \left\{ \overbrace{\frac{(N_S - 1)!}{(N_S - \bar{N}_2 - 1)!} \left(\prod_{v_{-i} \in \mathbf{v}_{-i}} (p \cdot f(v_{-i})) \right)}^{\text{succeeded, with given values}} \times \underbrace{\left(\frac{N_1 - \bar{N}_2}{N_S - \bar{N}_2} [p \cdot F(\min\{\mathbf{v}_{-i}, v\})]^{N_S - \bar{N}_2} \right)}_{\text{succeeded but drew lower values}} \times \underbrace{(1 - p)^{N_1 - N_S}}_{\text{did not succeed}} \right\}. \quad (1.1)$$

Phase III. Phase III is mechanical: values and costs are drawn in previous rounds and shared with the DOD, and the surplus is determined as a mechanical result of the Nash bargaining procedure described in Section 1.3.1.

Equilibrium. A type-symmetric equilibrium of this model consists of an effort function $t_{N_2}^*(v)$ for Phase II competitors (as a function of the realized number N_2 of competitors) as well as a Phase I probability of success p^* .

Focus on Phase II with N_2 entrants. Consider a firm with value v and beliefs with cdf $F(\cdot; v, N_2, p^*)$ about its opponents' values; note that these beliefs could depend on both the value of the competitor as well as the first-stage entry probability, as discussed above. Suppose its opponents follow an effort function $t_{N_2}^*(v)$. The firm's optimization problem is then given by

$$\arg \max_t \left\{ \eta \int_c^v \int_c^{v-c} (v - c - \max\{s, 0\}) dG(s; v, t_{N_2}^*(\cdot), p^*) dH(c; t) - t + t^{DOD}(v) \right\}, \quad (1.2)$$

where $G(s; v, t_{N_2}^*(\cdot), p^*)$ is the cdf of a type v competitor's beliefs about the highest surplus of its competitors. Note that the cdf of the surplus that a type v' firm generates is given by

$$S(s; v', t_{N_2}^*(\cdot)) = 1 - H(v' - s; t_{N_2}^*(v')) \quad (1.3)$$

³³In principle, firms could believe that each of their opponents exerted a *different* amount of effort. However, I will restrict to (type-)symmetric equilibria, and as such, I will restrict the notation at this point for brevity.

and the cdf of the maximum surplus of a type- v firm's opponents can be computed by combining (1.3) and (1.1) as

$$G(s; v, t_{N_2}^*(\cdot), p^*) \equiv \iint_{\mathbf{v}_{-i}} \left(\prod_{v_{-i} \in \mathbf{v}_{-i}} S(s; v_{-i}, t_{N_2}^*(\cdot)) \right) f_v(\mathbf{v}_{-i}) d\mathbf{v}_{-i}.$$

Let $\pi(v, N_2, p^*)$ denote the maximized value of (1.2). In Phase I, each firm chooses p to maximize the expected profits from Phase II, less the cost of Phase I effort. Since the expected profits from Phase II can be expressed as p times the profits conditional on success, we can write the firm's problem in Phase I as

$$p^* = \arg \max_{p \in [0,1]} \left\{ p \cdot \left[\sum_{N_S=1}^{\bar{N}_2} \binom{N_1-1}{N_S-1} (p^*)^{N_S} (1-p^*)^{N_1-N_S} \int_0^{\bar{v}} \lambda(v, N_S, \bar{N}_2) \pi(v, N_2, p^*) dF(v) \right] - \psi(p) \right\}, \quad (1.4)$$

where

$$\lambda(v, N_S, \bar{N}_2) \equiv \begin{cases} 1 & \text{if } N_S \leq \bar{N}_2 - 1 \\ \sum_{N_b=0}^{\bar{N}_2-1} \binom{N_S-1}{N_b} F(v)^{N_b} (1-F(v))^{N_S-1-N_b} & \text{otherwise} \end{cases}$$

is the probability that a successful firm with value v is allowed to enter Phase II if $N_S - 1$ other firms succeed. Collecting the equations in this section, we have that a type-symmetric Bayesian Nash equilibrium of the R&D contest is a p^* and a set of effort functions $\{t_{N_2}^*(\cdot)\}_{N_2 \leq \bar{N}_2}$ that simultaneously satisfy (1.2) and (1.4).

Assumption O. *The Phase I effort \hat{p} and Phase II effort schedule $\hat{t}(v)$ coincide with the type-symmetric Bayesian Nash equilibrium of the model of R&D contests, given by p^* and $\{t_{N_2}^*(\cdot)\}_{N_2 \leq \bar{N}_2}$, which satisfy (1.2) and (1.4).*

Assumption O(ptimality) states that the amounts spent on research—i.e., the amounts that determine the probability of success in Phase I and the distribution of cost draws in Phase II—are chosen by the firm. When taking the model to the data under Assumption O, I will assume that the Phase II research award coincides with this firm-optimal amount as well, so the DOD reimburses the cost of effort. In the case of Phase II, for instance, this amounts to saying that $t^{DOD}(v) = t^*(v)$. While there is admittedly a tension in assuming that the DOD transfer is the firm-optimal amount, this assumption is justifiable in this empirical setting. In practice, the firm submits a detailed cost proposal to the DOD for Phase II research, and the DOD can approve the funding amount or propose modifications to this amount.³⁴ Because the DOD has full information about the value of the particular firm's project, it can compare this proposed amount to the firm-optimal amount. First note that the DOD would be hesitant to offer the firm more funding than the optimal amount: these firms often have multiple ongoing projects and contracts, and given that the DOD can only imperfectly monitor how the firms spend the money, the firms can redirect

³⁴Since the DOD SBIR solicitation guidelines explicitly state that requested Phase II funding is not a factor in deciding which projects get funding, the firms need not be strategic about this amount.

some excess resources to other projects. One can conceptualize this process as the DOD giving an unconditional lump-sum transfer to the firm via the Phase II research contract and the firm then being able to choose the optimal amount to spend on research. Secondly, the DOD actively tries to encourage firms to participate in the defense industrial base through this program, and as such, it would like to limit ex-post losses. Were the DOD to award less than the firm-optimal amount, the firm would try to use money from other sources and suffer losses if the project does not enter Phase III. Even in a setting in which firms may have positive expected profits, Phase III is sufficiently rare that firms may have to enter many contests before realizing a payoff.³⁵

Finally, I note that Assumption O, in many cases, implies Assumption M. The following proposition formalizes this idea.

Proposition 1.1 (Monotonicity of Effort). *If each firm's beliefs about its opponents' values are independent of its own value, then $t_{N_2}^*(\cdot)$ is weakly increasing in v , and strictly so if effort is larger than the minimum possible value of effort.*

Proof. I will show that the maximand of (1.2) is strictly supermodular, and the proof will follow from a standard monotone comparative statics argument. We can write the first term of the maximand as

$$\begin{aligned} & \eta \int_{\underline{c}}^v \left[\int_0^{v-c} (v-c-s) dG(s) + (v-c)G(0) \right] dH(c; t) \\ &= \eta \int_{\underline{c}}^v \left[-(v-c)G(0) + \int_0^{v-c} G(s) ds + (v-c)G(0) \right] dH(c; t) \\ &= \left(\int_0^{v-c} G(s) ds \right) H(c, t) \Big|_{\underline{c}}^v + \int_{\underline{c}}^v G(v-c)H(c, t) dc = \int_{\underline{c}}^v G(v-c)H(c, t) dc. \end{aligned}$$

The cross partial with respect to v and t is

$$G(0) \frac{\partial H(v, t)}{\partial t} + \int_{\underline{c}}^v g(v-c) \frac{\partial H(c, t)}{\partial t} dc,$$

and each term is strictly positive. □

The intuition for Proposition 1.1 is that higher-value firms have both a higher probability of winning as well as a higher surplus conditional on winning. Moreover, the marginal winner is the one whose incremental contribution to surplus is exactly zero, and this firm earns zero profits. These two observations are key for the monotonicity result. However, if we do allow firms' beliefs about opponents to vary with values, as in the case with selection, then there is an additional effect that firms with weaker values tend to believe their opponents are weaker as well. This could

³⁵That firms would substitute internal funds for SBIR funding is consistent with the results of Wallsten (2000). Furthermore, other auxiliary evidence provided in this chapter suggests that the DOD is reasonably generous to firms throughout this process. For instance, the estimated bargaining parameter in Section 1.5 indicates that firms capture three-fourths of the surplus.

encourage them to exert more effort than firms with higher values, and the proof of Proposition 1.1 does not apply.³⁶

1.4 Identification and Estimation

In this section, I discuss identification of the model under Assumptions M and O. I then incorporate both observed and unobserved heterogeneity into the model from Section 1.3 and present an estimation procedure that is based on this identification argument.

1.4.1 Identification

Suppose that in the model in Section 1.3, we observe the numbers of players N_1 and \bar{N}_2 along with the realized number of Phase II players N_2 , the Phase I research effort ($\psi(p^*)$), the distribution of Phase II research efforts $t_{N_2}^*$ for $N_2 \leq \bar{N}_2$ (as long as $p^* \in (0, 1)$), and the Phase III contract amount (if the project enters Phase III).³⁷ The primitives we wish to identify are the cost function $\psi(\cdot)$, the value distribution V , the cost distribution $C(t)$ as a function of Phase II research efforts, and the bargaining parameter η . We will identify the Phase II and III primitives (i.e., everything except ψ) using (i) a selection equation that stipulates implementation in Phase III occurs if and only if the winner's value exceeds his cost, (ii) monotonicity of the Phase II effort in the value to recover values from effort (Assumption M), and (iii) a first-order condition that ensures that Phase II research effort is set optimally, with knowledge of η (Assumption O). The identification argument I provide is constructive, and I present it in two parts. The first part rests on the weak assumptions (i) and (ii) that are likely to have analogues in many different models, and I show that most of the primitives are identified given these assumptions. The second part applies to the specific model with Assumption O.

Identification Under Assumption M

Consider the model timing model described in Section 1.3.1 and suppose that Assumption M is satisfied. Restrict attention to contests where the realized number of Phase II competitors is $N_2 = 1$. Such auctions must exist in the data generating process dictated by the model as long as $\hat{p} \in (0, 1)$ (or $\bar{N}_2 = 1$ if $\hat{p} = 1$). Consider the distribution of the Phase III transfers conditional on a particular value t_2 of Phase II research. If Assumption M holds, this amounts to conditioning on some (yet unknown) value $v(t_2) = \hat{t}^{-1}(v)$, given by the inverse of the effort function. The transfer is $\eta v(t_2) + (1 - \eta)c$, where $c \sim C(t_2)$ if $c \leq v(t_2)$ and unobserved otherwise. Thus, the largest observed value of the Phase III transfer for a particular value of t_2 occurs when $c = v(t_2)$, and thus the maximum observed value of the transfer identifies $v(t_2)$. Varying t_2 identifies the entire function

³⁶I have not been able to find a counterexample where the computed equilibrium is nonmonotone, however, even if there is selection.

³⁷Note that research efforts are dollar amounts, measured as the Phase I and II contract amounts. Note as well that I discuss identification with and without knowledge of $\psi(p^*)$, because this amount is set institutionally and exhibits little variation, and thus it may be unrepresentative of the true expenditures on research in this empirical setting.

$v(\cdot)$ and thus the distribution of the values of competitors who enter Phase II nonparametrically. If there is no selection into Phase II (i.e., if $\bar{N}_2 > 1$ or $N_1 = 1$), then this distribution is simply the distribution of V . Otherwise, we can simply correct for selection to recover the distribution of V , as discussed in Appendix A.5.1.

Now suppose that η is known to the researcher. The next observation is that (part of) the distribution of costs $H(\cdot; t)$ is identified *as a function of this known η* . This is a simple function of the distribution of the Phase III transfer. With knowledge of the value $v(t_2)$, we can invert the observed distribution of $\eta v + (1 - \eta)c$ to determine the cost cdf as a function of η (but only for $c \leq v(t_2)$). For brevity, denote this cdf by $H(\cdot; t_2, \eta)$ and its associated pdf by $h(\cdot; t_2, \eta)$ to make this dependence on η explicit.

Note that the probability \hat{p} of success in Phase I is observed in the data. Truncation due to \bar{N}_2 is not an issue for identification: even when $\bar{N}_2 = 1$, the probability that Phase II does not occur is $(1 - \hat{p})^{N_1}$.³⁸ Since the Phase I research effort is also observed, we identify the single point $\psi(\hat{p})$. Variation that affects \hat{p} but not ψ can identify the entire cost function. We will be more explicit about the source of this variation with an explicit model for research efforts, such as Assumption O; one may expect that contests that are known to have different value distributions without having different Phase I cost functions would have different values of \hat{p} and thus different observed values of $\psi(\hat{p})$.

The following proposition summarizes this identification argument.

Proposition 1.2. *Suppose we have data on distributions of Phase III transfers, Phase I and II research efforts, and the realized number of Phase II competitors for a set of contests with a single (N_1, \bar{N}_2) . If Assumption M holds and η is known, then*

- (i) V is nonparametrically identified (and does not depend on η);
- (ii) $H(c; t)$ is nonparametrically identified on $[0, v(t)]$;
- (iii) and a single point on $\psi(\cdot)$ is identified, and variation in \hat{p} identifies $\psi(\cdot)$ entirely.

Furthermore, Assumption M gives information about a lower bound on the firm's bargaining parameter η from a combination of the failure rate as a function of research effort and the stochastic dominance condition on the cost distributions as a function of research effort. Since the estimation procedure in this chapter utilizes an optimality condition to recover information about the bargaining parameter (see Proposition 1.3, below) instead of exploiting this partial identification argument, I relegate the discussion of the identification of this lower bound from Assumption M to Appendix A.3.1.

³⁸Throughout this chapter, I maintain the assumption that successes in Phase I are uncorrelated. This assumption is mainly due to a data restriction, as most contests in the dataset have $\bar{N}_2 = 1$. Note, however, that with enough data on contests with $\bar{N}_2 > 1$, this assumption is testable: departures from the binomial distribution on N_2 will point towards correlation in successes. In particular, if certain projects are physically infeasible for all firms, we would expect a larger mass point at $N_2 = 0$ than would be expected from the remainder of the distribution.

Identification Under Assumption O

Suppose further that the research efforts are set optimally for the firm, as per Assumption O. Then, the bargaining parameter is identified as well. To see this, note that we know that the firm sets t_2 in response to its first-order condition, so that

$$\eta \int_{\underline{c}}^{v(t_2)} (v(t_2) - c) \frac{dh}{dt}(c; \eta, t_2) dc = 1. \quad (1.5)$$

The intuition is that (1.5) is an equation in a single variable (because $h(\cdot; \eta, t_2)$ is identified, albeit as a function of η) and thus identifies η . The full argument is slightly more involved and is based on rearranging (1.5) in terms of observables and quantities that have already been identified. I relegate it to Appendix A.5.1.

Optimality of the first-stage effort also gives us more information about the cost function $\psi(\cdot)$ than simply under Assumption M. In fact, $\psi'(\cdot)$ can be identified within a single parameter family of functions without observing Phase I expenditures in the data. From $H(\cdot; \cdot)$, V , and η , we can compute $\pi(v, N_2, p)$ for all values v , realizations of N_2 , and p . These quantities then allow us to compute the expected profit conditional on success for any p ; denote this $\pi(p)$. Since the distribution of N_2 is a truncated binomial with parameters N_1 and success probability p^* (truncated at \bar{N}_2), p^* is directly identified from the data. From the firm's first-order condition associated with (1.4) in Phase I, we have that $\psi'(p^*) = \pi(p^*)$. This equation lets us identify the marginal cost of Phase I research at one point. Furthermore, $\psi(p^*)$ is the equilibrium expenditure on Phase I research, and this is seen directly in the data. Thus, $\psi(\cdot)$ can be identified parametrically (within a one-parameter family of functions for $\psi'(\cdot)$), or we can exploit variation in p^* orthogonal to shifts in $\psi(\cdot)$. Note that without the assumption of optimality (i.e., in the baseline model), we could not recover information about the marginal cost and would have to rely exclusively on variation in \hat{p} to recover the cost function.

The following proposition extends Proposition 1.2 and summarizes the arguments in this section.

Proposition 1.3. *Suppose we have data on distributions on Phase III transfers, Phase I and II research efforts, and the realized number of Phase II competitors for a set of contests with a single (N_1, \bar{N}_2) . If Assumptions M and O hold,*

- (i) η is identified;
- (ii) V is nonparametrically identified;
- (iii) $H(c; t)$ is nonparametrically identified on $[0, v(t)]$; and
- (iv) $\psi(\cdot)$ is identified within a single-parameter family of functions for $\psi'(\cdot)$, and variation that continuously shifts the equilibrium probability of success in Phase I without shifting Phase I costs can identify $\psi(\cdot)$ nonparametrically.

Note that identification of $\psi'(\cdot)$ within a single-parameter family of functions does not require data on Phase I research efforts. However, identification of $\psi(\cdot)$ does require either such data or an

assumption akin to $\psi(0) = 0$. I will leverage such an assumption in the empirical model described in Section 1.4.2.

Discussion of the Identification Result

The identification argument for values and costs is at its heart based on a selection rule. This selection happens on a two-dimensional set of Phase II research efforts and Phase III transfers instead of being simply based on Phase III transfers, and the point at which selection occurs is informative of values.

This empirical setting also allows for a novel source of identification for the bargaining parameter that could be applicable to other settings with R&D. I identify the bargaining parameter off an *ex-ante investment*: the firm sets marginal costs equal to marginal returns, and we have information about both—modulo the bargaining parameter—from the joint distribution of contract amounts.³⁹ This identification argument is slightly different from ones used in other empirical papers involving Nash bargaining, and it is worth comparing this argument to those in related papers. Utilizing a different source of identification in the setting of business-to-business transfers, Grennan (2013) identifies the bargaining parameter roughly by comparing distributions of transfers that are generated by different value distributions but similar cost distributions: if the transfer distributions change dramatically, then the effect of the value on the transfer—governed by η in this model—would be high. In my setting, there is in principle an analogous source of identification: different realizations of N_2 shift the value associated with each Phase II effort amount (by shifting the effort function) without shifting the cost associated with each effort amount. However, note that such variation is discrete, and it can be unavailable when $\bar{N}_2 = 1$. Crawford and Yurukoglu (2012) use an identification argument that is slightly more similar to mine in spirit. They identify bargaining parameters by matching the model-implied outcomes to estimated outcomes with auxiliary knowledge about one of the components of the transfer.⁴⁰ I do not have similar auxiliary knowledge, because delivery costs are nonzero and unobserved in my setting, but unlike Crawford and Yurukoglu (2012), I can leverage the optimality of the ex-ante investment that I do see.

Note further that I will use estimates from this model to decompose the effect of increasing N_1 and \bar{N}_2 , and this identification argument lends itself to using information simply within a particular level of competition. The natural endogeneity concern, discussed in Section 1.2.3, is that contests with different numbers of Phase I competitors could be unobservably different from each other. As such, using cross- N_1 restrictions for identification and estimation would be at odds with this source of endogeneity. The benefit of this identification procedure is that it depends solely on contests with a particular (N_1, \bar{N}_2) . All parameters could vary flexibly with (N_1, \bar{N}_2) .⁴¹ In practice, I have

³⁹One can think of this identification strategy as leveraging the holdup problem: if the firm is underinvesting by a large margin, we would expect that it is unable to recover much of the generated surplus.

⁴⁰In their setting, these “outcomes” correspond to channel input costs, which are negotiated in their setting. The auxiliary knowledge is that the true marginal cost is zero.

⁴¹The exception is the distribution of unobserved heterogeneity, discussed in Section 1.4.2, which could of course not be estimated separately when there is a single informative data point in a contest. An example would be when $\bar{N}_2 = 1$.

to constrain costs and the bargaining parameter to be constant across N_1 , but I let the value distribution vary flexibly with N_1 .

While the argument presented in this section is specialized to this model, the identification is robust in many senses. Appendix A.3 provides a number of extensions of this result. Proposition 1.3 can be extended almost directly to models with asymmetric firms. It extends to models with certain forms of unobserved heterogeneity, such as the one considered in the empirical model in Section 1.4.2, by utilizing methods of Fourier deconvolutions to extract information from the failure rate as a function of research effort. Finally, note that because the first-order condition (1.5) holds at all points t_2 , it embeds a number of overidentifying restrictions. Relaxing these restrictions will allow for identifying models where firms receive benefits from effort not directly tied to the Phase III contract (e.g., by developing intellectual property).⁴²

1.4.2 Empirical Model

Before discussing the empirical specification I will use to take the model in Section 1.3 to the data, I discuss the map from observables to quantities in the data. For each contest, I observe the realized number N_1 and N_2 of contestants in each Phase I and II, along with whether a firm was awarded a Phase III contract. I infer \bar{N}_2 from the 40% rule of thumb provided by the DOD SBIR program and discussed in Section 1.2.1.

Note that because all research efforts specified in the model are monetary, the map to the data is clear. The Phase III contract amount is also observed and maps directly to the bargaining transfer of $v + \eta(v - c - s)$ in the model. The Phase I and II contract amounts are mapped to $\psi(p^*)$ and $t_{N_2}^*(v)$ in the model, respectively, as described in Section 1.3.2.⁴³

I add two components to the model in Section 1.3 to take it to the data: (i) observed covariates that affect values, costs, and the costs of research and (ii) heterogeneity unobserved to the econometrician that affects all these quantities. In particular, each contest j is characterized by a set of covariates X_j and an unobserved shifter $\theta_j \sim \Theta$, where $\log \Theta$ is normalized to have mean zero. A particular firm i in a contest j has value v_{ij} , cost of Phase I research $\psi_j(p)$, and delivery cost c_{ij} given by

$$\begin{aligned} v_{ij} &\equiv \tilde{v}_i \cdot \theta_j \cdot \exp(X_j \beta), \text{ where } \tilde{v}_i \sim \tilde{V}; \\ \psi_j(p) &\equiv \theta_j \cdot \exp(X_j \beta) \cdot \tilde{\psi}(p); \text{ and} \\ c_{ij} &\equiv \theta_j \cdot \exp(X_j \beta) \cdot \tilde{c}_i, \text{ where } \tilde{c}_i \text{ has cdf } \tilde{H}(\cdot; t/(\theta_j \cdot \exp(X_j \beta))). \end{aligned} \tag{1.6}$$

The primitives to be estimated are then \tilde{V} , $\tilde{\psi}(\cdot)$, the cdf $\tilde{H}(\cdot; \cdot)$, and η . I allow \tilde{V} to depend on N_1 to control for potential endogeneity in N_1 : the DOD may choose a larger number of Phase

⁴²I also conjecture that focusing on contests with $N_2 = 1$ is not necessary either, so the argument could apply to more general models where the outcome $N_2 = 1$ need not have positive probability (e.g., for similar contests that begin in Phase II but always have at least two competitors). In addition, a combination of the overidentifying restrictions in (1.5) and the Fourier deconvolutions used with unobserved heterogeneity can accommodate certain types of shocks to the value between Phases II and III. Future versions of this analysis will provide a more formal treatment of identification in these settings.

⁴³As noted below, however, I will not directly use the observed values of $\psi(p^*)$ in the estimation procedure, because they exhibit very little variation.

I competitors for projects that have higher value (or perhaps even more uncertain value). I set $\tilde{\psi}'(p) \equiv \alpha p$ and estimate α . Finally, I restrict η to be constant across contests.⁴⁴

The empirical specification (1.6) induces a correlation between values, implementation costs, and costs of research: certain projects are more valuable to the DOD but are also more costly to implement and conduct research on. Controlling for (θ_j, X_j) , however, the residual values \tilde{v} are still mutually independent among successful firms, and the residual costs \tilde{c} are still independent of \tilde{v} (controlling for the effective expenditure on research $t/\exp(\theta_j \cdot X_j \beta)$). Thus, one interpretation of the empirical specification is that the “vertical” heterogeneity across projects, which would intuitively make more valuable projects more expensive as well, is controlled by (θ_j, X_j) . The residual heterogeneity encapsulated in \tilde{v} comes from heterogeneous match quality with the DOD, and it is thus orthogonal to the research and implementation costs. Adding unobserved heterogeneity also “softens” the hard constraint induced by the fact that Phase III does not happen if $v \leq c$: especially large transfers need not signify that values are high; rather, they may signify that the particular contest in question had an especially large value of θ_j .

The multiplicative specification in (1.6) allows me to control for heterogeneity in the estimation procedure in a structured way, as shown in the subsequent proposition.

Proposition 1.4 (Scaling). *Suppose $(p^*, \{t_{N_2}^*(\cdot)\}_{N_2 \leq \bar{N}_2})$ is an equilibrium of the R&D contest with primitives $\psi(\cdot)$, V , $C(t)$, and η . Consider a scaled model with primitives $\tilde{\psi}(\cdot) = \gamma \cdot \psi(\cdot)$, $\tilde{V} = \gamma \cdot V$, $\tilde{C}(t) = \gamma \cdot C(t/\gamma)$ (i.e., so that $\tilde{H}(c, t) = H(c/\gamma, t/\gamma)$), and $\tilde{\eta} = \eta$. Then, $(p^*, \{\gamma \cdot t_{N_2}^*(\cdot)\}_{N_2 \leq \bar{N}_2})$ is an equilibrium of the scaled contest.*

Proof. The result follows from direct substitution into the equilibrium conditions (1.2) and (1.4). \square

Proposition 1.4 is reminiscent of scaling properties of auction models and allows for a simple method of controlling for heterogeneity, although unlike in auction settings—in which there is a single dimension of heterogeneity across competitors—I consider cases where costs and values both scale.⁴⁵ By Proposition 1.4 and the specification in (1.6), we have that in equilibrium,

$$t_{N_2}^*(v_{ij}; X_j, \theta_j) = \theta_j \cdot \exp(X_j \beta) \cdot \tilde{t}_{N_2}(\tilde{v}_i) \quad (1.7)$$

for some effort function $\tilde{t}_{N_2}(\cdot)$. As will be discussed in Section 1.4.3, this specification effectively allows us to control for heterogeneity by regressing Phase II efforts on covariates.

Distributional Assumptions. While much of the model is nonparametrically identified, I place parametric restrictions to assist in estimation in finite samples. In particular, I assume that

⁴⁴In principle, the dependence of these quantities on X_j can be replaced by a general function $f(X_j)$ instead of simply $\exp(X_j \beta)$ with no change in the estimation procedure. I keep the linear notation for simplicity. Furthermore, the other parameters (such as α and η) could depend on quantities like N_1 that are coarse once again without affecting the estimation procedure.

⁴⁵See Krasnokutskaya (2011) for an example in previous work.

- V is lognormal with location parameter μ_{N_1} and scale parameter σ_{N_1} ;⁴⁶
- $H(\cdot; t)$ is a lognormal with mean parameter $\mu(t)$, which is a decreasing function of t (and the particular parameterization is discussed in detail in Appendix A.7.1) and scale parameter σ_C ; and
- $\psi(p) = \alpha p^2/2$.

I place no parametric restrictions on the distribution of θ .

I make two comments about the choice of the parametrization. First, in this model, Phase II failure is rationalized by a large draw from the cost distribution, so we would expect to estimate distributions with long upper tails to rationalize high failure rates.⁴⁷ Second, note that the identification discussion in Section 1.4.1 showed that we can identify α from the fact that $\psi'(p) = \alpha p$ purely from information about the optimality of the Phase I research effort and without any knowledge of the level of $\psi(p)$. In the empirical section, I choose *not* to use any information about the observed Phase I contract amount in the data, instead estimating the first-stage cost function based on a parametric assumption on $\psi'(\cdot)$ and the assumption that $\psi(0) = 0$.⁴⁸ I make this decision because, unlike the Phase II contract amount, the Phase I contract amount is set essentially institutionally in the DOD SBIR program and shows very little variation across projects. Thus, the Phase I contract amount may not be an accurate representation of the amount of Phase I research the firm conducts.⁴⁹ I will instead rely on the parametric assumption and compare the implied research expenditures from the model with the institutionally specified Phase I contract amount of \$80,000.

1.4.3 Estimation Procedure

One main difficulty with estimation is that the model is computationally intensive to solve, and a full-solution approach is unwieldy. However, the identification argument given in Section 1.4.1 is constructive and lends itself to a transparent estimation procedure: the identification argument highlights the *upper bound* of Phase III transfers as a function of Phase II research efforts as an object that can be directly parameterized. I embed this intuition in an MLE procedure described in this section.

With the distributional assumptions given in Section 1.4.2, I can employ a maximum likelihood approach to estimation. The overview is to (i) estimate the dependence on X_j in a first-stage regres-

⁴⁶In the case where $N_1 = 2$ and $\bar{N}_2 = 1$, there is selection into Phase II. I parameterize V as a lognormal in this case as well, but when I compute the likelihood in Step 3 of the estimation procedure described in Section 1.4.3 below, I note that the distribution of values in Phase II is a mixture between V and the maximum of two draws of V . The mixing probabilities are a function of the probability p^* of success in Phase I, which I can estimate directly.

⁴⁷I have experimented with alternate specifications in which the cost distribution is a mixture of a lognormal and a mass point at ∞ , which has probability $\gamma(t)$. This mass point is to rationalize a failure rate without necessarily resorting to a large standard deviation of the cost distribution. In practice, these specifications tend to place somewhat low mass on this “outright” failure rate and not change the estimated cost distribution appreciably.

⁴⁸I could instead use a functional form such as $\psi(p) = \alpha_0 p^2/2 + \alpha_1$, for instance, if I were interpreting the Phase I contract amounts in the data as $\psi(p)$. Note that the functional form assumption does not affect the estimates of the value or delivery cost distributions or the bargaining parameter.

⁴⁹Since Phase I contract amounts are lower than Phase II amounts, firms may be more able and willing to use internal funds to finance shortfalls in research.

sion, (ii) estimate the distribution of Θ nonparametrically using the residual correlation in Phase II bids within-contest, (iii) estimate the cost and value distribution using MLE by integrating out the estimated distribution of unobserved heterogeneity, and (iv) choose the bargaining parameter by minimizing the distance between the effort implied by the estimated parameters and the solution of the model. I restrict the sample to settings in which there is guaranteed to be no selection (i.e., I drop all contests with $(N_1, N_2) = (3, 2)$ or $(N_1, N_2) = (4, 2)$) so that Assumptions M and O are guaranteed to hold and that searching for a monotone effort function is internally consistent with the equilibrium model. Below, I spell out the steps in detail.

Step 1 (Partialling out Covariates). Taking logs of (1.7) gives

$$\log t_{N_2j}^*(v_{ij}; X_j, \theta_j) = X_j\beta + \log \theta_j + \log \tilde{t}_{N_2j}(\tilde{v}_i).$$

Thus, a regression of the log of Phase II effort on contest-level covariates returns the “normalized bids” plus the unobserved heterogeneity

$$\nu_{ij} \equiv \log \tilde{t}_{N_2j}(\tilde{v}_i) + \log \theta_j \equiv \log \tilde{t}_i + \log \theta_j,$$

along with an estimate $\hat{\beta}$ of the impact of the covariates. I then residualize the Phase III transfer by dividing by $\exp(X_j\hat{\beta})$.

Step 2 (Estimating Θ). In this step, I use a deconvolution argument standard in the auctions literature (developed by Li and Vuong (1998) and applied by Krasnokutskaya (2011)) to estimate the distribution of Θ as well as the distribution of the normalized efforts \tilde{t} for each (N_1, N_2) combination. In particular, consider pairs (ν_{i_1j}, ν_{i_2j}) from the same contest j . Since $\nu_{ij} = \tilde{t}_i + \theta_j$, with \tilde{t}_{i_1} , \tilde{t}_{i_2} , and θ_j all mutually independent and the distribution of θ_j normalized to mean zero, Kotlarski (1967) shows that the distributions of θ_j and \tilde{t}_i are identified from the joint distribution of (ν_{i_1j}, ν_{i_2j}) .

I follow Krasnokutskaya (2011) to estimate these distributions, taking into account that the empirical model in this chapter assumes that certain distributions are identical. For each pair (N_1, N_2) with $N_2 \geq 2$, I estimate the joint characteristic function of (ν_{i_1j}, ν_{i_2j}) , as well as the derivative with respect to its first argument, as the empirical means

$$\hat{\Psi}_{(N_1, N_2)}(t_1, t_2) = \frac{1}{n_{(N_1, N_2)} \cdot N_2(N_2 - 1)} \sum_{j: (N_{1j}=N_1, N_{2j}=N_2)} \sum_{i' \neq i''} \exp(it_1\nu_{i'j} + it_2\nu_{i''j})$$

$$\hat{\Psi}'_{(N_1, N_2)}(t_1, t_2) = \frac{1}{n_{(N_1, N_2)} \cdot N_2(N_2 - 1)} \sum_{j: (N_{1j}=N_1, N_{2j}=N_2)} \sum_{i' \neq i''} i\nu_{i'j} \exp(it_1\nu_{i'j} + it_2\nu_{i''j}),$$

where $n_{(N_1, N_2)}$ is the number of contests with a particular value of N_1 and N_2 , and thus $n_{(N_1, N_2)} \cdot N_2(N_2 - 1)$ is the number of pairs of observed research efforts that correspond to these auctions.

From these estimates, I recover the characteristic functions of Θ , from this subset of the data, as

$$\hat{\Phi}_{\Theta, (N_1, N_2)}(t) = k \cdot \exp \left(\int_0^t \frac{\hat{\Psi}'_{(N_1, N_2)}(0, u)}{\hat{\Psi}_{(N_1, N_2)}(0, u)} du \right), \quad (1.8)$$

where $k \equiv i\hat{\Psi}(0, 0)/\hat{\Psi}'(0, 0)$ is a factor that ensures that Θ has mean zero. I then average (1.8) over all pairs (N_1, N_2) with $N_2 \geq 2$ to compute the characteristic function

$$\hat{\Phi}_{\Theta}(t) = \frac{\sum_{(N_1, N_2): N_2 \geq 2} n_{(N_1, N_2)} \cdot N_2(N_2 - 1) \cdot \hat{\Phi}_{\Theta, (N_1, N_2)}(t)}{\sum_{(N_1, N_2): N_2 \geq 2} n_{(N_1, N_2)} \cdot N_2(N_2 - 1)}. \quad (1.9)$$

For each pair (N_1, N_2) , including those with $N_2 = 1$, I also estimate the characteristic function of ν_{ij} as

$$\hat{\Psi}_{(N_1, N_2)}(t) = \frac{1}{n_{(N_1, N_2)} \cdot N_2} \sum_{j: (N_1, N_2)} \sum_{i'} \exp(it\nu_{i'j}).$$

Then, since $\nu_{ij} = \theta_j + \tilde{t}_i$, and the characteristic function of θ_j is given by (1.9), I can compute the characteristic function of \tilde{t}_i for a particular pair (N_1, N_2) as the ratio

$$\hat{\Phi}_{\tilde{t}_i, (N_1, N_2)}(t) = \frac{\hat{\Psi}_{(N_1, N_2)}(t)}{\hat{\Phi}_{\Theta}(t)}.$$

The densities of Θ and \tilde{t} can be recovered from the Fourier inversion formula

$$\begin{aligned} f_{\log \Theta}(\theta) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(-it\theta) \hat{\Phi}_{\Theta}(t) dt \\ f_{\log \tilde{t}_{(N_1, N_2)}}(u) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(-itu) \hat{\Phi}_{\tilde{t}_i, (N_1, N_2)}(t) dt. \end{aligned} \quad (1.10)$$

In practice, these integrals are approximated on a compact interval $[-T, T]$ chosen in a data-driven fashion. See Appendix A.7.2 for details. Finally, densities for θ and \tilde{t} (for each (N_1, N_2)) can be recovered from transforming the recovered densities for $\log \theta$ and $\log \tilde{t}$.

Step 3 (Maximum Likelihood Estimation of Phase II Parameters). The next step involves maximizing the likelihood of observing the Phase II and III data, integrating out over the distribution of unobserved heterogeneity estimated in Step 2. In particular, I maximize over the distributions of \tilde{V} and the cost distribution $\tilde{H}(\cdot; \cdot)$, fixing the bargaining parameter η .

For each candidate value of these parameters, I first approximate an implied effort function by appealing to Proposition 1.1: because efforts are one-to-one with values, fixing (N_1, N_2) , a firm with a value in the q^{th} quantile of the distribution of \tilde{V}_{N_1} will exert effort in the q^{th} quantile of the distribution of $\tilde{t}_{(N_1, N_2)}$, which was estimated in Step 2. Thus, for a candidate value of the distribution of values \tilde{V} , I can approximate the inverse effort function $\tilde{v}(\cdot)$ without solving the model directly, for $\theta_j = 1$.

Fix a particular contest j and guess a θ_j . The map \tilde{v} allows one to compute $\tilde{v}_{ij}(\theta_j)$ as $\tilde{v}(\nu_{ij}/\theta_j)$ for all firms i . The likelihood of drawing these values is

$$L_{\text{values},j}(\theta_j) = \prod_i f_{\tilde{v}}(\tilde{v}_{ij}(\theta_j)) \cdot \frac{\tilde{v}'_{ij}(\theta_j)}{\theta_j},$$

where the second term takes into account the Jacobian of the transformation. If the contest does not enter Phase III, then it must be that all firms drew costs larger than their values. Thus, the likelihood of observing this outcome is

$$L_{\text{Phase III},j}(\theta_j) = \prod_i [1 - H(\tilde{v}_{ij}(\theta_j); \nu_{ij}/\theta_j)].$$

If instead we do observe a Phase III transfer for firm i^* , then we can compute the likelihood of observing this transfer as

$$L_{\text{Phase III},j}(\theta_j) = \int_{\underline{s}}^{\tilde{v}_{i^*j}(\theta_j) - t_{3j}/\theta_j} \frac{1}{\theta_j \cdot (1 - \eta)} h\left(\frac{t_{3j}/\theta_j - \eta \tilde{v}_{i^*j}(\theta_j) + \eta \max\{s, 0\}}{1 - \eta}; \nu_{i^*j}/\theta_j\right) f_{\bar{S},j}(s) ds dc,$$

where $f_{\bar{S},j}(s)$ is the pdf of the maximum value of the surplus for all firms other than i^* , which is computed based on the observed values of ν_{ij} and the posited candidate $\tilde{H}(\cdot; \cdot)$. Note that if for the posited θ_j , the value of the Phase III transfer exceeds the implied $\tilde{v}_{ij}(\theta_j)$, then the likelihood is zero. Finally, if a Phase III transfer is observed that is implausibly low (less than \$1 million in this specification), I assume that the project succeeded but that the actual value of the Phase III transfer is unobserved. In this case, if i^* is awarded the contract, then it must be that i^* drew a cost less than its value and that the surplus generated by all other competitors is less; the likelihood $L_{\text{Phase III},j}(\theta_j)$ can be computed accordingly.

We can then compute the log likelihood over observing this outcome as

$$\log \int L_{\text{values},j}(\theta) \cdot L_{\text{Phase III},j}(\theta) \cdot f_{\theta}(\theta) d\theta,$$

integrating out against the distribution of unobserved heterogeneity. We maximize the sum of this log likelihood across all contests that enter Phase II. Computational details are given in Appendix A.7.

Step 4 (Estimation of the Bargaining Parameter). So far, estimation has only relied on Assumption M and Proposition 1.4. The identification argument, however, noted that information on the bargaining parameter comes from the firm's first order condition. In this step, I impose the firm's first order condition by solving the model explicitly. I do so at each value of η on a fine grid, at the estimated parameters from Step 3.⁵⁰ I then use a simulated method-of-moments procedure

⁵⁰While it is infeasible to use a full solution approach during the MLE procedure—and while it imposes more structure than would be necessary for estimation of the baseline model—solving the model is possible on a fine grid of η once the MLE estimates have been computed for each of those values of η . The estimation as well as the model

to match properties of the observed data with simulated values from each of the solved models for the various values of η .

In particular, for each contest j , let \hat{f}_j be an indicator for whether the contest failed before entering Phase III. Let \hat{t}_{3j} be the observed Phase III transfer. Since this quantity is undefined for contests that do not enter Phase III, I instead define the moment \hat{t}'_j to be 0 if the contest fails and \hat{t}_{3j} if it does not. I match these to the empirical counterparts, which are the simulated probability of failure $\widetilde{\Pr}(\text{failure}; \eta, \theta^*(\eta))$, where $\theta^*(\eta)$ are the MLE estimates conditional on η from Step 3, and the (partial) expectation of the observed transfer, $(1 - \widetilde{\Pr}(\text{failure}; \eta, \theta^*(\eta))) \cdot \widetilde{\mathbb{E}}[t_3; \eta, \theta^*(\eta)]$. I match these moments conditional on (N_1, N_2) . Thus, for a contest j , the relevant set of moments is

$$g_j(\eta) = \begin{pmatrix} \hat{f}_j - \widetilde{\Pr}(\text{failure}; \eta, \theta^*(\eta)) \\ \hat{t}'_j - (1 - \widetilde{\Pr}(\text{failure}; \eta, \theta^*(\eta))) \cdot \widetilde{\mathbb{E}}[t_3; \eta, \theta^*(\eta)] \end{pmatrix} \otimes \mathbb{1}_{(N_1, N_2)} \equiv \hat{g}_j - \bar{g}_j(\eta),$$

where $\mathbb{1}_{(N_1, N_2)}$ is a vector that contains a 1 in the element corresponding to (N_1, N_2) and zeros elsewhere.

For brevity, replace \hat{g}_j by the Kronecker product of \hat{g}_j and the dummy vector $\mathbb{1}_{(N_1, N_2)}$. Then, the optimal method-of-moments procedure to estimate η corresponds to

$$\eta^* = \arg \min_{\eta} \left(\sum_j g_j(\eta) \right)' \cdot \hat{\Omega}^{-1} \cdot \left(\sum_j g_j(\eta) \right), \quad (1.11)$$

where

$$\hat{\Omega} = \sum_j (\hat{g}_j - \bar{g})(\hat{g}_j - \bar{g})'$$

and \bar{g} is the empirical mean of \hat{g}_j . In practice, I evaluate η on the fine grid and pick the minimum value of the objective in (1.11).⁵¹

Step 5 (Estimation of the Phase I Parameter). For each value of (N_1, \bar{N}_2) , I use maximum likelihood to estimate the probability $\hat{p}_{(N_1, \bar{N}_2)}$ of a particular contestant succeeding when there are N_1 contestants in Phase I and a limit of \bar{N}_2 on Phase II. In particular, I estimate a censored binomial model in which for each contest j , the unobserved number of successes N_{Sj} is such that $N_{Sj} \sim \text{Binomial}(N_1, \hat{p}_{(N_1, \bar{N}_2)})$, but the observed quantity is

$$N_{2j} = \begin{cases} N_{Sj} & \text{if } N_{Sj} \leq \bar{N}_{2j} \\ \bar{N}_{2j} & \text{if } N_{Sj} > \bar{N}_{2j} \end{cases}.$$

solution can be parallelized conditional on η , whereas a full solution approach as part of an optimization procedure that involves η might have to be run in sequence.

⁵¹As noted in Step 3, there are a set of contests that I treat as a success with the Phase III amount unobserved. When constructing moments, I account for these contests when computing failure rates, but I ignore them when computing the means of the transfer distribution.

Upon estimating $\hat{p}_{(N_1, \bar{N}_2)}$, I compute the profits from Phase II by solving the model using the estimated parameters from Step 4 for all values of N_1 at the estimated $p_{(N_1, \bar{N}_2)}^*$. I then use the FOC associated with (1.4) as the estimating equation for α . In particular, I set

$$\alpha^* = \arg \min_{\alpha} \sum_{(N_1, \bar{N}_2)} w_{(N_1, \bar{N}_2)} \left[\psi' \left(\hat{p}_{(N_1, \bar{N}_2)}; \alpha \right) - \hat{\pi} \left(N_1, \bar{N}_2; \eta^*, \theta(\eta^*) \right) \right]^2,$$

where $\hat{\pi}(N_1, \bar{N}_2; \eta^*, \theta(\eta^*))$ is the expected profit conditional on success and $w_{(N_1, \bar{N}_2)}$ is a weighting function. I set the weight equal to the number of contests with (N_1, \bar{N}_2) .

1.5 Structural Estimates

In this section, I discuss the structural estimates of the value and cost variation, the research production functions, and the bargaining parameter in the model. These estimates together will allow me to summarize the type of heterogeneity that governs the outcome of these contests. I will then briefly present some information about the fit of the model to the data. Table 1.5 reports the parameter estimates for the equilibrium model of the R&D contests, following Steps 1–5 of the procedure outlined above.⁵² Appendix A.2.2 provides estimates of the Phase II parameters conditional on particular values of η , using only Assumption M and an analogue of the scaling property of Proposition 1.4, for comparison.

Panel (a) of Table 1.5 shows quantiles of the distribution of unobserved heterogeneity θ . The distribution is fairly concentrated around 1, suggesting that unobserved heterogeneity does not play an especially large role in the data. A contest in the 10th percentile of the data has values and costs that are about 70% of the median contest, and a contest in the 90th percentile has values that are about 35% larger than median. There is a somewhat large range, however: moving from the 2.5th percentile to the 97.5th percentile increases values and costs by a factor of 5.

Panel (b) of Table 1.5 shows the mean as well as measures of the variance in the value distributions as a function of N_1 . While the structural estimation procedure estimates the distribution of \tilde{v} , I scale these estimates by the mean value of the estimates of $\exp(X_j\beta)$ from Step 1 as well as the estimated mean for the distribution of θ from Step 2 to express these numbers in millions of dollars. The first observation is that average projects have mean values of around \$11.0–\$15.0 million dollars. Projects in which the DOD selects a larger number of Phase I competitors tend to have larger values, although the difference in the values is somewhat imprecisely estimated. Note that these values are for all projects, so the ones that do result in Phase III contracts will be selected from the upper tail of this distribution. The second observation is that these value distributions are fairly narrow. One measure of this variation is the standard deviation, which is estimated to be

⁵²I construct standard errors by a nonparametric bootstrap. I sample with replacement from the dataset, making sure that the distribution of contests with a particular (N_1, \bar{N}_2) remains fixed across bootstrap samples. I then repeat Steps 1–3, conditional on the value of η^* picked in Step 4. In future work, I will compute a standard error on η^* as well, but the analogous bootstrap procedure would require me to estimate the model on a fine grid of η for each bootstrap sample, which is especially time consuming.

Percentile	2.5%	10%	25%	50%	75%	90%	97.5%
θ	0.387 (0.090)	0.728 (0.060)	0.876 (0.039)	1.012 (0.018)	1.165 (0.039)	1.346 (0.103)	1.938 (0.321)

(a) Quantiles of the distribution of unobserved heterogeneity Θ .

Values (\$M)	$N_1 = 1$	$N_1 = 2$	$N_1 = 3$	$N_1 = 4$
Mean	10.98 (4.09)	11.96 (2.76)	13.20 (2.88)	14.94 (2.90)
Standard Deviation	0.34 (0.13)	0.36 (0.09)	0.40 (0.09)	0.46 (0.09)
95% Range	1.32 (0.51)	1.41 (0.34)	1.55 (0.37)	1.79 (0.36)

(b) Moments of the value distribution, in millions of dollars

$\Pr(c < v)$		$\mathbb{E}[c c < v]$ (\$M)		Quantiles (\$M)			
Value	Semi-Elasticity	Value	Elasticity	1%	5%	10%	Elasticity
0.071 (0.010)	0.012 (0.004)	6.85 (0.91)	-0.016 (0.005)	2.85 (0.40)	9.27 (1.30)	17.39 (2.43)	-0.161 (0.046)

(c) Moments of the cost distributions, averaged over both the observed distribution of N_1 and efforts as well as the estimated distribution of unobserved heterogeneity.

Firm Bargaining Parameter (η)	0.73
Phase I Marginal Cost (α)	0.208 \$M
Average Phase I Cost	0.027 \$M

(d) Phase I and bargaining parameters

Table 1.5: Structural estimates

about \$350,000–\$450,000. Given the lognormal distribution, these estimates correspond to a “95% range,” i.e., the difference between the 97.5th percentile and the 2.5th percentile, or about \$1–\$2 million. Thus, the extent of the variation in values is approximately 12% of the mean.

The identification argument in Section 1.4.1 can shed some light on the moments in the data that influence these estimates. Most of the observed Phase III transfers lie below the 95th percentile of the estimated values (as seen in Figure 1-1(b), for instance), and in this sense, the Phase III values serve as an upper bound for the transfer distribution: points beyond this upper bound are explained by the heterogeneity encapsulated by X and θ . The slope of this “soft” upper bound provides information about the variance in the value distribution: the fact that even projects with low levels of Phase II funding tend to occasionally have reasonably high Phase III contract

amounts suggests that these projects have reasonably high values as well. Of course, due to the parametric assumptions and the introduction of heterogeneity, the estimates of values are influenced by matching the failure rate as well, which depends on the cost estimates I discuss in the next part of this section.

Panel (c) of Table 1.5 shows the estimates related to the delivery cost distributions. Since the delivery cost depends on research effort, which varies across the sample, I aggregate across all data points to report these numbers.⁵³ In particular, I fix a value of unobserved heterogeneity θ and compute moments of the cost distribution at the implied value of $\tilde{t}_{2ij} = t_{2ij}/\theta$ for *each* contestant i in each contest j . I compute the moments of interest for all these data points; I then average across all these data points and integrate out over θ . I scale the estimates to bring the appropriate ones into units of millions of dollars.

Just as values are positively selected, the cost draws are negatively selected, conditional on success; because so few contests succeed in Phase II, the mean cost draw is irrelevant for observables. I instead report (i) the probability that the cost draw is less than an independent value draw (for the associated value of N_1), (ii) the conditional expectation of cost draws that are less than value draws, and (iii) some relevant quantiles of the cost distribution. The probability that costs are less than values is about 0.07, which is slightly lower than the observed success rate. The mean of these cost draws is about \$6.9 million. The 1st percentile of the *unconditional* cost distribution is about \$2.9 million, the 5th percentile about \$9.3 million, and the 10th percentile about \$17.4 million. I also report elasticities of these quantities, which are estimated to be rather low. The elasticity of the quantiles with respect to research effort is about 0.2: if research efforts increase by 1%, the quantiles of the delivery cost distribution decrease by 0.2%.⁵⁴ This value translates to an elasticity of about 0.016 for the conditional expectation of costs and a semi-elasticity of 0.012 for the probability that the cost draw is less than the value draw.⁵⁵

Conditional on the bargaining parameter, the cost distributions are estimated from two main patterns in the data. First, the failure rate decreases with research effort, and the rate of this decrease—after accounting for the increase in the value estimated above—as well as the failure rate itself, affect the distribution and the elasticity. At the same time, the observed transfers do increase with the Phase II amount, which must be due to the increase in the values; because a decrease in the cost would counteract this effect, the estimated elasticity cannot be so high as to cause the observed transfers to drop.

Finally, Panel (d) of Table 1.5 provides a few remaining statistics related to the model. First, the firms' bargaining parameter is estimated to be 0.73, meaning the DOD gives the winning firm about three-fourths of the (incremental) surplus generated from the project. This estimate

⁵³Standard errors do not account for this variation across the sample.

⁵⁴It is a property of the lognormal, together with the fact that the research effort only parameterizes the mean, that this elasticity is uniform across quantiles.

⁵⁵Note that $\Pr(c < v)$ is not exactly a failure rate, since I compare the cost draw to a generic draw from the value distribution. Similarly, the elasticities are lower in magnitude than the rate of change of the failure rate (conditional expectation of costs) with respect to the research efforts, since v would increase as well. Rather, these moments are descriptive features of the cost distributions themselves, and I relegate quantities such as actual failure rates to the discussion of model fit.

directly uses information about the firm choosing research efforts optimally.⁵⁶ It is determined by fitting the equilibrium transfers and failure rates. Roughly, a larger value of η would overpredict the transfers (by bringing them closer to the value of the project) and reduce the failure rate by increasing the incentives to conduct research. Panel (d) also reports the estimate of α (in dollars per unit probability), averaged across values of N_1 . A one percentage-point increase in the probability of success costs roughly \$2,000, an estimate that is obtained directly from equating the marginal cost of research to the expected profits at the observed success rates. Using the additional functional form assumption that $\psi(p) = \alpha p^2/2$, the Phase I expenditure amounts to approximately \$27,000; the estimate is slightly higher when restricting to contest with $N_1 = 1$ (about \$43,000) or $N_1 = 4$ (about \$66,000). While these values are slightly lower than the DOD-specified amount of \$80,000, they are nevertheless in the right ballpark. This agreement provides suggestive evidence in favor of the model, especially given that the estimation uses absolutely no information about the Phase I contract amount. The lower model-implied estimates may be due to misspecification of the functional form, or perhaps a fixed cost of research should be included in this function; alternatively, the SBIR program may simply wish to set an institutional amount that is guaranteed to cover costs for a wide range of projects. Throughout the rest of the chapter, I will maintain this functional form, with the caveat that I may be underestimating the cost of Phase I research slightly.

Figure 1-2 plots two further measures of model fit. Panel (a) plots the observed failure rate (from Phase II to Phase III) for bins of (N_1, N_2) against the model-implied failure rate. The observed and implied failure rates are close to each other, although the model does have a tendency of overpredicting failure. Panel (b) plots the observed contract amounts against the model-implied values; this fit is especially close, and the model predictions are within one standard error of the data in almost all the bins.

The parameter estimates in this section paint a picture of these contests as ones with moderate value (about \$11–\$15 million) but without especially large variation across competitors within contests. This result is consistent with the notion that these projects are well-specified ex-ante, and, while there is heterogeneity across contestants, there is not much room for innovation on the dimension of project quality. Delivery costs, however, are substantively different across firms, although the map from research effort to the realization of the cost is estimated to be rather flat. Finally, the DOD does allow the firms to capture a fairly large portion of the surplus they generate, giving them the incentives to conduct research throughout the contest.

1.6 Social Inefficiency in R&D Contests

In this section, I use the estimated values from Section 1.5 and ask whether the equilibrium of the R&D contests features underprovision or overprovision of R&D from a social standpoint. In doing so, I will discuss the potential sources of inefficiency, and this discussion will help interpret the

⁵⁶Of course, this bargaining parameter does feed into the value and cost estimates, but not the unobserved heterogeneity distributions, from above.




Figure 1-2: Observed and model-implied (a) failure rates and (b) Phase III contract amounts. The error bars indicate *one* standard error; the point in (b) without standard error bars corresponds to a bin with just one data point.

design counterfactuals studied in Sections 1.7 through 1.9. I also provide a measure of the optimal social surplus to quantify the surplus left on the table due to the current design of the contest, and the optimum provides a benchmark to which we can compare the surplus generated in these alternate designs. Note that social surplus is defined to be the maximum value of $(v - c)^+$ generated by the contestants in Phase II, less total research costs in Phases I and II. This section as well as Sections 1.7 through 1.9 focus on social surplus as the outcome of interest; Section 1.10 considers the profits of the DOD as well, and Appendix A.1 goes into more details about DOD profits.

To assess the social efficiency of research efforts in this contest, I conduct the following experiment. First, I compute the equilibrium of the R&D contest; denote the first-stage effort by p^* and the second-stage effort by $t_{N_2}^*(\cdot)$. I then compute the optimal second-stage effort function $\hat{t}_{N_2}(\cdot)$ of the form $\gamma \cdot t_{N_2}^*(\cdot)$; I vary γ and keep p^* fixed. An optimal value of $\gamma > 1$ would suggest that research is underprovided in the equilibrium of the R&D contest while a value of $\gamma < 1$ would suggest that there is excessive research (holding first-stage behavior fixed). In the next experiment, I compute the optimal first-stage entry probability \hat{p} , keeping the second-stage effort function fixed at the equilibrium level. I compare this value to p^* . Table 1.6 shows the social surplus in the equilibrium of the R&D contest, together with the optimal values of γ and \hat{p} and the social surplus at these values, for various values of (N_1, \bar{N}_2) .⁵⁷

What are the sources of inefficiency in Phase II? If $N_2 = 1$, we can note that the firm's problem and the social planner's problem coincide when $\eta = 1$. Setting $\eta = 1$ makes the firm the sole claimant to the generated surplus and effectively amounts to selling the project to the firm, maximizing social surplus.⁵⁸ Indeed, the only source of inefficiency in Phase II with $N_2 = 1$ is the *holdup problem*, because the party that invests in research only receives part of the surplus. Thus,

⁵⁷I use the parameters for the associated value of N_1 .

⁵⁸It is in fact possible to show that social surplus is monotonically increasing in η in this case as well.

N_1	\bar{N}_2	Baseline		Phase I			Phase II			Optimum	
		p^*	SS	\hat{p}	SS	%	γ	SS	%	SS	%
1	1	0.40	0.012	0.35	0.013	4.2%	1.50	0.013	8.1%	0.014	11.7%
2	1	0.42	0.028	0.31	0.031	10.2%	1.47	0.031	9.5%	0.033	18.1%
3	2	0.63	0.111	0.51	0.119	6.9%	1.44	0.119	7.2%	0.125	13.0%
4	2	0.77	0.246	0.54	0.299	21.5%	1.40	0.258	5.0%	0.310	25.9%

Table 1.6: Baseline social surplus (SS) and first-stage effort p^* , along with socially optimal levels of first stage effort \hat{p} and scaling factor γ for second-stage effort for various values of (N_1, \bar{N}_2) . Values of $\hat{p} < p^*$ imply that Phase I research is socially excessive in the R&D contest, and values of $\gamma > 1$ suggest that Phase II research is underprovided in equilibrium. The final columns report the optimum surplus, in which $\eta = 1$ and p is chosen to maximize surplus subject to $\eta = 1$. I also report percent improvements in surplus relative to the baseline design of the contest.

we would expect Phase II research efforts to be underprovided in the R&D contest when $N_2 = 1$. Accordingly, Table 1.6 notes that in the cases where $\bar{N}_2 = 1$ (so that N_2 must equal 1 when Phase II occurs), the socially efficient level of R&D is about 47–50% larger than the equilibrium level of R&D. Across contests with different (N_1, \bar{N}_2) , this amounts to a gain in social surplus of around 5–10% relative to the R&D contest, which can be interpreted as a ballpark estimate of the “cost of holdup” in this setting.⁵⁹

A less obvious implication of the model is that a similar conclusion holds for $N_2 > 1$ as well: the social planner’s optimum is supportable by the firms in equilibrium if $\eta = 1$. The key observation driving this result is that in Phase II, the winning firms’ profit (ignoring research costs) is η times the difference between the surplus from his project and the surplus from the next-best project. When $\eta = 1$, this difference is exactly the winner’s marginal contribution to the social surplus, meaning the firm is rewarded in a manner that coincides with the social planner’s objective function. Thus, for Phase II, the social planner would always prefer to set $\eta = 1$.⁶⁰ I codify this argument in the following proposition, whose proof is in Appendix A.5.2.⁶¹

Proposition 1.5. *Consider a contest that begins in Phase II. The social planner’s solution (when the social planner is constrained to choose effort schedules that depend only on an individual competitor’s value) can be supported by a competitive equilibrium when $\eta = 1$. Moreover, if there is exactly one competitor, the social surplus is monotonically increasing in η .*

Table 1.6 also shows that values of γ for cases where $\bar{N}_2 = 2$. In these situations, we would have

⁵⁹As a point of comparison, we could consider an alternate experiment where we subsidize Phase II research by τ so that when the firm invests t in R&D, it gets an additional τt . The firm internalizes this subsidy when making R&D decisions. Then, if $1 + \tau = 1/\eta$, we would have the efficient level of investment. At the estimated value of $\eta = 0.73$, this corresponds to a 37% subsidy.

⁶⁰That the social surplus is maximized at $\eta = 1$ with $N_2 > 1$ does of course depend on the adopted bargaining procedure for Phase II. In an alternate bargaining procedure in which the firm with the highest v is approached and given the transfer $c + \eta(v - c)$ regardless of the outcomes of other firms, there would be a business-stealing effect.

⁶¹See Hatfield, Kojima, and Kominers (2016) for a discussion of this issue in general models. Note further that not all equilibria of the contest at $\eta = 1$ need to coincide with the social planner’s problem. However, I have no numerical evidence of multiple equilibria.

instances where $N_2 = 1$ or $N_2 = 2$. Since in both cases we expect underinvestment, it is expected that $\gamma > 1$ here as well. I find magnitudes similar to the instances when $\bar{N}_2 = 1$: socially efficient research efforts would be about 40–45% larger than the ones in the R&D contest, and the social surplus would increase by about 5–7.5% off the baseline in the R&D contest.

A different story emerges when considering the full contest, starting at Phase I. First, when exerting effort in Phase I, the firm internalizes the fact that its Phase II research efforts will be refunded by the DOD contract. As such, even at $\eta = 1$ for $N_1 = 1$, the social planner’s problem does not coincide with the firm’s. The *reimbursement effect*, in which later-stage research expenditures are not internalized when early-stage expenditures are decided, would lead to overprovision of Phase I research efforts. The second effect—which is arguably more robust and present in general models of R&D—is analogous to a *business-stealing effect* from Mankiw and Whinston (1986b): when setting research efforts, a firm does not internalize the loss to its rival when it displaces it from entering into Phase II. Of course, this business-stealing effect only exists for $\bar{N}_2 < N_1$. This effect would also lead to overprovision of R&D. Finally, we have the *holdup effect* that also exists in Phase II; this would point towards underprovision of R&D in Phase I. The net effect thus depends on the cumulative magnitude of these three effects and is in principle ambiguous.

Comparing the equilibrium p^* in the R&D contest to the socially optimal \hat{p} in Table 1.6 suggests that there is overprovision of R&D in equilibrium. The sum of the reimbursement effect and the business-stealing effect (or just the reimbursement effect when $N_1 = 1$) outweighs the holdup effect. In all cases, $\hat{p} < p^*$, and the social surplus given \hat{p} is between 4% (for $N_1 = 1$) and 22% (for $N_1 = 4$) larger than under the equilibrium. The final two columns of Table 1.6 show the optimal social surplus, subject to the information constraints that the agents face. In particular, I set $\eta = 1$, so that social surplus is maximized in Phase II, and then simultaneously choose the effort p in Phase II to maximize the total surplus generated in the entire contest.⁶² These columns provide a measure of the surplus left on the table due to the current design of the contest and provide a benchmark against which the design counterfactuals discussed in the subsequent sections can be compared. Social efficiency can be improved by between 11% and 26% by setting this optimal design.

1.7 The Effect of Early- and Late-Stage Competition

The addition of a competitor to an R&D contest has two effects. First, there is a *direct* effect of getting another draw from the pot, albeit at some additional cost of research. Second, there is an indirect *incentive* effect in that the equilibrium effort exerted by the firms changes. Due to both the cost of research and to this incentive effect, it may be optimal to limit entry into R&D contests, and stylized theoretical models sometimes show that it even can be optimal to restrict competition to two competitors. In this section, I quantify the effect on social surplus of adding competitors in

⁶²This corresponds to a situation where the social planner chooses effort as a function of value and has beliefs about other agents’ values that coincide with those of the agents. One could also compute a “first best,” in which the planner can condition research efforts on the vector of realizations of values. The first best does not increase the surplus appreciably for these parameters, and I focus on the “second best” shown in Table 1.6 throughout the chapter because it is more closely related to the current design of the contest.

	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	-0.024	0.129			$N_1 = 2$	0.022	0.188		
$N_1 = 3$	-0.022	0.099	0.247		$N_1 = 3$	0.025	0.232	0.369	
$N_1 = 4$	-0.019	0.102	0.218	0.354	$N_1 = 4$	0.029	0.235	0.418	0.545

(a) Change in social surplus (Baseline is 0.144 \$M)

(b) Change in total research costs (Baseline is 0.195 \$M)

Table 1.7: Total effects of moving from a baseline of $N_1 = \bar{N}_2 = 1$ to various values of (N_1, \bar{N}_2) on (a) social surplus and (b) total research costs. Each entry in the table lists the *change* from the baseline value, and the baseline values are listed in the respective captions. All values are in millions of dollars.

both the early (Phase I) and late (Phase II) stages of the program. I then decompose this effect into direct and incentive effects to quantify the contribution of the equilibrium change in incentives to social surplus.

1.7.1 Changing N_1 and \bar{N}_2

I first compute the *total* effect of changing the number of competitors in the contest, using $N_1 = \bar{N}_2 = 1$ as a baseline. Table 1.8 presents the results.⁶³ At these parameters, the expected social surplus generated per contest is \$144,000. About \$200,000 of this is due to R&D cost reimbursements, so each contest generates about \$340,000 of surplus when ignoring research costs.

Table 1.7(a) shows the total effect on social surplus of going from a contest with $N_1 = \bar{N}_2 = 1$ to different values of N_1 and \bar{N}_2 . Social surplus drops by approximately \$24,000 when keeping $\bar{N}_2 = 1$ but increasing N_1 to 2, and by slightly less if instead increasing it to 3 or 4.⁶⁴ While I will discuss these numbers in more detail in Section 1.7.2, some rough intuition is as follows: increasing N_1 without increasing \bar{N}_2 reduces each individual competitor's incentive to exert Phase I effort (which is socially beneficial, as Section 1.6 shows that there is overprovision of R&D in Phase I) but does lead to larger *total* Phase I effort expenditures.⁶⁵ However, much of the failure rate is due to failure in Phase II, and limiting entry to exactly one competitor in Phase II only leverages the benefit of having a larger value draw. This benefit, especially given the fairly narrow estimated value distributions, is not large enough to counteract the additional cost of Phase I research.

Increasing the limit \bar{N}_2 on Phase II competition improves the chances of success in Phase II, albeit at the cost of more research. Whether this increase is socially beneficial depends on the extent to which two competitors in Phase II are “ex-ante” substitutes. Since Phase II failure rates are so high in this setting, firms are effectively *not* substitutes; the two firms would only be substitutable in the event that they both succeed, which is very unlikely. Thus, we would expect that if inviting one firm to Phase II is socially beneficial (as it is because the social surplus is positive

⁶³I use the parameter estimates with $N_1 = 4$ in this section.

⁶⁴Fixing \bar{N}_2 , social surplus need not be monotone in N_1 , as seen from this example. Fixing N_1 , social surplus also need not be monotone in \bar{N}_2 , although that is not clear in these numerical examples.

⁶⁵Moving from $N_1 = 1$ to $N_1 > 1$ does introduce a business-stealing effect in Phase I as well, which leads to further overprovision of Phase I R&D.

when $N_1 = \bar{N}_2 = 1$), inviting more firms would be socially beneficial as well. Accordingly, we see that social surplus increases (almost) linearly when we increase both N_1 and \bar{N}_2 by 1, starting from $N_1 = \bar{N}_2 = 1$: moving from $N_1 = \bar{N}_2 = 1$ to $N_1 = \bar{N}_2 = 2$ increases social surplus by \$130,000 (slightly less than the base of \$144,000). Adding one more competitor to each stage increases it by the slightly lower value of \$118,000. This slight decrease is due to the fact that the firms become slightly more substitutable as competition increases.⁶⁶ In addition, there are effects on equilibrium incentives that I will discuss in Section 1.7.2, but the fact that research efforts increase almost linearly (see Table 1.7(b)) suggest that they are quite small.

I only show the total effect for relatively small contests, but a takeaway message from this counterfactual is that the social planner would want to invite more firms to enter *both* phases of the competition. Indeed, the optimal numbers of Phase I and Phase II contestants at these estimated parameters are both larger than 4.⁶⁷

1.7.2 Decomposing the Effect of Competition

I now decompose the total effects presented in the previous subsection into the direct effect of adding competitors and the indirect effect that competitors can change their equilibrium effort. Note that in a multistage contest like the one considered in this chapter, the design variables are the number N_1 of competitors in the first stage and the *limit* \bar{N}_2 on the competitors in the second stage. For concreteness, I will define direct and incentive effects relative to $N_1 = 1$ and $\bar{N}_2 = 1$ to make the decomposition comparable to the computations in the previous subsection.

Consider a contest characterized by an arbitrary (N_1, \bar{N}_2) and consider any outcome denoted $S(N_1, \bar{N}_2, p, \{t_{N_2}(\cdot)\}_{N_2 \leq \bar{N}_2})$, defined as a function of the number N_1 of Phase I participants, the limit \bar{N}_2 of Phase II participants, effort p in Phase I, and the effort function $t_{N_2}(\cdot)$ for all possible realizations of N_2 . In equilibrium, the firms would exert the effort level $p_{(N_1, \bar{N}_2)}^*$ and the effort functions $t_{N_2}^*(\cdot; p_{(N_1, \bar{N}_2)}^*)$. The decomposition I conduct in this section of the total effect of moving from a contest with just one contestant to one with N_1 Phase I contestants and \bar{N}_2 Phase II

⁶⁶Note for reference that the case $N_1 = \bar{N}_2$ does not feature a business-stealing effect in the first stage, so there is one less force towards R&D being excessive in Phase I.

⁶⁷Computing the equilibrium for $\bar{N}_2 \geq 3$ and $\bar{N}_2 < N_1$ is increasingly cumbersome because beliefs require integrating over dimension a joint density of dimension $\bar{N}_2 - 1$. Furthermore, the exact optimal number would be influenced by whether or not $\psi(\cdot)$ has a fixed cost, for instance. I thus do not search for the optimal (N_1, \bar{N}_2) and instead note that the robust conclusion is that the social planner prefers to increase competition.

	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	-0.083	-0.083			$N_1 = 2$	-	0.215		
$N_1 = 3$	-0.175	-0.175	-0.175		$N_1 = 3$	-	0.232	0.429	
$N_1 = 4$	-0.270	-0.270	-0.270	-0.270	$N_1 = 4$	-	0.239	0.453	0.636
(a) Direct (Phase I)					(b) Direct (Phase II)				
	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	0.059	0.000			$N_1 = 2$	0.000	-0.002		
$N_1 = 3$	0.153	0.045	0.000		$N_1 = 3$	0.000	-0.002	-0.007	
$N_1 = 4$	0.250	0.135	0.042	0.000	$N_1 = 4$	0.000	-0.002	-0.007	-0.012
(c) Incentive (Phase I)					(d) Incentive (Phase II)				

Table 1.8: Decomposition of the total change in social surplus from changing the number of competitors in Phase I (N_1) and the limit on the number of competitors allowed to enter Phase II (\bar{N}_2), following (1.12). All values are in millions of dollars, and the baseline value of social surplus (at $N_1 = \bar{N}_2 = 1$) is \$144,000.

contestants is

$$\begin{aligned}
& \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_{N_2}^*(\cdot; p_{(N_1, \bar{N}_2)}^*)\}_{N_2 \leq \bar{N}_2}\right) - S\left(1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{total effect}} \\
&= \underbrace{S\left(N_1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right) - S\left(1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{direct effect of Phase I competition}} \\
&+ \underbrace{S\left(N_1, \bar{N}_2, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right) - S\left(N_1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{direct effect of Phase II competition}} \\
&+ \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_1^*(\cdot)\}\right) - S\left(N_1, \bar{N}_2, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{incentive effect from Phase I competition}} \\
&+ \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_{N_2}^*(\cdot; p_{(N_1, \bar{N}_2)}^*)\}_{N_2 \leq \bar{N}_2}\right) - S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_1^*(\cdot)\}\right)}_{\text{incentive effect from Phase II competition}}. \quad (1.12)
\end{aligned}$$

In words, the direct effect of Phase I competition simply considers the effect of adding Phase I competitors without changing any equilibrium efforts. The direct effect of Phase II competition subsequently increases the maximum allowed Phase II competition, again without any change in equilibrium efforts. In the cases in which multiple competitors enter Phase II, I assume they all exert effort following the schedule $t_1^*(\cdot)$; in this way, I separate the impact of competition on Phase II outcomes. The incentive effect from Phase I competition then allows for firms to readjust their research efforts *in Phase I* to the final equilibrium effort given by the new competitive structure. Finally, the incentive effect from Phase II competition allows firms to readjust their Phase II efforts

and finally arrives at the new equilibrium.

Table 1.8 quantifies these four effects. Panel (a) shows the direct effect of adding Phase I competitors; note that this quantity is definitionally independent of \bar{N}_2 . As discussed above, increasing N_1 without increasing \bar{N}_2 simply increases total Phase I expenditures and increases the value of the Phase II competitor slightly, but it does not improve the probability of success in Phase II appreciably. Thus, the direct effect of adding a Phase I competitor is negative, and it scales (almost) linearly with N_1 , at approximately \$83,000 for social surplus. Panel (b) shows the direct effect of increasing entry into Phase II. This effect is large and positive for social surplus, except for when $\bar{N}_2 = 1$, when it is definitionally zero. Once again, the low chance of Phase II success means that firms are not close substitutes in Phase II; thus, the benefit of an additional draw is not dampened by substitutability, and each additional draw outweighs the cost (even ignoring all effects on effort). The *net* direct effect is thus often positive for social surplus.

Panel (c) of Table 1.8 shows the incentive effect for Phase I. Phase I effort decreases with N_1 and increases with \bar{N}_2 .⁶⁸ Note that Phase I effort is socially excessive for these parameters, so decreases in this effort from more intense competition will tend to improve social surplus. Thus, the Phase I incentive effect on social surplus, which is large and positive, is increasing as N_1 increases but decreasing as \bar{N}_2 increases. Finally, the incentive effect for Phase II trades off savings in the cost of effort with higher cost draws. This effect is, unsurprisingly, estimated to be rather small. A firm factors in competition when determining its research effort only to the extent that it expects to influence its marginal surplus; because the probability that one's opponent succeeds is so low, this event does not influence incentives much.

1.8 The Effect of the Bargaining Parameter

What proportion η of the surplus should the firms receive? The bargaining parameter provides a second way to control the level of competition within a contest without resorting to finding more competitors—which may be impossible if the set of firms capable of conducting specialized research is small, or costly for unmodeled reasons. In this section, I take the parameter estimates from Section 1.5 as fixed and compute social surplus as a function of η and identify to what extent we can rectify the social inefficiency purely by changing the rewards the firms earn from the procurement phase.

Is it possible for the social planner to face a nonmonotonicity in η ? This question is related to the discussion of social inefficiencies provided in Section 1.6, along with the tradeoff between the holdup problem and the business-stealing and reimbursement effects. Increasing η ameliorates the holdup problem by giving the firm a greater claim to the surplus generated. Proposition 1.5 notes that increasing η is unambiguously beneficial for social surplus in Phase II, because there is no business-stealing in Phase II. However, larger values of η increase both the business-stealing and reimbursement effects in Phase I. Research is already overprovided in Phase I, and the cost of

⁶⁸For these parameters, the Phase I effort for $N_1 = \bar{N}_2$ is at a corner solution, which I set to be 0.99, and the Phase I incentive effect is thus zero for these parameters since Phase I effort does not adjust due to the boundary condition.




Figure 1-3: Social surplus as a function of η for two different levels of competition.

N_1	\bar{N}_2	η^*	Social Surplus			
			Baseline	η^*	% Increase to η^*	% Increase to Opt
1	1	0.56	0.012	0.013	8.3%	11.7%
2	1	0.50	0.028	0.029	3.5%	18.1%
3	2	0.57	0.111	0.122	10.3%	13.0%
4	2	0.59	0.246	0.271	10.0%	25.9%

Table 1.9: Optimal values of η from the perspective of social surplus. This table also reports social surplus (in millions of dollars) at the baseline value of $\eta = 0.73$ as well as at $\eta = \eta^*$. The second-to-last column reports the percent increase in social surplus from changing η to its optimal value. The final column repeats Table 1.6 and reports the percent increase in social surplus from changing to the optimum.

holdup in Phase II is not especially large for many parameters. Thus, increasing η could exacerbate the business-stealing and reimbursement effects to the point where they overshadow the gain from addressing the holdup problem.

Figure 1-3 shows social surplus as a function of η for two levels of competition, and Table 1.9 provides summary statistics for all four values of N_1 . First note that at the estimated value of $\eta = 0.73$, a marginal increase in η reduces social surplus: increasing η increases Phase II research (which is socially beneficial, as discussed in Section 1.6) but also Phase I research efforts (which is socially harmful). The latter cost is larger at the estimated value of the bargaining parameter. Social surplus is maximized at bargaining parameters between 0.5 and 0.6.

Table 1.9 suggests that social surplus would increase by 4–11% if the DOD switches to the socially optimal value of η . In addition, this change would of course reduce firm profits.⁶⁹ To

⁶⁹While it is possible to invoke the language of direct and incentive effects to discuss Figure 1-3, doing so is trivial in this setting. The direct effect of changing η on social surplus is identically zero, as it corresponds simply to changing the transfer between the DOD and the firms. The entire effect on social surplus comes through the incentive effect,

put this gain into perspective, we can compare the social surplus from optimizing η in Table 1.9 with the optimal values in Table 1.6. We see that the planner can achieve a surplus relatively close to the optimum when $(N_1, \bar{N}_2) = (1, 1)$ or $(N_1, \bar{N}_2) = (3, 2)$. When the business-stealing effect is stronger—i.e., when $(N_1, \bar{N}_2) = (2, 1)$ or $(N_1, \bar{N}_2) = (4, 2)$ —there are significant gains from being able to adjust Phase I effort separately without exacerbating the holdup effect by lowering η . Accordingly, adjusting η still keeps surplus relatively far from the frontier.

I briefly make one point related to procurement outcomes I have not discussed yet—in particular, profits. This section, like the other sections on design counterfactuals, has focused on total social surplus and *social* efficiency. However, in R&D procurement contests, there is a natural question of whether rewards are *Pareto* efficient from the perspective of the procurer and the firms. If the firm is not promised any part of the surplus (i.e., if $\eta = 0$), then the firm has no incentives to exert effort, and there will be very little surplus generated in the R&D contest.⁷⁰ On the other hand, if the firm is promised the entire surplus (i.e., if $\eta = 1$), then even though there may be a large amount of social surplus generated, the DOD will capture very little of it; indeed, the DOD will likely run a negative profit in this setting if it accounts for refunding the firms’ research efforts out of pocket as well.⁷¹ We can thus expect an inverted-U curve in the space of η on the horizontal axis and DOD profits on the vertical.⁷² Bargaining powers on the right side of this curve guarantee the procurer the same surplus while giving the firm a larger share of the social surplus, and they thus Pareto-dominate bargaining powers on the left side of the curve. The shape of this “Laffer” curve and where we stand on it is an empirical question studied in Appendix A.1. The conclusion is that the estimated value of $\eta = 0.73$ is on the Pareto-efficient side of this Laffer curve. Thus, while the value of η may be inefficient from the standpoint of total social surplus, there are no changes that are Pareto-improving for the DOD and the firms.

1.9 Decoupling Research from Delivery: Prizes and IP Sharing

In the current setup of the SBIR program, the incentives to conduct research come entirely from the possibility of a Phase III contract. In this sense, research and delivery are bundled. The DOD provides neither separate incentives for Phase I research nor the opportunity for firms in Phase II to develop research ideas generated by other Phase I competitors. In this section, I address the empirical relevance of these issues by running two related counterfactuals. First, I allow the DOD to modulate competition in Phase I separately from incentives in Phase II by giving a prize in Phase I for a successful innovation. Second, I consider the setting of full unbundling of tasks, in which this prize is to compensate the firm for the DOD buying the research plans developed in Phase I and sharing this intellectual property with other competitors.

so the derivative of the social surplus function with respect to η is the incentive effect.

⁷⁰As long as exerting no effort corresponds to failure with probability 1, there will indeed be zero surplus generated through the R&D contest.

⁷¹In cases where $N_2 = 1$, the DOD will earn identically zero from the procurement phase. When $N_2 > 1$, the DOD could still capture some of the surplus due to the competition embedded in the bargaining procedure.

⁷²In principle, there could be other nonmonotonicities in the intermediate region.

The effect of a Phase I prize is clear: it has no direct effect on social surplus but can stimulate research in the first stage at some cost to the DOD (and some extra social cost of R&D, through these stronger incentives). In this sense, prizes can counteract the holdup effect, albeit *only in Phase I*, whereas increasing η can remedy the holdup effect in both phases. The effect of IP sharing is somewhat more complicated. In the current setup of the contest, it happens that certain firms work during Phase II to develop ideas that have strictly worse value than their opponents. There is a direct social benefit for firms with higher draws of v to share these plans with competitors.⁷³ However, there are countervailing incentive effects again: while an otherwise weaker firm given access to a higher-quality idea may have more of an incentive to exert effort, the firm with the high-quality idea would shade its effort below its level in the equilibrium without information sharing. Furthermore, firms in Phase II naturally face more competition, as the DOD could share successful plans even with firms that are not successful in Phase I on their own. The net effect on social surplus is ambiguous. Moreover, the net effect on firms' profits in Phase II is also ambiguous, and this in turn affects Phase I research efforts; the incentives to generate ideas in Phase I are of course influenced by the fact that these ideas will be shared with competitors.⁷⁴

I omit the details for the model with only Phase I prizes: each firm is awarded K by the DOD for a success in Phase I, and the equations in Section 1.3.2 can be modified immediately. In Section 1.9.1, I present a model of IP sharing. I comment briefly on the issue of mandatory vs. optional IP sharing as well. Section 1.9.2 presents the value of IP sharing in the empirical setting of the DOD SBIR program.

1.9.1 A Model of IP Sharing

To develop a model of IP sharing, I consider the same timeline as in Section 1.3.2. To it, I add a prize at the end of Phase I that firms can (or, in the baseline case, *must*) accept in return for making the plans of their project public. In Phase I, each of N_1 firms exerts effort p_i at cost $\psi(p_i)$ to generate an idea of value v_i with probability p_i , with $v_i \sim F$ independently across i . All successful firms are given a prize K by the DOD and *must* make their plans public. (I eventually relax this assumption of mandatory IP sharing.) If no one generates a successful project, then the contest fails. Otherwise, the DOD shares the highest- v plan with \bar{N}_2 firms. For concreteness, and to maximize the incentives to exert effort in Phase I, I take the stance that the DOD shares the plans with successful firms first, and then the unsuccessful firms, breaking ties arbitrarily.⁷⁵ Thus,

⁷³Throughout this section, I will maintain the assumption that two firms working on the same idea generated in Phase I still get independent draws of delivery cost in Phase II. I comment on this assumption briefly at the end of this section.

⁷⁴Note that IP sharing corresponds to sharing intermediate breakthroughs that may make competitors stronger in future stages of the competition. Such "interim information sharing" has been considered in a number of stylized models, such as Bhattacharya, Glazer, and Sappington (1990), d'Aspremont, Bhattacharya, and Gérard-Varet (1998), and d'Aspremont, Bhattacharya, and Gérard-Varet (2000). In these models, firms each have a Poisson rate of developing a successful innovation, and information sharing is modeled as the firm with the higher rate increasing its competitor's rate while (potentially) reducing its own rate. To incentivize this IP sharing, firms are compensated by shares of the final surplus generated.

⁷⁵For instance, the DOD could invite the \bar{N}_2 firms that received the highest draws of v . If fewer than \bar{N}_2 firms were successful, it could pick the ones that are successful and then randomly choose among the firms that are not

if the contest enters Phase II, *exactly* \bar{N}_2 firms enter and they have identical values v . As before, they each draw costs $c_i \sim H(\cdot; t_i)$. In Phase II, the DOD chooses the firm with the lowest cost draw (as long as it is larger than v) surplus and pays it an amount equal to its implementation cost c_i plus a fraction η of the incremental surplus it generates above the next-best firm.

Equilibrium. As before, I search for a symmetric equilibrium. In Phase II, firms choose efforts to solve the problem

$$t^*(v) = \arg \max_t \left\{ \eta \int_{\underline{c}}^v \int_{\underline{c}}^c (c' - c) dH_{\bar{N}_2-1:\bar{N}_2-1}(c'; t^*(v)) dH(c; t) - t \right\}, \quad (1.13)$$

where $H_{\bar{N}_2-1:\bar{N}_2-1}(\cdot; t)$ is the minimum of \bar{N}_2-1 draws from $H(\cdot; t)$. Let $\pi_s(v)$ denote the maximized value of this problem. In Phase I, firms compute the profit conditional on success, from both the Phase I prize and the potential for profits from Phase II, as $\mathbb{E}[\pi_{\text{success}}(v; p^*, K)]$, where

$$\begin{aligned} \pi_{\text{success}}(v; p^*, K) \equiv & \underbrace{K}_{\text{Phase I prize}} + \underbrace{(1-p^*)^{N_1-1} \pi_s(v; \bar{N}_2)}_{\text{Phase II profit if sole success}} + \sum_{N_S=2}^{N_1} \underbrace{\binom{N_1-1}{N_S-1} (p^*)^{N_S-1} (1-p^*)^{N_1-N_S}}_{\text{probability of } N_S \text{ successes}} \\ & \times \underbrace{\Pr(\text{selected to Phase II if value is } v) \cdot \int_{v'} \pi_s(\max\{v, v'\}; \bar{N}_2) dF_{N_S-1:N_S-1}(v')}_{\text{Phase II profit}}, \end{aligned}$$

and where $F_{N_S-1:N_S-1}(\cdot)$ is the distribution of the highest of N_S-1 draws from F . Note that in this model with IP sharing, firms also have the potential of earning profits even when they do not generate a successful innovation, if they are selected to enter Phase II and use plans generated by a different firm. In this case, the firm expects to earn

$$\pi_{\text{failure}}(p^*) \equiv \underbrace{\sum_{N_S=2}^{\bar{N}_2-1} \binom{N_1-1}{N_S-1} (p^*)^{N_S-1} (1-p^*)^{N_1-N_S}}_{\text{probability of } N_S \text{ successes}} \cdot \underbrace{\frac{\bar{N}_2 - N_S}{N_1 - N_S}}_{\text{probability of selection}} \cdot \underbrace{\int_{v'} \pi_s(v'; \bar{N}_2) dF_{N_S:N_S}(v')}_{\text{Phase II profit}}.$$

Then, in equilibrium, p^* satisfies

$$p^* = \arg \max_p \{p \cdot \mathbb{E}[\pi_{\text{success}}(v; p^*, K)] + (1-p) \cdot \pi_{\text{failure}}(p^*) - \psi(p)\}. \quad (1.14)$$

An equilibrium of the R&D contest with IP sharing is a pair $(p^*, t^*(\cdot))$ such that $t^*(\cdot)$ satisfies (1.13) and p^* satisfies (1.14).

The setup assumes that firms are forced to share research breakthroughs at the end of Phase I, perhaps because the DOD can commit to not allow successful firms to enter Phase II if they choose not to share IP. This requirement may be difficult to enforce, and it may have long-term

successful. Since Phase I effort is decoupled from the realization of the value, and all successful contestants are awarded the Phase I prize, the actual tie-breaking rule is irrelevant.

repercussions by reducing the number of firms interested in participating in the SBIR program in the first place. Another option, therefore, is to make the prize $K(v)$ depend on v and contingent on sharing IP: all successful firms can still enter Phase II, subject to the limit \bar{N}_2 , but they forgo the payment $K(v)$ if they keep their breakthroughs private. I compute the incentive-compatible prize schedule $K(v)$ for comparison as well; details are in Appendix A.4.

1.9.2 The Costs of Prizes and IP Sharing

The total effect of prizes and IP sharing on outcomes of interest is an empirical quantity. In this section, I take the estimated parameters and compute the social surplus from the contest with Phase I prizes and with IP sharing described in Section 1.9.1. I focus on cases with $\bar{N}_2 > 1$ (because otherwise IP sharing is irrelevant), using parameter estimates corresponding to the associated value of N_1 .

Figure 1-4 shows social surplus as a function of the prize, in the cases without IP sharing and with mandatory IP sharing. The plot illustrates the case with $N_1 = 2$ and $\bar{N}_2 = 2$, using the estimates for $N_1 = 2$. Focusing on the setting without IP sharing, we see that offering a small prize (of about \$26,000 for a success) can increase social surplus, but it does so only by a modest 4%. This small prize increases the probability that an individual firm succeeds in Phase I by about 10 percentage points, and the added benefit from this cancels out the additional induced cost of effort. Note that in this case, we do not have a business-stealing effect in the first stage, so the equilibrium R&D effort in Phase I is an outcome of the countervailing reimbursement and holdup effects; without the business-stealing effect, the holdup effect dominates and makes the equilibrium R&D effort less than optimal (slightly). A small prize can remedy this situation and improve social surplus.

Consider next the plot of social surplus in the case of IP sharing. The first observation is that IP sharing by itself is not beneficial; in fact, it drops social surplus to essentially zero. This can be traced back to a severe drop in the probability of success in the first stage. First, firms are guaranteed to face a competitor in Phase II, which dissuades them from exerting effort in Phase I; however, because the probability that one's opponent is successful is small in Phase II, this is likely a small effect. Second, there is a free-riding effect with IP sharing in this mechanism: firms have a chance to enter Phase II if the DOD shares a successful rival's breakthrough. Thus, the return to effort is lower in Phase I. It is unsurprising, therefore, that research is underprovided in Phase I, and adding prizes can improve outcomes. Indeed, when coupled with Phase I prizes, IP sharing can be considerably more beneficial to social surplus than simply adding prizes. A prize of approximately \$62,000 increases social surplus to about \$93,000 from a base of \$54,000 with no prizes and no IP sharing. The benefit comes from added effort in the first stage together with an additional draw of an equally strong project in the second stage.

Table 1.10 reports summary statistics of this analysis for other values of N_1 and \bar{N}_2 .⁷⁶ The first few columns assume no IP sharing and list outcomes for no prizes (corresponding to the baseline

⁷⁶For each configuration of competition, I use the values estimated for the associated value of N_1 .




Figure 1-4: Social surplus as a function of the Phase I prize, without IP sharing and with mandatory IP sharing. I use the estimated parameters for $N_1 = 2$ and use $N_1 = \bar{N}_2 = 2$.

N_1	\bar{N}_2	No IP Sharing			Mandatory IP Sharing				
		$K = 0$	K_{SS}^*		$K = 0$	K_{SS}^*		K_{IC}	
		SS	K	SS	SS	K	SS	$\mathbb{E}[K(\cdot)]$	SS
2	2	0.054	0.026	0.055	0.004	0.062	0.093	0.131	-0.001
3	2	0.111	0.000	0.111	0.162	0.000	0.162	0.099	0.047
3	3	0.139	0.000	0.139	0.013	0.073	0.264	0.083	0.185
4	2	0.246	0.000	0.246	0.268	0.000	0.268	0.041	0.212
4	3	0.362	0.000	0.362	0.400	0.000	0.400	0.002	0.396
4	4	0.498	0.000	0.498	0.647	0.001	0.650	0.001	0.602

Table 1.10: Prizes K (and expected prizes $\mathbb{E}[K(\cdot)]$ in the case of the incentive-compatible schedule) and social surplus for various levels of prizes both without and with IP sharing. I report outcomes for (i) no prizes, (ii) the socially optimal prize, and, in the case of IP sharing, (iii) the minimum prize to make IP sharing incentive-compatible. All values are in millions of dollars.

contests studied through the rest of the chapter) as well as the socially optimal value of $K(\cdot)$. The takeaway from these columns is that the social planner usually gets no benefit from introducing prizes in this setting: research is already overprovided in the R&D contest (see the discussion in Section 1.6) for most of these parameters. Subsidizing Phase I research further is suboptimal.

The remainder of the table assumes IP sharing is mandatory, as in the model in Section 1.9.1. Note that prizes are beneficial in this setting (i) to ensure that there is *at least one* success (because as long as one person succeeds, entry into Phase II is fixed at \bar{N}_2) and (ii) to improve the quality of the best Phase II competitor. For these parameters, the social planner mostly prefers to not have a prize (or set a small one, as is the case with $N_1 = \bar{N}_2 = 4$); the exceptions are when p^* drops especially sharply without a prize, as in $N_1 = \bar{N}_2 = 2$ or 3. The next-to-last column also reports the expectation of the minimum prize schedule needed to make IP sharing incentive-compatible. This prize is always larger than the socially optimal prize. It decreases with both N_1 and \bar{N}_2 : accepting

the best draw from the opponents is more beneficial when there are more draws (N_1 is larger), and the expected benefit of holding out decreases when there is more competition (\bar{N}_2 is larger). Note that while the social planner would prefer to set a lower prize, the incentive-compatible prizes still do increase social surplus relative to the case of no prizes (which results in the benchmark with no IP sharing) for certain parameter values.

Finally, comparing the two parts of the table shows that introducing mandatory IP sharing without a prize need not be beneficial to social surplus, although it is for many of the parameters considered. For all parameters, IP sharing with a prize improves social surplus relative to the current design of the contest;⁷⁷ in some cases, this prize can be made incentive-compatible as well.

As a note, the computations in this counterfactual assumed that two firms working on the same project in Phase II (after IP is shared in Phase I) still receive *independent* cost draws. While such an assumption is justifiable in the baseline model in which firms independently generated their ideas, we may expect cost draws in a model with IP sharing to be correlated. Such correlation would obviously reduce the benefit of another draw. However, depending on the details of the correlation structure, it could intensify competition by introducing a setting like a “Bertrand trap,” in which the marginal benefit of exerting effort beyond the competitor’s level becomes especially large because it guarantees a lower cost draw. I leave a more detailed analysis of correlated cost structures to future work.

1.10 DOD Profits Under Alternate Contest Designs

How does the DOD evaluate the alternate contest designs proposed in Sections 1.7 through 1.9? In this section, I supplement the analysis from these previous sections by considering two potential objective functions for the DOD. The first one is a natural measure of DOD profits: this measure is the value of the project the DOD acquires in Phase III, less the delivery cost, less total research costs in Phases I and II. In settings where the DOD pays out prizes, these prizes are subtracted from the profits as well. To fix ideas, note that if $\bar{N}_2 = 1$, then this measure is $(1 - \eta) \cdot (v - c)^+$ less total research costs. Furthermore, DOD profits in this way are defined so that these profits plus firm profits equals social surplus. This measure is my preferred measure of profits. I include a second measure of profits that I call “Phase III DOD Profits.” This measure is simply the value of the product less delivery cost, and it ignores research costs and prizes. It provides an interesting point of comparison for institutional reasons: the DOD is institutionally constrained to spend approximately a fixed proportion of its R&D budget on Phase I and Phase II research, and thus the surplus generated in delivery (Phase III) may be independently of interest.⁷⁸

Table 1.11 lists outcomes—social surplus, DOD profits, and the Phase III DOD profits—for various contest designs, using parameters for $N_1 = 4$. In this table, I fix N_1 and consider η , \bar{N}_2 ,

⁷⁷Note that this is not a forgone conclusion. It is possible in this model for IP sharing to be socially suboptimal relative to the current design of the contest—for any level of the prize—purely because it induces adverse incentive effects in Phase II.

⁷⁸One could in principle consider other measures, such as a weighted average of DOD profits and firm profits, or the return on Phase I and Phase II investment.

Design	Social Surplus (\$M)	DOD Profits (\$M)	Phase III DOD Profits (\$M)
Baseline	0.246	-0.238	0.191
IP Sharing	0.268	-0.245	0.203
$\eta = \eta^*$	0.271	-0.064	0.241
$N_2 = \bar{N}_2^*$	0.498	-0.361	0.378
Social Optimum	0.521	-0.797	0.077

Table 1.11: Social surplus, DOD profits, and DOD profits in Phase III (i.e., ignoring research and prizes) for various contest designs. These numbers are computed for parameters with $N_1 = 4$ with $(N_1, \bar{N}_2) = (4, 2)$. “IP Sharing” refers to mandatory IP sharing with socially optimal prizes. “ $\eta = \eta^*$ ” refers to the socially optimal value of η with $(N_1, \bar{N}_2) = (4, 2)$. “ $N_2 = \bar{N}_2^*$ ” refers to $(N_1, \bar{N}_2) = (4, 4)$ and η at the estimated value. The social optimum is implemented by setting $\eta = 1$, setting $\bar{N}_2 = 4$, and imposing fees for entry into Phase II to set the equilibrium Phase I effort to the social optimum.

and whether there are prizes and IP sharing to be design choices. I report outcomes for the baseline design, in which $\bar{N}_2 = 2$ and $\eta = 0.73$, as well as the socially optimum design, in which $\bar{N}_2 = 4$, $\eta = 1$, and fees are imposed by the DOD on Phase II competitors to induce the socially optimal effort in Phase I. I also separately consider the three design counterfactuals studied in Sections 1.7 through 1.9: changing \bar{N}_2 , changing η , and allowing for prizes while mandating IP sharing.

The first column summarizes the results regarding social surplus of the previous sections: information sharing and choosing the optimal η increases social surplus slightly, but the largest gains are from allowing more Phase II competitors. Increasing the cap \bar{N}_2 increases surplus by 102%, compared to the socially optimal increase of 112%.

The second column shows the effect of these design changes on DOD profits. Note that in the baseline case, the DOD runs a loss of approximately \$238,000 per contest. About \$430,000 of this is due to reimbursement of research expenditures, so the Phase III profit is about \$191,000 per contest. Note that the DOD runs losses in the program because, like the social planner, it internalizes the full costs of research. However, unlike the social planner, it only internalizes (about) one-quarter of the surplus generated in the delivery phase.

At the social optimum, firms have a larger incentive to exert effort in Phase II because $\eta = 1$, so the DOD pays a larger amount to reimburse these research efforts. More importantly, it does not recover much of the surplus generated in Phase III. If only one firm succeeds in Phase III, the DOD earns nothing from the delivery process because it pays the firm its value. In the considerably less likely scenario that multiple firms succeed, the DOD recovers the inframarginal surplus generated by the firms, but the winning firm captures the entire incremental surplus it generates. Accordingly, a DOD that only cares about Phase III profits would not want to move from the baseline design to the socially optimal design, as these profits decrease by about 60%. If the DOD fully internalizes the costs of research in its objective, the result is even more stark: losses increase more than three-fold, so the DOD would not have an incentive to move to this implementation of the socially optimal design.

In general, socially beneficial design changes need not be beneficial to the DOD. Changing $\bar{N}_2 = 2$ to the socially optimal value of $\bar{N}_2 = 4$ increases the DOD's losses by about 40% when accounting for research expenditures. IP sharing reduces DOD profits by a small amount: the improvement in values when entering Phase II is already low due to the low variance of the value distribution. Moreover, the DOD only captures a small fraction of the improvement, and it is insufficient to counteract the increase in research costs.⁷⁹ The Phase III profits of the DOD, however, are often aligned with the social planner's objective. Increase \bar{N}_2 almost doubles the DOD's profits from Phase III, due to a combination of a significantly larger probability that someone succeeds in Phase II and (less importantly) to competition between multiple firms that succeed. IP sharing increases Phase III profits slightly as well: \bar{N}_2 firms enter Phase III in more instances (i.e., as long as at least one firm succeeds in Phase I), the competitors have slightly higher values, and the DOD has a slightly larger effective bargaining parameter due to multiple competitors.

The one design change that is beneficial to the DOD under either objective is reducing η to the socially optimal one: this cuts the DOD's losses by about 75%. Recall that the social benefit of reducing η comes from reducing business-stealing and the reimbursement effect, at the cost of exacerbating the holdup problem. The DOD benefits from reducing business-stealing and the reimbursement effect more than the social planner does because its objective places more weight (relatively) on saving effort costs. Moreover, reducing η has a direct benefit of allowing the DOD to capture a larger portion of the surplus. In fact, the Phase III profits of the DOD increase by about 20%: even though less surplus is generated in Phase III, the fact that the DOD captures a larger portion of it increases its Phase III profits. However, while reducing η is beneficial to both the social planner and the DOD, it does reduce the firms' profits.

The results of this section can be summarized by noting that simple design changes that are beneficial from the perspective of social surplus are almost always harmful from the perspective of the DOD, because the DOD refunds research costs while only capturing a somewhat small portion of the generated surplus. Of course, this is not to say that the baseline design of the contest is close to optimal for the DOD. If given the option of choosing the parameters within each class of design changes, the DOD would select starkly different ones. In particular, it would set $\bar{N}_2 = 1$, choose a much lower level of η , and sometimes choose *not* to mandate IP sharing. Details are provided in Appendix A.1.

1.11 Conclusion

In this chapter, I proposed a model of R&D contests that are incentivized by procurement contracts. I showed how to identify the distribution of values of the projects to the DOD, the delivery costs to the firms, and the costs of research from data on research efforts and the final delivery contract. Identification rests on two fairly weak assumptions: that firms that are working on higher-value

⁷⁹As shown in Appendix A.1.3, there are some combinations of N_1 and \bar{N}_2 where it is possible to improve both social surplus and DOD profits by a combination of IP sharing and prizes. In particular, this is true with N_1 and \bar{N}_2 are both 3 or both 4. However, these correspond to "nonstandard" parameters that do not obey the 40% rule of the DOD.

projects also spend more on R&D, and that the DOD does not procure projects whose delivery costs exceed value. Adding information about the optimality of research efforts allows me to identify the bargaining parameter in the final delivery process as well.

By applying this methodology to the case of the DOD SBIR program, I quantified the relative contribution of value and cost variation to the final outcomes. The SBIR program focuses on projects that are moderately valuable to the DOD, amounting to about \$11-\$15 million for an average project. Within a contest, there is a rather small amount of variation across contestants in the value of the projects that they bring to the table. Much of the variation in outcomes is instead attributed to variation in delivery costs, which are determined later in the R&D process. Moreover, I provided evidence that firms are able to capture a large portion of the surplus generated through this program.

These estimates allowed me to quantify the social inefficiency in the R&D contest. I show that, due to the holdup effect, late-stage research is underprovided in the R&D contests; however, because the firms still earn a reasonably large portion of the surplus, social surplus can only be improved by 5–10% by changing late-stage research efforts, which provides an estimate of the cost of holdup. In the early stage of the contest, there is also a business-stealing effect and a reimbursement effect that firms do not internalize future R&D efforts, and the net impact is that R&D is overprovided in equilibrium. Social surplus can be improved by over 20% by modulating this effort.

I then analyzed three design counterfactuals to modulate competition and effort. First, I studied the effect of adding competitors. I found that the social planner prefers to encourage a large amount of entry into the contests, while the DOD—which captures a somewhat small part of the surplus—prefers to restrict entry. The direct benefit of having another draw in Phase II—the phase in which most of the failure occurs—is especially large for the social planner, and it comes at a somewhat small cost. The indirect incentive effects, which arise from the equilibrium readjustment of effort, are large and *beneficial* (due to savings in research costs) for early-stage competition but moderately small for late-stage competition. Second, I considered the impact of changing the surplus given to the firms. While it is not possible to improve both the firms' profits and the DOD's profits, reducing the bargaining power of the firms can increase overall social surplus by a small amount. Finally, I considered the impact of decoupling research and development by incentivizing firms to share interim research breakthroughs with each other, and I showed that this mechanism can improve social surplus but is not guaranteed to improve DOD profits. These results suggest that the DOD and the social planner would prefer significantly different design changes to the contest.

Considering both the theoretical interest in them as well as their empirical relevance, contests have been understudied in the structural literature. Indeed, many settings—including ones that are not explicitly structured as contests—lend themselves to be conceptualized as such. For instance, the auctions with entry literature has studied the impact of value discovery and bid preparation costs on auction participation, but it does so almost exclusively in a setting where these entry costs do not directly impact values and costs. One may naturally wonder whether in large construction projects, for instance, there is a direct relationship between constructing a more careful plan and drawing a lower cost draw; firms that are ex-ante better candidates to complete the project may have an

incentive to develop better plans. A very different example involves FDA trials for pharmaceuticals, which have a multistage structure as well. The size and scope of later phases of a trial can depend on how promising a drug seems in earlier phases, and the decision to take the drug to the market occurs only if the expected benefit exceeds any cost of commercialization. One main advantage of the methodology developed in this chapter is that it is likely flexible enough to be adapted to these diverse situations—and doing so is an especially promising avenue for the future.

Chapter 2

Regulating Bidder Participation in Auctions

Regulating bidder participation in auctions can potentially increase efficiency compared to standard auction formats with free entry. We show that the relative performance of two such mechanisms, a standard first-price auction with free entry and an entry rights auction, depends non-monotonically on the precision of information that bidders have about their costs prior to deciding whether to participate in a mechanism. As an empirical application, we estimate parameters from first-price auctions with free entry for bridge-building contracts in Oklahoma and Texas and predict that an entry rights auction increases efficiency and reduces procurement costs significantly.

2.1 Introduction

Given the importance of public procurement for many economies around the world (e.g., 12% of GDP is spent on public procurement in OECD countries (OECD (2011))), there is natural interest in improving the efficiency of the procurement process and reducing procurement costs. Under the assumptions, which we maintain, of symmetric and independent private costs, quite standard auction formats, such as first-price (low-bid) auctions with suitably chosen reserve prices will be optimal when the number of bidders is fixed. Assuming a fixed set of bidders will not be appropriate, however, when suppliers have to spend significant resources in “due diligence” to understand how much it will cost them to complete the project, and in most of the procurement settings that have been studied, these types of entry costs appear to be significant (e.g., Li and Zheng (2009) and Krasnokutskaya and Seim (2011)). In this article we examine the performance of a standard first-price auction with costly but unregulated, or “free”, entry relative to an alternative mechanism

This chapter is joint with James Roberts and Andrew Sweeting, and it was published in the Winter 2014 issue of the *RAND Journal of Economics*. It is reproduced here with permissions. We are very grateful to Dakshina De Silva, Timothy Dunne, Anuraddha Kankanamge, and Georgia Kosmopoulou for sharing their data with us. Ali Hortaçsu, Tim Hubbard, Jon Levin, Rob Porter, Paulo Somaini, Steve Tadelis, John Turner, John Asker (the co-editor), seminar participants, and three referees provided feedback that improved the article. Bhattacharya would like to thank the Davies Fellowship from the Duke University Department of Economics for financial support. Roberts and Sweeting acknowledge National Science Foundation support (grant SES-1260876) for this research. Any errors are our own.

where the procurer regulates entry by using an initial auction to allocate a fixed number of rights to participate in a second auction for the contract (an “entry rights auction”).

When entry is endogenous, a standard first-price auction is usually thought of as a two-stage game. In the first stage, a set of potential suppliers (for example, a set of local contractors who purchase or are issued with specifications for the contract) simultaneously and non-cooperatively decide whether to pay the cost of entering the auction based on any private information (signals) that they have about their costs of completing the project. In the second stage, the entrants find out their true costs of completing the project and submit bids. The contract is allocated to the firm with the lowest bid, as long as this is less than any reserve price, at a price equal to its bid. We will call this standard model a “first-price auction with free entry” (FPAFE). In the equilibrium of an FPAFE, the marginal entrant in the first stage will be indifferent between entering the auction, which involves paying the entry cost but allows the supplier to potentially win the contract, and staying out.

A natural question is whether these entry decisions are efficient, in the sense of being consistent with the minimization of the expected sum of the winning bidder’s cost of completing the project and total entry costs (social costs). As shown theoretically by Levin and Smith (1994) and Gentry and Li (2012) under different assumptions about the private information that potential suppliers have about their costs prior to entry, entry decisions will be efficient in the sense that a planner who had to choose whether a potential supplier should enter based only on that firm’s private information would choose the same entry rule that the potential supplier would choose in the symmetric equilibrium.¹ However, inefficiencies in entry can still arise from the fact that the entry decisions of different potential bidders are not coordinated, so that there can be outcomes where too many or too few firms enter. For example, when potential suppliers have no private information about their costs when taking entry decisions, as assumed by Levin and Smith (1994) (LS), the symmetric equilibrium involves potential suppliers mixing over entry. As a result, the number of entrants will be stochastic, and Milgrom (2004) shows that efficiency would be increased if the procurer instead randomly selected a fixed number of potential suppliers to enter the auction.

The assumption that potential suppliers have no private information about their costs before entering seems implausible for many procurement settings, as firms are likely to know at least something about their free capacities and their costs of securing some of the necessary inputs. When their private signals are informative about their costs (Gentry and Li (2014) and Roberts and Sweeting (2013b)), free entry will tend to lead to the most efficient firms entering and an arrangement where the auctioneer randomly selected entrants without using this private information would have a natural disadvantage. However, there could still be scope for the auctioneer to regulate

¹The typical Mankiw and Whinston (1986a) result that free entry results in excess entry in a homogeneous goods market does not hold because an entrant only takes business from other firms when it is socially optimal for it to do so. The Levin and Smith and Gentry and Li results are conditional on the auctioneer setting a reserve price that is equal to the cost of completing the project outside of the auction if there is no winner. This is the auctioneer-optimal reserve in the Levin and Smith model, but not the Gentry and Li model, although it is the natural “non-strategic” reserve price in that setting as well. As pointed out by a referee, one interpretation of these results is that the procurer cannot improve efficiency by choosing a reserve strategically.

the number of entrants by using an entry rights auction mechanism (ERA) where, based on their private information, the potential suppliers first participate in an auction for a fixed number of rights to perform due diligence and enter the auction for the contract. As long as first-stage bids are monotonic in the bidders' private signals, the ERA will also tend to select the most efficient firms, but the ERA will have the advantage that the number of entrants will be controlled.

We show, using an example, that whether an FPAFE or an ERA is more efficient can depend in a non-monotonic way on the precision of the private information that suppliers have about their costs. When signals are imprecise, the ERA is more efficient, as one would expect given the informational assumptions made by LS and Milgrom in deriving the result mentioned above. When signals are very informative, in which case the model tends towards the Samuelson (1985) (S) model where potential entrants know their costs exactly, the ERA is also more efficient because it is optimal to have only the supplier with the best (lowest cost) signal paying the entry cost; this is the outcome that the ERA ensures will happen, whereas no supplier or more than one supplier may enter in an FPAFE. However, when potential bidders' signals are moderately informative about their costs, an FPAFE can be more efficient. The reason is that with free entry the number of entering suppliers will depend on the private information available to the suppliers but not the auctioneer, and this can be desirable. For example, although on average it might be desirable to have two entrants, three entrants might be preferred if they all have signals that indicate that their costs are likely to be low.

Given this ambiguity in which mechanism is more efficient, we estimate our model using data on a sample of contracts for highway projects involving work on bridges let by the Departments of Transportation (DoTs) in Oklahoma and Texas using FPAFEs in order to understand which mechanism might perform best in a real-world context. To do so, we develop methodologies for both solving and estimating a parametric model of FPAFEs where potential bidders have imperfect information about their costs when deciding whether to enter. We solve the model using the Mathematical Programming with Equilibrium Constraints approach proposed by Su and Judd (2012).² We estimate our model using a simulated method of moments estimator where the moments are computed using importance sampling, following Akerberg (2009). This approach enables us to allow for observed and unobserved heterogeneity in the parameters across auctions, unlike the one previous attempt to estimate this type of model of which we are aware (Marmer, Shneyerov, and Xu (2011)).³

We estimate that entry is moderately selective for these contracts. However, given the other parameters of our model, our model predicts that overall efficiency would have been increased by using the ERA format. For example, for the representative (average) auction in our data, we

²Hubbard and Paarsch (2009) use this approach to solve first-price auction models under the S model assumption that potential bidders know their values when deciding whether to enter. We extend their solution method to allow for noisy signals about values (in our case, costs), creating imperfect selection of low cost bidders into the auction.

³Some earlier work, such as Krasnokutskaya and Seim (2011), estimates models of endogenous entry into first-price auctions under the informational assumptions of LS. This implies that there is no selection in the entry process, which simplifies estimation because the distribution of values for bidders will be the same as the distribution of values in the population of potential suppliers.

predict that the sum of the expected cost of the winning bidder and total entry costs would be 2.4% lower in the ERA than the FPAFE and that the total cost to the DoT of letting the contract (procurement cost) would fall by the same proportion. This fall in DoT procurement costs is much greater than what we predict would be achieved by using a more standard design adjustment to an FPAFE, such as adding an optimal reserve price, which for our representative case would only reduce costs by 0.03%, or by adding additional potential entrants.

We are not aware of any previous attempts in the literature to compare ERAs and auctions with free entry using real-world parameters. A related theoretical literature has considered so-called “indicative bidding” schemes where potential bidders submit non-binding indications of what they are willing to bid, which the seller may use to select a subset of firms to participate in an auction.⁴ In contrast, in an ERA, first-round bids are binding and can result in payments by firms that do not win the contract. Indicative schemes are often used by investment banks and as part of the complicated processes by which defense equipment is purchased (Quint and Hendricks (2013), Foley (2003), Welch and Fremond (1998)). Ye (2007) argues that, generically, indicative schemes will not induce strictly monotonic first-stage bids in equilibrium and so will not tend to guarantee the selection of the bidders that are most likely to have the highest values or the lowest costs in the second stage. However, he shows that certain types of ERA schemes, like the one that we consider, will induce strictly monotonic first-stage bidding and so guarantee an efficient selection of entrants. Recently Quint and Hendricks (2013) have shown that, when indicative bids are limited to a discrete set and either entry costs or the number of bidders is large, the unique indicative bidding equilibrium can involve weakly monotonic bid functions, and that, in computed examples, indicative schemes can outperform standard auctions with free entry. By guaranteeing fully efficient selection of entrants, ERAs should, in theory, be more efficient than indicative schemes when the number of selected entrants is the same. At the end of the article, we discuss several possible reasons why ERAs may be less common than indicative schemes or FPAFEs in practice.

The article is also related to our earlier work on selective entry auction models (Roberts and Sweeting (2013a), Roberts and Sweeting (2013b)).⁵ In those articles we consider selective entry into second-price auctions, where the simpler form of equilibrium bidding eases computation and estimation relative to the first-price auctions considered in this article. FPAFEs are the auction format that is most widely used in practice.⁶ In Roberts and Sweeting (2013b), we compare the

⁴There are other designs that also have the flavor of using an auction to select a small number of firms into another round. For example, in the Anglo-Dutch auction described by Klemperer (2002), bidders first participate in an ascending auction that selects a smaller set of bidders who compete in a first-price sealed bid auction. This design has been discussed primarily in settings where bidders are asymmetric (for example, where there are well-known strong incumbents) and where valuations are likely to have some common component. In our model, suppliers are symmetric and we assume IPV.

⁵Other work on selective entry in auctions includes Gentry and Li (2014), who focused on identification, and Marmer, Shneyerov, and Xu (2013), who consider non-parametrically testing different models using first-price auction data in the absence of unobserved auction heterogeneity.

⁶As we assume potential suppliers are symmetric, it is possible to compute expected efficiency and procurement costs for first-price auctions by exploiting the fact that these outcomes should be the same in the second-price format. However, to use the information contained in the distributions of submitted bids and winning bids in estimation it is necessary to be able to solve first-price auctions with selective entry.

performance of a free-entry auction with a sequential bidding process which is more efficient, and can improve the seller's revenues for a wide range of parameters, even though it allows early bidders to deter entry. Here we consider a different type of design change, which retains the simultaneous bidding feature of most auction formats. This may be attractive, as sequential procedures have the potential weakness that the auctioneer may favor certain bidders when determining the order.

Four comments are in order about both our model and the nature of the results. First, when considering the ERA, we model the first-stage auction for entry rights as an all-pay auction. Following Ye (2007), the advantage from a modeling perspective of using the all-pay format is that equilibrium first-stage bids are guaranteed to be strictly monotonic in the signal that potential bidders have about their costs when entry costs are subsidized, which we also assume. This guarantees that the suppliers that are most likely to have the lowest costs are selected. When the more common discriminatory or uniform price formats are used to sell entry rights, this property may not hold because the expected benefit of being the marginal entrant allowed into the final auction, which determines the bid in these formats, is not guaranteed to be monotonic in the signal. In fact, it is straightforward to find examples where the relationship is not monotonic for the distributions that we consider. However, as Ye shows, when the relationship is monotonic, the expected costs and efficiency under all three first-stage schemes (all-pay, discriminatory and uniform) will be identical.

Second, we are comparing our results to a specific form of ERA, where the seller announces an exact value for the number of bidders who will be able to enter the second-stage auction in advance, rather than considering the optimal mechanism or even the optimal entry rights auction, where the number of allowed entrants into the second stage might well depend on the bids that are announced in the first stage. We do this partly because it is not straightforward to characterize the procurer-optimal form of entry rights auction in our model, but also because even if it were known, the optimal mechanism would likely have such a complicated form that a public agency would likely not want to use it in practice.⁷ In any case, because our bottom-line conclusion is that there may be significant cost and efficiency advantages to using ERAs, this conclusion would only be strengthened if we considered the optimal form of the ERA instead. In discussing our results we do, however, compare the efficiency of our ERA mechanism with the efficiency of two hypothetical mechanisms where we allow a planner to have access to the private information of the bidders and to use these signals to choose how many firms enter. Although these hypothetical alternatives are naturally more efficient, we show that the ERA actually performs almost as well as they do for the parameters that we estimate.

Third, our results are numerical. When entry is partially selective, it is very difficult to solve models analytically outside of some special cases.⁸ Moreover, as we show, efficiency and cost comparisons will often depend on the distributions and parameters that are assumed, so it is

⁷Lu and Ye (2013) characterize the optimal entry rights auction format in a setting where there is heterogeneity in potential bidders' entry costs, but they have no private information on their values prior to entering. Cremer, Spiegel, and Zheng (2009) characterize the optimal mechanism in this case, which involves a complicated sequential search process.

⁸For example, Ye (2007) considers a case where bidders are uncertain about a component of costs that can only take on discrete high or low values.

important to try to be as realistic about these elements as possible. For this reason, we focus our results on parameters that are estimated from real-world data and emphasize the magnitude as well as qualitative direction of our results.

Finally, although we allow for partially selective entry and a model of auction heterogeneity, we maintain the independent private costs/values assumption that has been typically made in both empirical studies of public procurement (e.g., Krasnokutskaya and Seim (2011)) and theoretical work on models of two-stage entry rights or indicative bidding procedures (e.g., Quint and Hendricks (2013)). Understanding the effects of trying to regulate entry where costs have both common value and private value components is left to future work. Similarly we assume that firms would act competitively in both FPAFEs and ERAs, whereas in some settings it might be important to consider which mechanism would be more susceptible to collusion.

The article proceeds as follows. Section 2.2 presents the model of imperfectly selective entry for FPAFEs and ERAs, and explains how we solve these models. Section 2.3 describes the data. Section 2.4 explains our estimation method. Section 2.5 discusses the estimates of the model’s parameters and compares the performance of the FPAFE and the ERA. Section 2.6 concludes. The Appendices contain some additional details on our numerical routines and present some Monte Carlo evidence on the performance of our estimator.

2.2 A Model of Endogenous Bidder Participation in Auctions

We introduce the general structure of costs and information in our model of procurement and then describe the FPAFE and the ERA. We also present a numerical example illustrating how an efficiency comparison of the two mechanisms will depend on the degree of selection in the entry process.

A procurement agency, which we will call the “procurer” in what follows, wishes to select one of N risk-neutral suppliers to complete a project. In our empirical setting, a project might involve repairing or constructing a bridge, and we will allow for cross-auction variation in both N and the other parameters. The agency may require that it pay no more than a reserve price r to complete the project, and if no firm submits a bid below this amount, the agency incurs a cost c_0 of procuring the work outside of this specific auction (e.g., this could be the expected cost from re-running the auction at a later date, or negotiating one-on-one with a particular contractor). Supplier i can complete the project at cost C_i distributed $F_C(\cdot)$ with compact support $[\underline{c}, \bar{c}]$ that admits a continuous density $f_C(\cdot)$. We assume that suppliers are ex-ante symmetric and have independent private costs. To participate in the project-allocation auction of any mechanism,⁹ a supplier must pay an entry cost K . Entrants learn their true costs of completing the project. One interpretation of K is that it includes the cost of research involved in finding out the true cost, but it will also include any costs associated with preparing a bid. An important assumption that we maintain throughout is that a firm cannot participate in the project-allocation stage of

⁹This will be the auction in the FPAFE and the second-stage auction in an ERA.

any mechanism without paying K .¹⁰ Prior to deciding whether to enter, each supplier observes a signal S_i that is correlated with its true cost C_i and that is not correlated with any other supplier's cost. Signals are affiliated with costs in the standard sense: the cost distribution conditional on a signal s first-order stochastically dominates the cost distribution conditional on a signal $\tilde{s} < s$. We assume that the marginal distribution of signals admits a continuous density on a compact interval $[s_{\min}, s_{\max}]$.

The LS and S models are limiting cases of this model. If $S_i = C_i$, then signals are perfectly informative of costs, and the setup reduces to the S model. When S_i is independent of C_i , signals contain no information about costs, and the setup reduces to the LS model. For many empirical settings, it seems plausible that buyers will have some, but imperfect, information about their costs prior to conducting costly research, consistent with S_i being positively, but not perfectly, correlated with C_i .

2.2.1 Selective Free Entry into First-Price Auctions

The first model of procurement is the FPAFE, which we view as describing the procedure used by the DoTs in our data and most other government agencies. In the first stage, potential suppliers observe their imperfectly informative cost signals. Based on these private signals, they take independent, simultaneous entry decisions and pay K if and only if they enter. In the second stage, the set of entrants compete in a first-price (low-bid) auction with reserve price r . We assume that firms that do not pay K cannot participate in the auction and that agents bid without knowing how many of their competitors actually entered the auction, as is done in Li and Zheng (2009). If an entrant learns that his true cost exceeds the reserve price, he will not bid.

We solve for entry decisions and bid functions that define the unique symmetric Bayesian Nash equilibrium with monotone bidding behavior.¹¹ Potential bidders enter using a cutoff strategy; that is, a supplier enters if and only if he observes a signal $s_i < s'^*$ for some critical value of s'^* (Gentry and Li (2012)). Define $H(c)$ to be the probability (not conditional on S_i) that a given supplier either (i) enters the auction and has a cost no smaller than c or (ii) does not enter the auction; thus, $H(\cdot)$ depends on the value of the signal cutoff used for the entry decision. The equilibrium bid functions $\beta^*(\cdot)$ are given by the solutions to the optimization problem

$$\beta^*(c) \equiv \arg \max_b (b - c) \left[H \left(\beta^{*-1}(b) \right) \right]^{N-1}. \quad (2.1)$$

¹⁰As pointed out by a referee, if K is interpreted as containing an information acquisition cost then there may be incentives for some suppliers to bid without paying the information acquisition cost. To get around this objection, one can either assume that K is simply a cost that must be paid in order to be responsive in the auction (in which case we could allow non-entering suppliers to find out their costs as well once they have decided not to pay K) or we can assume that there are some penalties outside of our model that deter this type of uninformed bidding. For example, there may be penalties that arise from defaulting on the contract or incurring financial losses, or it may be harder for an uninformed supplier to find suitable sub-contractors.

¹¹Gentry and Li (2012) establish uniqueness of the symmetric equilibrium in a class of affiliated signal models. As usual in entry games, asymmetric equilibria may exist.

The first order condition associated with this optimization problem gives the differential equation

$$1 + \beta^{*-1'}(b) \left(b - \beta^{*-1}(b) \right) (N - 1) \left[\frac{H' \left(\beta^{*-1'}(b) \right)}{H \left(\beta^{*-1'}(b) \right)} \right] = 0, \quad (2.2)$$

with the upper boundary condition

$$\beta^*(r) = r. \quad (2.3)$$

The equilibrium critical cutoff values s^* are determined by the indifference condition that any potential supplier who receives a signal of s^* must be indifferent between entering the auction or not paying the entry cost at all. This zero-profit condition is thus written

$$\int_{\underline{b}}^r \left(b - \beta^{*-1}(b) \right) f_C \left(\beta^{*-1}(b) | s^* \right) \left[H \left(\beta^{*-1}(b) \right) \right]^{N-1} db = K, \quad (2.4)$$

where $\underline{b} \equiv \beta(\underline{c})$, and $f_C(\cdot|s)$ is the conditional density of a supplier's costs, computed using Bayes' Rule, given he receives a signal s .

To solve for the bid functions, we use the Mathematical Programming with Equilibrium Constraints (MPEC) approach (Su and Judd (2012)). In a manner similar to that outlined by Hubbard and Paarsch (2009), who use MPEC to solve an FPAFE model where bidders know their values when they decide whether to enter, we express the inverse bid function as a linear combination of the first P Chebyshev polynomials (we use $P = 25$), scaled to the interval $[\underline{b}, r]$. The choice variables in our programming problem are, therefore, P Chebyshev coefficients, the signal cutoff, and the value of the low bid (\underline{b}). We pick a fine grid $\{x_j\}_{j=1}^J$ with $J = 500$ points on the interval $[\underline{b}, r]$. Then we solve for the bid functions (more precisely, the Chebyshev coefficients) and the signal cutoff using

$$\arg \min_{\{\underline{b}, \beta^{*-1}, s^*\}} \sum_{j=1}^J g \left(\beta^{*-1}(x_j) \right)^2 \text{ s.t. (2.3) and (2.4),} \quad (2.5)$$

where $g \left(\beta^{*-1}(b) \right)$ is defined to be the left-hand side of equation (2.2). This nonlinear programming problem is solved using the SNOPT solver called from the AMPL programming language. We give more details about the numerical methods in Appendix B.1.¹²

2.2.2 Entry Rights Auctions

When entry is regulated and suppliers have some information about what their costs are likely to be, a straightforward way to choose the firms that will be allowed to bid for the contract is to use an initial auction to allocate a fixed number of rights to participate in a second-stage auction for the contract (Ye (2007)). Before the first stage, the procurer announces the number of suppliers that will be selected to participate in a (second-stage) auction for the contract and the reserve price

¹²We assume that suppliers are ex-ante symmetric, but this procedure can also be used to solve for equilibria when suppliers belong to certain discrete types. However, a procedure for finding all of the equilibria in an asymmetric FPAFE model remains a topic of ongoing research.

in that auction. In the first stage of our ERA, each supplier i receives its signal s_i and submits and pays a non-negative bid, given by a function $\gamma^*(s_i)$, for the right to participate in the second-stage auction. The highest n bidders are selected to participate in the second stage.¹³ In the second stage, the procurer pays each of the selected entrants K to cover their entry costs, and the selected firms incur these costs and find out their true costs of completing the contract. The procurer also publicly announces the value of the $(n + 1)^{\text{st}}$ highest first-stage bid.¹⁴ The second-stage auction uses a first-price format with the pre-announced reserve price, r . However, as discussed by Ye (2007) any standard auction format used in the second stage (i.e., one where the firm with the lowest bid wins the contract and a firm with the highest possible cost pays nothing), will generate the same expected payoffs for all parties, and therefore the same incentives in the first stage, as long as entrants have symmetric beliefs about their rivals and the first stage has selected the firms with the lowest costs. We assume a first-price format because it reveals an interesting difference in bidding behavior between the FPAFE and the second stage of an ERA with at least two entrants.¹⁵

The first-stage auction is therefore a *binding all-pay auction with an entry subsidy*. Some explanation for why we model the ERA in this way is required. As shown by Ye (2007) in his Proposition 5, in an all-pay auction, there will be a unique, symmetric pure strategy equilibrium where the first-stage bid function, $\gamma^*(s)$, is a strictly decreasing function of the cost signal, so that a firm with a lower cost signal will bid more, when the entry subsidy is sufficiently large.¹⁶ This monotonicity property is important because it guarantees that the selection of entrants into the second stage will be efficient. An entry subsidy equal to the due diligence cost K is sufficient for this property to hold, and, by Ye's Proposition 6, the level of the subsidy does not affect either the procurer's or the bidders' total payoffs from the mechanism as long as equilibrium first-stage bids are strictly decreasing, as more generous subsidies are exactly offset by higher first-stage bids. A corollary of this argument is that none of our efficiency or procurement cost comparisons would change if we considered slightly lower subsidies that also gave rise to monotonic first-stage bid functions.

In contrast, a strictly decreasing equilibrium first-stage bid function might not exist if we considered a first-round format that operated as either a uniform price auction (e.g., the n selected

¹³We assume that there is no reserve price in the first-stage auction, as, given the all-pay nature of the auction, a reserve price might lead some suppliers being unwilling to participate in the mechanism at all. However, it is possible that a first-stage reserve price could be used to improve the performance of the ERA beyond the gains that we currently identify.

¹⁴Publicly revealing the $(n + 1)^{\text{st}}$ highest bid in the first stage, but not the winning bids, ensures that bidders enter the second stage with symmetric beliefs about the distribution of their rival entrants' costs, which helps to ensure the efficiency of the second-stage allocation and the revenue equivalence of different second-stage formats. Otherwise beliefs about rivals' costs will depend on an entrant's own signal and so they will not be symmetric. As pointed out by a referee, it is unclear whether the procurement cost or efficiency performance of the ERA would change significantly if this information was not revealed, but it would certainly complicate our analysis.

¹⁵We exploit payoff equivalence so that we can verify our solutions are accurate by considering a second-stage auction that operates in a second-price format where we do not need to solve for Chebyshev-approximated bid functions.

¹⁶Fullerton and McAfee (1999) also show how an all-pay auction with an entry subsidy induces efficient entry into the related mechanism of a research tournament. See Ye (2007) for an extensive discussion of the differences between the Fullerton and McAfee model and the type of model considered here.

entrants would each pay the $(n + 1)^{\text{st}}$ highest bid) or a discriminatory auction where only the selected entrants pay their bid. As shown by Ye (2007) in his Proposition 3, first-round bid functions will be strictly decreasing in these formats only if, in the first stage, a bidder's expected gain from participating in the second stage conditional on being the marginal (i.e., n^{th}) selected entrant is strictly decreasing in the cost signal. Whether this will be the case will depend on both n and the specifics of the cost and signal distributions considered, and we have found ranges of signals over which this monotonicity property does not hold for some of the parameters that we consider in this article. However, if the property does hold, then both efficiency and the expected payoffs of all parties should be the same as under the all-pay format, so it is still insightful to use the all-pay format even if it is used less frequently than other formats in the real-world.

We are interested in the symmetric Bayesian Nash equilibrium where the first-round bidding strategy $\gamma^*(\cdot)$ is monotone in the signal and the second-round bid function $\beta^*(\cdot; \bar{s})$ is monotone in the cost for every possible value of the revealed $(n + 1)^{\text{st}}$ highest first-round bid. Note that in such an equilibrium, revealing the $(n + 1)^{\text{st}}$ highest bid is equivalent to revealing the $(n + 1)^{\text{st}}$ lowest signal \bar{s} . Thus, entrants in the second round use Bayes' Rule to compute the distribution of the costs of any one of their opponents as $F_{C|S \leq \bar{s}}(\cdot)$, with density $f_{C|S \leq \bar{s}}(\cdot)$.

The second-stage bid function $\beta^*(\cdot; \bar{s})$ solves the differential equation

$$\beta^{*\prime}(c; \bar{s}) = (n - 1) (\beta^*(c; \bar{s}) - c) \left[\frac{f_{C|S \leq \bar{s}}(c)}{1 - F_{C|S \leq \bar{s}}(c)} \right],$$

with the boundary condition $\beta^*(r; \bar{s}) = r$. To solve for the bid function $\gamma^*(\cdot)$, note that the profit of a bidder with cost c who is invited to enter the auction, when the $(n + 1)^{\text{st}}$ signal is \bar{s} , is, in the case of a second-round auction operating in a first-price format,

$$\Pi(c; \bar{s}) = (\beta^*(c; \bar{s}) - c) [1 - F_{C|S \leq \bar{s}}(c)]^{n-1},$$

where we are exploiting the fact that the entry cost is being fully subsidized by the procurer.

The first-stage bid function $\gamma^*(s)$ should solve

$$\gamma^*(s) \equiv \arg \max_g \int_{\gamma^*(s)}^{s_{\max}} \int_{\underline{c}}^{\bar{c}} \Pi(c; \bar{s}) dF_{C|S=s}(c) dF_{S^{(n:N-1)}}(\bar{s}) - g,$$

where $F_{C|S=s}(\cdot)$ is the conditional distribution of the cost given a signal of s and $F_{S^{(n:N-1)}}(\cdot)$ is the distribution of the n^{th} highest signal of the remaining $N - 1$ bidders, with pdf $f_{S^{(n:N-1)}}(\cdot)$. Solving this maximization problem and imposing that in equilibrium $g = \gamma^*(s)$ gives the differential equation

$$\gamma^{*\prime}(s) = - \int_{\underline{c}}^{\bar{c}} \Pi(c; s) f_{S^{(n:N-1)}}(s) dF_{C|S=s}(c),$$

with boundary condition $\gamma^*(s_{\max}) = 0$. As entry is subsidized, the integrand on the right-hand side is always positive (a selected firm cannot lose money in the second stage because the entry subsidy covers the entry cost), so the first-stage bid function will be monotonically decreasing in

s. The total procurement cost will be equal to the price paid in the second round *plus* the sum of the entry subsidies for the n entrants *less* the sum of the first-stage bids.

The value of n , the number of firms selected for the second stage, plays an important role in the ERA, and is chosen by the procurer. So that the first stage is meaningful, it is natural to assume that $1 \leq n < N$.¹⁷ If $n = 1$, then the procurer selects one firm in the first stage and effectively makes that firm a take-it-or-leave-it offer to complete the project at the second-stage reserve price r (we assume that the chosen firm would still have to pay the entry cost even if it faces no competition in the second stage). While this case may seem somewhat trivial, this is exactly what an efficiency-maximizing procurer would like to do when bidders are well-informed about their costs, because the firm with the best signal will be the one that should be selected to complete the project with high probability. Therefore for the rest of this section we allow for the possibility that $n = 1$, but in the empirical application we will restrict ourselves to allowing only $2 \leq n < N$. As we find that the ERA outperforms the FPAFE, changing this constraint could only strengthen our conclusions.

2.2.3 Efficiency of Bidder Entry

We now compare the efficiency of the ERA and the FPAFE using an example to illustrate how the advantage of fixing the number of entrants depends on how well informed bidders are about their values before they pay the entry cost. We say that a mechanism is more efficient when the expected social costs of completing the project (the cost of the auction winner or c_0 if there is no winner, plus the sum of entry costs) is lower than in the other mechanism. Recall from the discussion in the Introduction that although entry strategies in the FPAFE are optimal from an efficiency perspective whether or not signals are affiliated with true costs (Levin and Smith (1994), Gentry and Li (2012)), if a potential bidder's entry strategy is only allowed to depend on that firm's private information, there may be an efficiency advantage to fixing the number of entrants, for example through an ERA. Doing so will help overcome the problem that conditioning only on a bidder's own private information will result in uncoordinated entry, and possibly too many or too few entrants.

It is straightforward to show that an ERA must be more efficient when either signals are completely uninformative, as in the LS model, or completely informative, as in the S model. In the former case, Milgrom (2004) p. 225–227, shows this result.¹⁸ In the case of the S model, efficiency requires that only the firm with the lowest cost signal enters (assuming this signal is less than c_0), and this is achieved through an ERA with $n = 1$ but cannot be guaranteed in any mechanism

¹⁷Considering a case where all suppliers are invited to participate in the second stage ($n = N$) would involve no competition in the first stage, so all bids would be zero. In essence, the procurer would subsidize the entry costs of all the potential suppliers and then run a standard FPA in which everyone knows their true costs. For the parameters that we look at in Section 2.5 choosing $n = N$ would not be optimal.

¹⁸The results reflects the fact that when signals are uninformative, the sum of the expected cost of the lowest cost entrant plus total entry costs will be convex in the number of realized entrants either in an FPAFE or when a fixed number of entrants is chosen randomly by the procurer. Under the LS assumptions, the expected cost of the lowest cost firm among n will be the same in both of these cases.

where entry decisions are not coordinated.

The situation is more complicated when signals are partially informative. An ERA, where the number of entrants is fixed in advance, does not allow the number of entrants to depend on the private information that suppliers have about their costs, which is correlated with how valuable their entry is likely to be from an efficiency perspective. To illustrate, consider a specific example where there are four suppliers with costs distributed lognormally with location parameter $\mu_C = -0.09$ and scale parameter $\sigma_C = 0.2$ (so that the mean cost is 0.93 and the standard deviation is 0.19). We assume that c_0 and the reserve in the FPA and the second round of the ERA are equal to 0.85. We truncate the costs to the interval $[0, 4.75]$. $K = 0.02$ and $S_i = C_i \cdot \exp(\epsilon_i)$, where $\epsilon_i \sim N(0, \sigma_\epsilon^2)$. Initially we assume that $\sigma_\epsilon = 0.2$.

For these parameters, the equilibrium signal threshold in the FPAFE would be 0.829 and, on average, 1.46 suppliers would enter. An ERA designed to maximize expected efficiency, without knowledge of what suppliers' signals are, would select a single entrant. To illustrate why the FPAFE might be more efficient, suppose that the signals of the four potential entrants are $\{0.78, 0.79, 0.80, 0.90\}$. *Given these signals*, expected efficiency would be maximized by having the three firms with the lowest signals enter (expected social costs are 0.812), which in this example is what happens in the FPAFE, rather than just the firm with the lowest signal, which is what happens in the ERA (expected social costs 0.823).

Given these parameters, the FPAFE also dominates in expectation (i.e., integrating out over all possible realizations of suppliers' cost signals). To show how this result relates to the precision of the suppliers' information, the top panel of Figure 2-1 shows the percent decrease in expected social costs from using an ERA rather than an FPAFE, allowing for either $n = 1, 2,$ or 3 in the ERA. Positive numbers reflect the ERA outperforming the FPAFE, and negative numbers reflect the FPAFE doing better. The solid line indicates the gain from using an ERA with n chosen to maximize efficiency, so that it marks the upper envelope of the lines for the individual ERAs. On the horizontal axis, we measure the precision of information by a parameter $\alpha \equiv \sigma_\epsilon^2 / (\sigma_\epsilon^2 + \sigma_C^2)$ which varies from 0 to 1.¹⁹ Holding σ_C fixed, as $\alpha \rightarrow 1$, the model will tend towards the informational assumptions of the LS model, and as $\alpha \rightarrow 0$, it tends towards the informational assumptions of the S model. The other parameters are the same as in the preceding example. The middle panel shows the expected cost of completing the project for each of these mechanisms (the cost is c_0 if no bidder submits a bid below the reserve price), and the bottom panel shows the expected number of entrants in the FPAFE and the efficiency-maximizing ERA.

When α is close to 0 or 1, an ERA is more efficient in expectation when n is chosen appropriately. However, for values of α between 0.45 and 0.7 the FPAFE is more efficient than any of the ERAs. Whether the FPAFE does better because it economizes on entry costs or lowers the expected cost of completing the project depends on the value of α , because this affects how many firms the ERA should select for the second stage. When $\alpha = 0.5$, the optimal n is 1 and we are in the case considered above where the additional entry under the FPAFE lowers the expected cost of

¹⁹If the support of the lognormal distribution were not truncated, then the conditional distribution of the cost given a signal s would be lognormal with location parameter $\alpha\mu_C + (1 - \alpha)\log(s)$ and scale parameter $\sigma_C\sqrt{\alpha}$.

completing the project by enough to lower social costs. When $\alpha = 0.6$, a different logic applies. In this case, the optimal ERA has $n = 2$, and there is higher expenditure on entry costs and a lower cost to complete the project than in the FPAFE. However, because the firm with the second-lowest signal will be less likely to win in the second-stage auction, especially if it would not have entered the FPAFE, the entry cost effect dominates and the FPAFE is again more efficient.²⁰ The fact that which mechanism is more efficient can depend on the parameters, including the degree of selection, motivates us to estimate our model using real-world data on procurement auctions.




Figure 2-1: Performance of an ERA (with $n = 1, 2$, or 3) and an FPAFE. We assume $N = 4$, $F_C(c) \sim LN(-0.09, 0.2)$, $K = 0.02$, $r = c_0 = 0.85$, and $S_i = C_i \cdot \exp(\epsilon_i)$. $\alpha \equiv \sigma_\epsilon^2 / (\sigma_\epsilon^2 + \sigma_C^2)$. In the top panel the y-axis shows the percentage reduction in social costs from using an ERA rather than an FPAFE. Outcomes are calculated using 500,000 simulations.

²⁰As illustration, the supplier with the second lowest cost signal wins in the second stage of the ERA with probability 0.314 if its cost signal would have led it to enter the FPAFE and with probability 0.178 if it would not have entered the FPAFE.

2.3 Data

We estimate our model using a sample of FPAFE procurement auctions for bridge construction projects conducted by the Oklahoma and Texas Departments of Transportation (DoTs) from March 2000 through August 2003. This data is part of the sample originally collected and used by De Silva, Dunne, Kankanamge, and Kosmopoulou (2008), and we detail how our sample was chosen in Appendix B.2. As they describe, the data comes from all areas of Oklahoma, but only the North Texas and Panhandle regions of Texas, so that the geographic, construction, and economic conditions are likely to be similar across the two states.

For each project, the data contain information on the number of “planholders”, which are the construction companies that were interested enough in the project to purchase, at a price of about \$100, the detailed plans for the project developed by the state’s engineer. The list of planholders is publicly available prior to the auction and we use the number of planholders to define N , the commonly known number of potential suppliers for the project.

	Variable	Mean	Std. Dev.	25 th -tile	50 th -tile	75 th -tile
OK # obs. = 262	Potential Suppliers	6.882	1.997	5	7	8
	Number of Entrants (Bidders)	4.344	1.570	3	4	5
	Unemployment	4.136	0.971	3.2	4.1	4.6
	Engineer’s Estimate	491,120	543,662	226,338	320,755	523,139
	Winning Bid	451,800	520,945	207,146	288,317	471,151
	Relative Winning Bid	0.917	0.133	0.826	0.888	0.980
TX # obs. = 154	Potential Suppliers	7.909	2.027	6	8	10
	Number of Entrants (Bidders)	4.844	1.657	4	5	6
	Unemployment	5.641	1.214	4.4	5.6	6.5
	Engineer’s Estimate	1,409,158	1,185,331	550,537	967,817	1,977,169
	Winning Bid	1,366,487	1,177,504	530,084	892,737	1,847,180
	Relative Winning Bid	0.976	0.134	0.892	0.965	1.038
ALL # obs. = 416	Potential Suppliers	7.262	2.067	6	7	9
	Number of Entrants (Bidders)	4.529	1.619	3	4	6
	Unemployment	4.693	1.291	3.8	4.4	5.7
	Engineer’s Estimate	830,971	949,131	258,485	446,273	914,745
	Winning Bid	790,410	936,689	240,553	406,155	887,277
	Relative Winning Bid	0.939	0.136	0.844	0.914	1.006

Table 2.1: Summary statistics for our sample of auctions, by state. Winning Bid and Engineer’s Estimate are in dollars, and the Relative Winning Bid is the Winning Bid divided by the Engineer’s Estimate.

Summary statistics for the sample are presented in Table 2.1. On average, the number of potential suppliers, the number of entrants, the level of unemployment, and the number of entrants (observed bidders in the FPAFE) are significantly (at the 1% level) higher in Texas than Oklahoma, but the entry rate (measured by the proportion of planholders that submit bids) is not

significantly different across the states. One feature of the data is that the number of entrants can vary significantly across auctions even when the number of planholders is fixed. For example, out of the 51 Oklahoma auctions with 6 planholders, there are 2, 6, 11, 13, 14 and 5 auctions with 1 through 6 observed bidders respectively. If entry is costly, which is suggested by the fact that around one-third of planholders ultimately chose not to submit bids, this type of variation in the number of entrants suggests that coordinating entry may help to increase efficiency.

Since we will apply the FPAFE model presented in Section 2.2 to these data, we briefly discuss how far its major assumptions are appropriate in this setting. The assumption of independent private costs is common in the literature studying highway procurement auctions (e.g., Krasnokutskaya and Seim (2011), Li and Zheng (2009) and Jofre-Bonet and Pesendorfer (2003)). We also follow much of the literature in assuming that bidders are ex-ante symmetric (see, for example, Li and Zheng (2009)).²¹ Lastly, our model assumes that there is a commonly known reserve price, which is required to prevent suppliers wanting to submit infinite bids when there is some probability that no other firms enter. Therefore we assume that the DoT would reject bids that were too high; specifically we assume that the DoT uses a reserve price equal to 1.5 times the engineer's estimate of the cost of the project, even though the DoTs actual policy for rejecting bids is not public information.²² We choose this relatively high price so that the number of auctions that we have to drop because the DoT accepted bids above this assumed reserve price is very small (2 in OK, 4 in TX).²³ In our data, we also observe that the DoTs reject 28 winning bids that are below this reserve price. However, many of these auctions have winning bids that are well within the range that we see the DoTs accepting; for example, the average rejected winning bid is 0.96 times the engineer's estimate, which is close to the average accepted winning bid in the data, and only 5 of the rejected winning bids are above 1.2 times the engineer's estimate. We assume that these must have been rejected for reasons other than the level of the bid, and so we also exclude these auctions from our estimation sample. In Section 2.5, we consider whether changing the assumed reserve price could alter our conclusions about how well the ERA and an FPAFE perform.

2.4 Estimation

In this section, we describe how we estimate the model and detail the additional parametric assumptions we make on the specification of the model in order to take it to the data. Our estimation approach is based on matching a set of moments predicted by our model, which we calculate using simulations and importance sampling, to a set of moments observed in our data (e.g., the propor-

²¹The De Silva, Dunne, Kankanamge, and Kosmopoulou (2008) data does contain an estimate of capacity utilization and bidder distance from the project site, but for the auctions that we look at these supplier-specific variables have almost no power when trying to predict which supplier wins the auction. For example, in a linear probability model with auction fixed effects, where the dependent variable is a dummy for the winning bidder, these variables are both individually and jointly insignificant and the within-auction R^2 is only 0.0005.

²²An alternative approach (e.g., Li and Zheng (2009)) is to assume that the government acts as an additional bidder if only one firm enters.

²³If our assumption about the reserve price is correct we might still expect to see a few bids submitted above the reserve price if a supplier decides to enter and finds out that it has a cost that is greater than the reserve price.

tion of auctions where 5 suppliers enter). Li and Zhang (2010) provide an earlier example of the estimation of a first-price auction model with non-selective entry using a moment-based approach.²⁴

2.4.1 Model Setup and Specification

We estimate a fully parametric version of the FPAFE model from Section 2.2. We normalize the cost parameters by the engineer’s estimate of costs, so that a supplier cost of 0.9 means 90% of the engineer’s estimate, and an entry cost of 0.03 means that it costs 3% of the engineer’s estimate for the entire project to enter the auction.²⁵ Allowing for the a subscript to denote a particular auction, we assume that costs in an auction a are drawn from a truncated lognormal distribution, $F_{C_a}(\cdot)$, with location parameter μ_{C_a} and scale parameter σ_{C_a} . We truncate costs to the interval $[0, 4.75]$, but, as the average winning bid is 0.94, this truncation has essentially no effect on the distribution of costs for plausible parameters. We assume $S_i = C_i \cdot \exp(\epsilon_i)$, where $\epsilon_i \sim N(0, \sigma_{\epsilon_a}^2)$ and ϵ_i independent across suppliers. However, rather than estimating the values of $\sigma_{\epsilon_a}^2$, we estimate α_a where $\alpha_a \equiv \sigma_{\epsilon_a}^2 / (\sigma_{\epsilon_a}^2 + \sigma_{C_a}^2)$.²⁶

As is implicit in the notation, we allow for parametric, auction-specific observed and unobserved heterogeneity in the structural parameters $\theta_a \equiv \{\mu_{C_a}, \sigma_{C_a}, \alpha_a, K_a\}$. Previous research (e.g. Krasnokutskaya (2011) and Bajari, Hong, and Ryan (2010)) has found both types of heterogeneity in the means of the cost distribution for road construction contracts, and it is plausible that the variance of these distributions and the level of entry costs are also heterogeneous. Note, however, that we assume that within an auction bidders are symmetric in that they draw their project costs and signal noise from the same distributions and that they have the same cost of entering the auction. The parameters for auction a are drawn from independent truncated normal distributions with means that depend on observed covariates X_a . Letting $TRN(\mu, \sigma^2, \underline{c}, \bar{c})$ denote a normal distribution with mean μ and variance σ^2 and truncation points $[\underline{c}, \bar{c}]$, we assume that

$$\begin{aligned}\mu_{C_a} &\sim TRN(X_a \beta_{\mu_C}, \omega_{\mu_C}^2, \underline{c}_{\mu_C}, \bar{c}_{\mu_C}) \\ \sigma_{C_a} &\sim TRN(X_a \beta_{\sigma_C}, \omega_{\sigma_C}^2, \underline{c}_{\sigma_C}, \bar{c}_{\sigma_C}) \\ \alpha_a &\sim TRN(X_a \beta_{\alpha}, \omega_{\alpha}^2, \underline{c}_{\alpha}, \bar{c}_{\alpha}) \\ K_a &\sim TRN(X_a \beta_K, \omega_K^2, \underline{c}_K, \bar{c}_K).\end{aligned}\tag{2.6}$$

In our setting, X_a is a vector consisting of a constant, the unemployment rate in the county where the project is located in the month that the auction is held (in percentage points), and a dummy variable equal to 1 for auctions in Texas. We denote the coefficients for these covariates as $\beta_{0,\times}$, $\beta_{1,\times}$, and $\beta_{2,\times}$, respectively, where \times can be μ_C , σ_C , α , or K . We only allow unemployment

²⁴Specifically Li and Zhang (2010) use an indirect inference method where they match parameters with those from an auxiliary model. Li (2010) provides additional information on this type of approach.

²⁵It is natural to assume that larger projects will require more due diligence and may also require the bidders to post larger bonds that guarantee that they will complete the work. Of course, the imposed linearity of the entry cost with the engineer’s estimate is an assumption.

²⁶Estimating the distribution of α_a means that we are directly estimating how the degree of selection differs across auctions. Understanding this variation is less straightforward when estimating separate distributions of σ_{ϵ_a} and σ_{C_a} .

to affect the location parameter of the cost distribution, setting $\beta_{1,\sigma_C} = \beta_{1,\alpha} = \beta_{1,K} \equiv 0$. The Texas dummy coefficient could capture differences in the procurement process across the states, which might affect entry costs or how much suppliers know about their costs, as well as the costs of completing the project.

To make our estimation approach work, we assume that we, the researchers, know the truncation points.²⁷ For the cost parameters we choose very wide bounds that should have little effect on the parameter distributions for plausible parameters. For example, we truncate K_a at 0.04% and 16% of the engineer's estimate, where the industry literature (cited in our discussion of the results below) indicates that values in the range 1–2% are likely. The parameters to be estimated are $\Gamma \equiv \{\beta_{\mu_C}, \beta_{\sigma_C}, \beta_{\alpha}, \beta_K, \omega_{\mu_C}^2, \omega_{\sigma_C}^2, \omega_{\alpha}^2, \omega_K^2\}$.

2.4.2 Importance Sampling

We estimate the model using a simulated method-of-moments estimator and we use importance sampling to approximate the moments predicted by a given set of parameters (Akerberg (2009)). The motivation behind using importance sampling is that it would be too costly to re-solve a large number of FPAFE models each time one of the parameters changes. Instead, with importance sampling, we can solve a very large number of games once and then only reweight these outcomes as we change the parameters.

To explain the way that we use importance sampling, let $y^e = f(X_a, \theta_a)$ denote an expected outcome of an auction, such as an indicator for whether the winning bid lies between 0.8 and 0.9 times the engineer's estimate. Calculating y^e in our setting is expensive because we have to solve the FPAFE and then simulate outcomes. The density of θ , given X_a and Γ is $\phi(\theta|X_a, \Gamma)$. The expected value of the outcome given X_a and Γ , which we can label $E(y^e|X_a, \Gamma)$, is

$$E(y^e|X_a, \Gamma) = \int f(X_a, \theta_a) \phi(\theta_a|X_a, \Gamma) d\theta_a.$$

As this integral does not generally have an analytic form, one could approximate it for a given value of Γ using simulation by drawing S samples of θ from the distribution $\phi(\theta|X_a, \Gamma)$ and computing $f(X_a, \theta_{as})$ for each draw θ_{as} so that

$$E(y^e|X_a, \Gamma) \approx \frac{1}{S} \sum_{s=1}^S f(X_a, \theta_{as}).$$

However, this calculation would require solving a new set of S auctions whenever one of the parameters in Γ changes. The importance sampling methodology exploits the fact that

$$\int f(X_a, \theta_a) \phi(\theta_a|X_a, \Gamma) d\theta = \int f(X_a, \theta_a) \frac{\phi(\theta_a|X_a, \Gamma)}{\psi(\theta_a|X_a)} \psi(\theta_a|X_a) d\theta,$$

²⁷As noted by Akerberg (2009), we would not be able to use importance sampling as part of our estimation approach if the truncation points depended on the parameters that we are estimating.

where $\psi(\theta_a|X_a)$ is an importance sampling density that has the same support as $\phi(\theta_a|X_a, \Gamma)$, but does not depend on unknown parameters. In this case we can calculate $f(X_a, \theta_{as})$ ²⁸ for a large number of draws taken from $\psi(\theta_a|X_a)$, and then calculate an approximation to $E(y^e|X_a, \Gamma)$, where we only have to re-compute the weights $\frac{\phi(\theta_a|X_a, \Gamma)}{\psi(\theta_a|X_a)}$ when Γ changes, i.e.,

$$E(y^e|X_a, \Gamma) \approx \frac{1}{S} \sum_{s=1}^S f(X_a, \theta_{as}) \frac{\phi(\theta_{as}|X_a, \Gamma)}{\psi(\theta_{as}|X_a)}. \quad (2.7)$$

To estimate the parameters, we group the auctions observed in the data into 32 groups, based on the interaction of the number of plan-holders ($N = 4, \dots, 11$), the state (OK or TX), and whether the unemployment rate in the county is above or below the median level for the state in the month that the auction was held. For each of these groups, we create moments that measure the difference between average observed outcomes (y_a) and expected outcomes ($E(y^e|X_a, \Gamma)$) for the auctions in the group. The outcomes consist of indicator variables for the number of firms submitting bids, indicators for whether the winning bid lies in one of 15 discrete bids (we divide the interval $[0, 1.5]$ into 15 equally sized bins) and the proportion of all suppliers' bids that lie in each of these bins (as some suppliers do not enter, the sum of the proportions across the bins for a given group will be less than 1).

We estimate Γ by minimizing the squared sum, across groups and moments, of these differences, which is a consistent method of moments procedure where the different moments and the different groups receive equal weighting. We use $\psi(\theta_a|X_a) \equiv \phi(\theta_a|X_a, \tilde{\Gamma})$, where $\tilde{\Gamma}$ is given in Table 2.2 and $\phi(\cdot)$ denotes the truncated normal model given in (2.6), as our importance sampling density. For each group of auctions, we take NS draws of θ from this distribution and solve the associated FPAFE model, calculating $y^e = f(X_a, \theta_a)$ using simulation, where NS is equal to 100 times the number of auctions in the group. We choose 50 of these simulated auctions for each of our auctions to calculate $E(y^e|X_a, \Gamma)$ using equation (2.7), as part of the estimation procedure. In Appendix B.3 (available online), we provide some Monte Carlo evidence that an importance sampling-based estimator can perform well for this model even using a smaller number of draws (as few as five) for each auction. As is usual when using importance sampling, accuracy is improved by using an importance sampling density which is close to the true density. In practice, we chose the starting parameters reported in Table 2.2 based on a large number of initial runs, using more diffuse importance sampling densities, which indicated that these parameters values would be close to the parameters that we would estimate in both states.

We calculate standard errors using a bootstrap procedure. For a given bootstrap replication, we re-draw auctions from each of our 32 groups with replacement. For each of these auctions, we choose a new set of 50 simulated auctions, from our large sample of NS for that group, to use in the importance sampling calculation. In this way, the reported standard errors should account for the fact that our estimates are dependent on the particular set of importance sampling draws that

²⁸For each solved auction (i.e., for a given draw of θ_{as}) we use a single simulation as an unbiased estimator of $f(X_a, \theta_{as})$. We could, of course, use more draws to increase efficiency.

	β_0	β_1	β_2	ω	\underline{c}	\bar{c}
	Constant	Unemployment	Texas	Std. Dev.	Lower Trunc.	Upper Trunc.
μ_C	-0.07	-0.005	0	0.012	-0.4	1.0
σ_C	0.05	0	0	0.025	0.0095	0.995
α	0.5	0	0	0.25	0.1	0.9
K	0.015	0	0	0.01	10^{-4}	0.16

Table 2.2: Importance sampling density parameters ($\tilde{\Gamma}$) for estimation. The parameter $\tilde{\Gamma}$ specifies the distribution, as given by the functional form in (2.6).

we use to calculate the moments.

2.4.3 Identification

Gentry and Li (2014) formally study non-parametric identification in a broad class of selective entry auction models where potential bidders first simultaneously decide whether to enter the model based on signals about their values and then compete in a standard auction, such as a first-price auction. With no unobserved heterogeneity, they show that exogenous sources of variation in the equilibrium level of entry can identify the model. When the sources of variation in entry are sufficiently rich (e.g., a continuous variable that shifts entry costs or a continuously varying reserve price with suitable support), one obtains point identification, whereas when there is only discrete variation (e.g., in the number of potential entrants), one obtains partial identification, although they show that the resulting bounds are often informative.

One intuition for why the model is identified with no unobserved heterogeneity comes from the fact that when there are few potential entrants or entry costs are very low, all potential bidders should choose to enter with high probability. In this case, we have exogenous entry and more standard results for the identification of value or cost distributions in auction models will hold (Athey and Haile (2002)). The average level of entry costs and the degree of selection (equivalently, the informativeness of suppliers' signals) will then be identified from how the amount of entry and the distribution of bids change as the number of potential entrants rises.²⁹ If not all potential entrants participate in the auction for larger values of N , entry costs must be nontrivial. Furthermore, as shown by Levin and Smith (1994) in the corollary to their Proposition 9, one implication of the random entry, discussed in the Introduction, that occurs with no selection ($\alpha = 1$) is that the expected procurement cost will rise with the number of potential entrants given optimal or non-strategic reserve prices. When $\alpha < 1$, the sign of the relationship will depend on the parameters (Samuelson (1985), p. 56, discusses the $\alpha = 0$ case). In our sample, controlling for state fixed effects and unemployment, there is a significant (at the 10% level) negative relationship between the log of the procurement cost relative to the engineers estimate and the number of potential

²⁹Variation in the number of potential entrants is a common source of identification in entry models (Berry and Tamer (2006)).

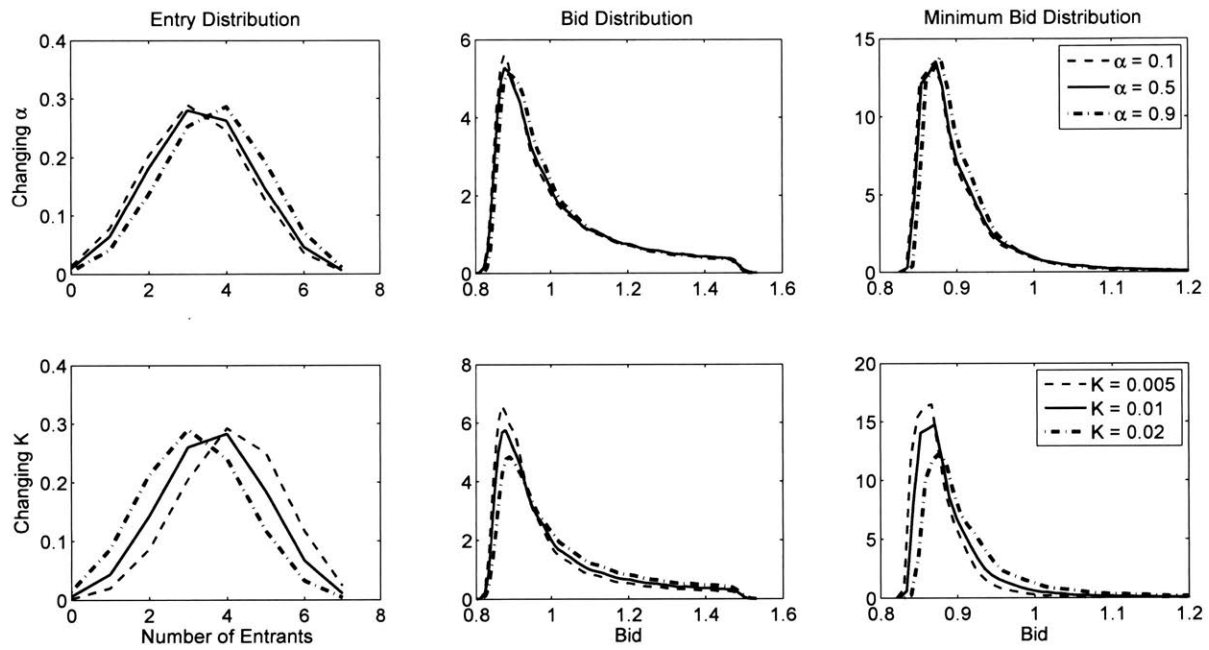


Figure 2-2: Simulated entry, bid, and winning bid distributions for a representative auction for various values of α and K . The top row plots these distributions as α moves from 0.1 to 0.5 to 0.9. The bottom row plots them as K moves from 0.005 to 0.01 to 0.02. We set the location of the value distribution equal to -0.0963 , the scale parameter equal to 0.0705 , α (when we vary K) equal to 0.4979 and K (when we vary α) equal to 0.0147 . Densities are plotted based on 100,000 simulations of each auction.

entrants, indicating that a model with no selection can be rejected. In estimation, the value of $\hat{\alpha}$ will be chosen to not only match this relationship, but also, given the estimated value of \hat{K} how the distribution of all bids and the distribution of the number of entrants change with both N and observables such as unemployment.

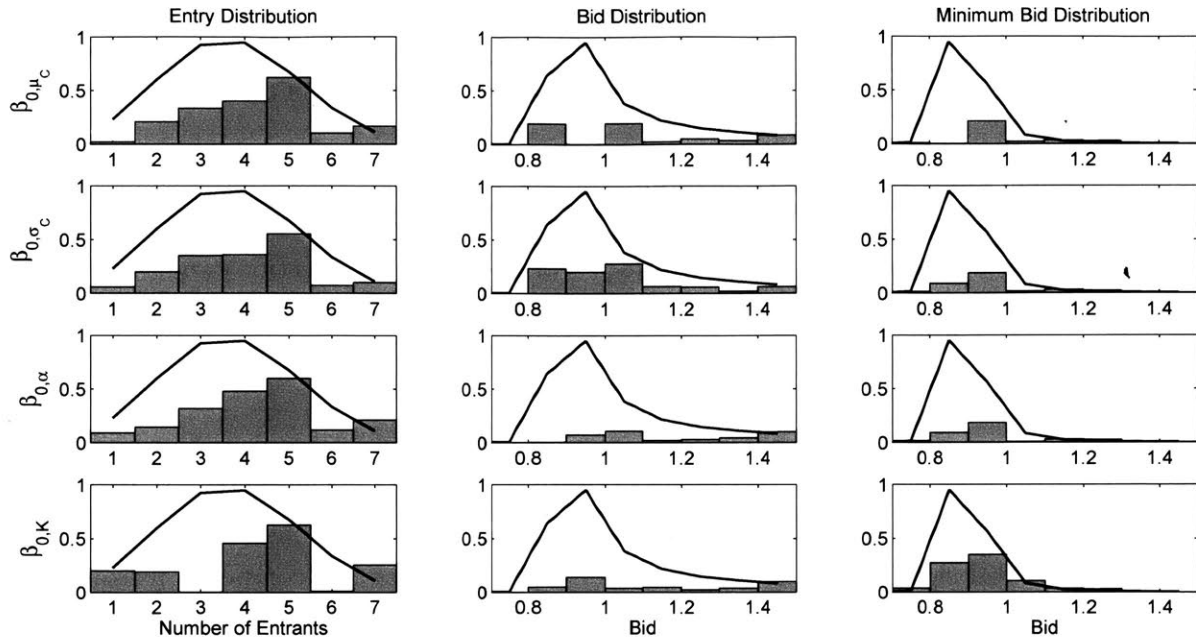
To illustrate how α and K jointly determine observed outcomes, the two rows of Figure 2-2 show how the distributions of the number of entrant, all bids and winning bids, change as we change α from 0.1 (a high degree of selection) to 0.5 to 0.9 (a low degree of selection) and K from 0.005 to 0.01 to 0.02, with fixed N . We draw these figures holding the other parameters fixed at their mean estimated values for auctions in Oklahoma with an unemployment rate of 3%, and we assume that there are seven potential entrants.³⁰

As α increases, there is more entry on average (potential suppliers place less weight on their signals when forming their posteriors so that suppliers with high (bad) signals are more willing to enter). However, suppliers with very low costs become less likely to enter (it is more likely that they receive a bad signal), and those that do enter tend to face less competition from other low cost entrants so that their markups tend to increase. This tends to shift the distribution of winning bids slightly to the right. The distribution of all bids also shifts slightly to the right, as more high cost suppliers tend to enter. As entry costs increase, fewer firms enter but, because of selection, the firms that do enter tend to have lower costs. On its own this would tend to shift the distribution of the bids to the left. However, because they expect to face less competition, entering suppliers will tend to submit higher markups. The increase in markups tends to shift the distribution of bids, as well as the distribution of winning bids, to the right, and, in this case, we can see that this effect actually tends to dominate. As we shall see in Section 2.5 below, the markups that suppliers with moderately high costs submit in an FPAFE play an important role in our counterfactual comparison with an ERA.

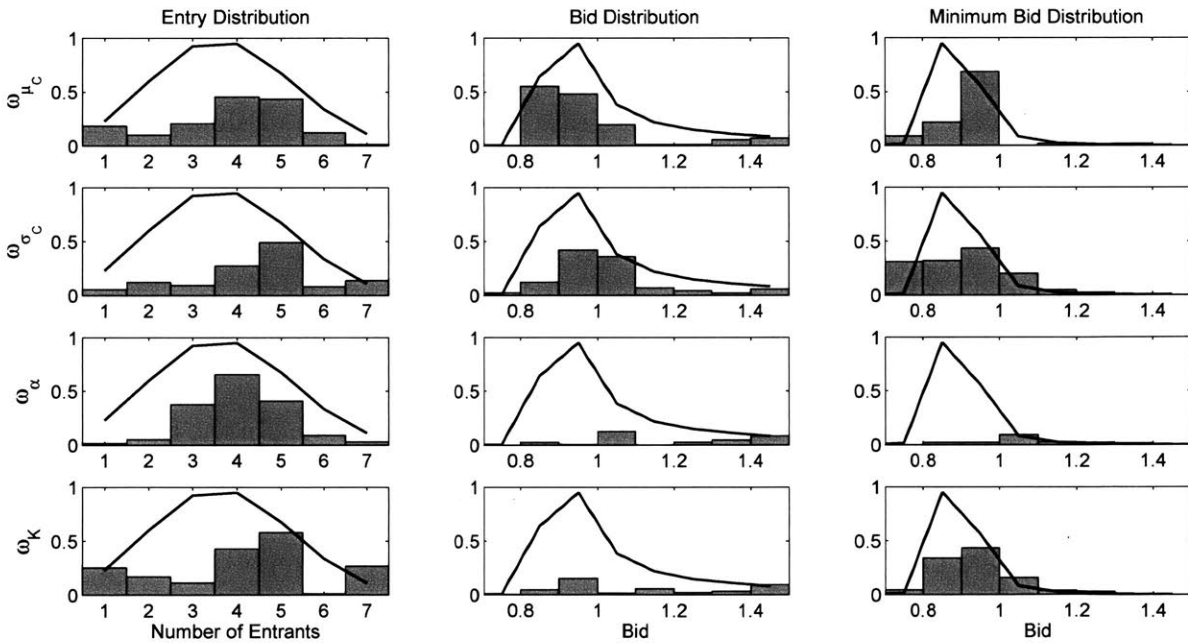
Gentry and Li (2014) show that their identification results extend to a case where there is unobserved heterogeneity that affects both the distribution of values/costs and the level of entry costs, as long as there is still some observed exogenous source of variation in the equilibrium level of entry. Our model, where we allow for independent sources of unobserved heterogeneity to affect the distribution of costs, entry costs, and the degree of selection is not included in the class of models that they consider and our parametric assumptions are likely to be important. To provide some additional transparency into how the moments that we use in estimation identify the parameters of our model, we follow a recent paper by Gentzkow and Shapiro (2013). They suggest a method for illustrating how sensitive (locally) estimated parameters are to particular groups of moments that are defined by a structural model.

The eight rows of Figure 2-3 show the absolute values of Gentzkow and Shapiro's scaled sensitivity parameters Λ for moments based on $N = 7$ potential suppliers and low unemployment in Oklahoma, and the parameter estimates that will be presented in Section 2.5. Gentzkow and

³⁰Specifically, we set the location of the value distribution equal to -0.0963 , the scale parameter equal to 0.0705 , α (when we vary K) equal to 0.4979 and K (when we vary α) equal to 0.0147 . The winning bid distribution is drawn conditioning on at least one potential bidder entering.



(a) Changing β_0



(b) Changing ω

Figure 2-3: Absolute values of the scaled sensitivity parameter from Gentzkow and Shapiro (2013) for Oklahoma auctions with low unemployment and 7 potential suppliers. Each row corresponds to a different parameter ($\beta_{0,x}$ in (a) and ω_x in (b)), and each column corresponds to a different distribution. The scaled sensitivity parameters for each moment (i.e., each bin of the distribution) are represented by the gray bars and the value of the scaled sensitivity parameter can be read from the vertical axis. The solid lines show the relative value of the different moments evaluated at the different parameters.

Shapiro define their sensitivity parameter as $\Lambda = \Sigma_{\theta\gamma} \Sigma_{\gamma\gamma}^{-1}$ where θ are the parameters, γ are the moments and Σ represents the variance-covariance matrix of the stacked vector of the parameters and the moments evaluated at the estimated parameters. This scaling means that we can interpret the vertical axis values as telling us how much a one standard deviation increase in a given moment would change the parameter relative to its asymptotic standard deviation.³¹ Figure 2-3(a) shows the scaled sensitivity parameters for the constant ($\beta_{0,\times}$) parameters, and Figure 2-3(b) shows them for the standard deviation (ω_{\times}) parameters. In each panel, we show the scaled sensitivity parameters for each of the moments. The first column considers the entry moments, with each bar corresponding to the moment that measures the probability of seeing a particular number of entrants. The second column considers the bid distribution, and each bar corresponds to a moment that measures the probability that a supplier's bid lies in the corresponding bin (e.g., 0.7–0.8, 0.8–0.9, etc.). The final column considers the distribution of the winning bid. Figure 2-3 also overlays the simulated values of the associated moments, at the estimated parameters, to put the horizontal axis into perspective; these moments are scaled so that they fit on the figures and the value of the moment should not be read from the vertical axis. All distributions are conditional on at least one entrant, as this is what we observe in the data.

The sensitivity parameters show two intuitive patterns. First, the parameters tend to be most sensitive to the moments that capture the most observations in the data (for example, 4 suppliers enter, rather than 1 or 7 suppliers). Second, all of the parameters are somewhat sensitive to all of the moments, reflecting how, as explained above, a parameter such as the level of entry costs will affect not only the cost distribution of the suppliers that enter and how much entrants markup their bids. That said, the figures suggest that for the four parameters related to α and K , the entry moments are particularly important in the sense that their sensitivity parameters are larger. On the other hand, for ω_{μ_C} and ω_{σ_C} , the moments describing the distributions of either all bids or winning bids play a relatively greater role. This is intuitive, as, for example, increasing both the number of auctions where suppliers tend to have high costs and the number of auctions where they tend to have low costs, which is what happens as ω_{μ_C} increases, should tend to spread out the distribution of the winning bid without necessarily affecting the average amount of entry. Of course, one could take these results further to try to identify more optimal moments for estimation (for example, interactions of the amount of entry and the winning bid), but we view this type of extension as a topic for further research.

2.5 Results

In this section we present the estimation results and our counterfactual analysis of the impact of switching from the observed FPAFE format to an ERA format.

³¹When interpreting the values on the y-axis it is important to remember that the combination of $N = 7$, OK, and low unemployment is only one of the 32 N -state-unemployment level combinations that we use in estimation. The patterns are broadly similar but not identical for the other groups.

2.5.1 Parameter Estimates

The left-hand columns of Table 2.3 present the parameter estimates. To interpret the location and scale parameters together, the right-hand columns report the average across the auctions in each state of (going down the rows) the mean of the cost distribution, the standard deviation of the cost distribution, the value of α , and the entry cost. The mean of the cost distributions are similar in the two states, between 91% and 92% of the engineer’s estimate. As the average *winning* bids are higher than this, it suggests that the markups may be quite substantial. We will return to why this happens, and how the ERA can help to reduce markups, in discussing our counterfactual results.

	β_0	β_1	β_2	ω	OK Mean	TX Mean
	Constant	Unemployment	Texas	Std. Dev.		
μ_C	-0.0868 (0.0324)	-0.0023 (0.0064)	0.0154 (0.0132)	0.0142 (0.0091)	91.1% (0.84%)	92.1% (1.89%)
σ_C	0.0687 (0.0196)		-0.0117 (0.0213)	0.0304 (0.0132)	6.45% (0.97%)	5.61% (0.87%)
α	0.4979 (0.0972)		0.1115 (0.0943)	0.1284 (0.0764)	0.4979 (0.0770)	0.6055 (0.0752)
K	-0.0018 (0.0406)		0.0105 (0.0191)	0.0189 (0.0109)	1.47% (0.26%)	1.88% (0.28%)

Table 2.3: Parameter estimates. Standard errors are in parentheses and are calculated using a non-parametric bootstrap described in the text. The state averages are (going down the rows) the average of costs, the mean standard deviation of the cost distribution, the mean α and the mean entry cost, relative to engineer’s estimate, for auctions in each state.

Mean entry costs are 1.5% of the engineer’s estimate of the cost of the project in Oklahoma and 1.9% in Texas. Although in many industries it is hard to know what entry costs are really reasonable, in this setting we are able to compare them with estimates from manuals that are used in the industry. Halpin (2005) estimates that the cost of researching and preparing bids is around 0.25% to 2% of total project costs. Park and Chapin (1992) estimate that these costs are typically about 1% of the total bid. Our estimates therefore appear quite reasonable.

Our estimates of entry costs are also lower than those that have been estimated based on models with no selection. For example, based on highway paving contracts in California, Krasnokutskaya and Seim (2011) estimate mean entry costs to be approximately 3% of the engineer’s estimate, and Bajari, Hong, and Ryan (2010) estimate them to be 4.5% of the engineer’s estimate. There are some intuitions for why allowing for selection will tend to lower estimated entry costs. In any model, the level of the entry cost will be identified from the fact that the marginal entrant must be indifferent between getting some expected surplus if it enters and not paying the entry cost. Holding the entry probability of other suppliers and the distribution of costs fixed, the expected surplus of the marginal entrant, and therefore the estimated entry cost, will tend to be lower with selection for two reasons. First, with selection, the other suppliers who choose to enter will tend to have lower costs than if their costs were randomly drawn, reducing the marginal entrant’s surplus. Second, as pointed out by one of our referees, because a bidder’s surplus is usually convex in its

own type, the marginal entrant's expected surplus will tend to be lower when it is conditioning on an informative signal than when it has no signal.

The estimates of α suggest that entry is partially selective in this setting, with a mean α of around 0.50 in Oklahoma and 0.61 in Texas; these mean values of α correspond to the scale parameter σ_ϵ of the error distribution being approximately σ_C in Oklahoma and 1.3 times σ_C in Texas. The degree of selection is therefore in the range where our example suggests that it is possible that the FPAFE may be more efficient than an ERA, even though entry into FPAFEs is clearly volatile in our empirical setting.

2.5.2 Regulating Bidder Entry with an Entry Rights Auction

In this subsection, we use our parameter estimates to quantify how procurement costs and efficiencies would change if the DoTs switched from using FPAFEs to using ERAs. Throughout what follows the procurement cost refers to the net amount that the procurer expects to pay to get the project completed. In an FPAFE this is just the expected winning bid, or c_0 if there is no winner. In the ERA it is the expected winning second-round bid (or c_0) plus the sum of entry costs less the expected sum of first-round bids. Social costs refer to the expected sum of entry costs plus the cost of the supplier completing the project (or c_0), and a mechanism is more efficient if and only if it has lower social costs. We assume that the cost of completing a project outside the auction is 1.5 times the engineer's estimate and, unless we state otherwise, this is used as the reserve price in the FPAFE and the second stage of the ERA.

In presenting the results we assume the DoTs would choose the number of entrants in the ERA to minimize the procurement cost (n_{cost}^*), but we will comment below on how the results are very similar if we assume instead that the number of entrants is chosen to minimize social costs (n_{eff}^*). The choice of n_{cost}^* balances several competing effects. In equilibrium, the procurer pays the full entry cost of an additional entrant, and increasing the number of entrants will reduce the amount that any supplier bids in the first-stage auction. On the other hand, when there are more entrants, a firm with lower costs may be selected to complete the project and bidding in the second-stage auction will tend to be more competitive. With the exception of the additional entry cost, the size of these different effects will depend on the degree of selection, as an additional entrant will be more likely to win when signals are less informative.

Tables 2.4 and 2.5 present the results. Table 2.4 compares the outcomes of interest (social costs, procurement costs, and the profits for an individual supplier), and Table 2.5 shows that the average cost of the winner, the average winning bid and the average amount of entry under each mechanism. For the ERA it also shows the average total expenditure on entry costs and the average revenue from first-stage bids. In both tables, we use a range of different parameter values. The baseline case (top row) corresponds to the mean parameters for auctions in Oklahoma (unemployment rate 4.14%) and approximately the mean number of potential suppliers in that state (7). In the remaining rows we change the number of potential suppliers up and down by approximately one standard deviation and each parameter, in turn, to be equal to either the 10th

or the 90th percentiles of its distribution. The parameter that is changed from the baseline is listed in italics.

For all of the parameters listed, we predict that an ERA would both be more efficient and lower procurement costs. For the baseline parameters, both social costs and procurement costs are 2.4% lower under the ERA. The magnitudes are broadly similar across the parameters, except when entry costs are low (*Low K*) when both the absolute and the proportional differences are small. Aggregating across all of the auctions in both states, by taking a draw of the parameters for each of the auctions in our sample, we predict that using ERAs would lower social costs by \$7.78 million (2.49%) and lower procurement costs by \$8.03 million (2.50%). For 3.5% of these draws, we predict that the FPAFE would be more efficient, but in all of these cases entry costs are very low, α has values between 0.39 and 0.77, the efficiency advantage of the FPAFE is small, and we also predict that the ERA would lower procurement costs, although the procurement cost differences are small in these cases.

While improvements of around \$8 million may seem fairly small, it is important to remember that the sample we are using (contracts involving bridge work in Oklahoma and a small section of Texas) are a very small fraction of highway contracts let by US DoTs in the years that we study: in 2007, the public sector spent \$146 billion to build, operate, and maintain highways in the US with three-quarters of this amount spent by state and local governments (Congressional Budget Office (2011)). Moreover, the proportional effects that we are finding are large relative to those produced by auction design changes that are often considered. For example, for the baseline parameters setting a strategic reserve to minimize procurement costs in the FPAFE would only lower procurement costs by 0.03%, whereas we predict that an ERA would lower procurement costs by 2.4%.³²

Interestingly, the ERA also reduces procurement costs by more than adding additional potential suppliers, which is often viewed as providing an upper bound on what an optimal design change can achieve.³³ This can be seen in Table 2.4 where adding two more potential suppliers in an FPAFE to the baseline case (changing it to the *High N* case) lowers procurement costs by 0.6% of the engineer's estimate, compared to a 2.4% reduction from switching to an ERA.

Table 2.5 shows that in most cases the ERA is more efficient because there is both less entry and the average cost of the winner is lower. The fact that it can have both of these advantages reflects the uncoordinated nature of entry in the FPAFE: even though more firms may enter an FPAFE on average, there may be instances where no firms, or only one firm, enter(s) and at least one additional entrant, who will always be in the ERA as $n_{\text{cost}}^* = 2$ for these parameters,

³²We calculate the optimal reserve price by computing procurement costs on a fine grid of reserve prices with 0.01 spacing, using 500,000 simulations to calculate expected costs in each case. Efficiency in the FPAFE falls when a reserve price other than 1.5 (the value of c_0) is used. We have also found that an optimal reserve price in the second stage of the ERA can only improve its performance by a trivial amount.

³³Bulow and Klemperer (1996) show that in an IPV model with a fixed number of symmetric bidders, an additional bidder is more valuable to the seller than being able to implement an optimal auction design. This theoretical result does not hold with endogenous entry, but its logic is often used to motivate the importance of trying to generate additional interest in the assets being auctioned (e.g., Klemperer (2002)).

Case	Parameters					Social Costs			Procurement Costs			100 × Avg. Supplier Profits		
	N	μ_C	σ_C	α	K	FPAFE	ERA	% Diff	FPAFE	ERA	% Diff	FPAFE	ERA	% Diff
<i>Baseline</i>	7	-0.0963	0.0705	0.4979	0.0147	0.886	0.865	-2.41%	0.914	0.892	-2.41%	0.399	0.391	-2.01%
<i>Low N</i>	5	-0.0963	0.0705	0.4979	0.0147	0.891	0.874	-1.91%	0.923	0.905	-1.95%	0.648	0.636	-1.85%
<i>High N</i>	9	-0.0963	0.0705	0.4979	0.0147	0.883	0.859	-2.74%	0.908	0.883	-2.75%	0.280	0.277	-1.07%
<i>Low μ_C</i>	7	-0.1145	0.0705	0.4979	0.0147	0.871	0.850	-2.47%	0.899	0.876	-2.56%	0.397	0.384	-3.27%
<i>High μ_C</i>	7	-0.0781	0.0705	0.4979	0.0147	0.901	0.880	-2.36%	0.930	0.908	-2.37%	0.404	0.397	-1.73%
<i>Low σ_C</i>	7	-0.0963	0.0335	0.4979	0.0147	0.926	0.902	-2.54%	0.940	0.915	-2.65%	0.201	0.192	-4.48%
<i>High σ_C</i>	7	-0.0963	0.1081	0.4979	0.0147	0.848	0.829	-2.21%	0.890	0.870	-2.25%	0.599	0.577	-3.67%
<i>Low α</i>	7	-0.0963	0.0705	0.3339	0.0147	0.884	0.861	-2.64%	0.915	0.891	-2.62%	0.441	0.439	-0.45%
<i>High α</i>	7	-0.0963	0.0705	0.6620	0.0147	0.889	0.871	-2.12%	0.914	0.893	-2.29%	0.348	0.328	-5.75%
<i>Low K</i>	7	-0.0963	0.0705	0.4979	0.0024	0.839	0.837	-0.21%	0.870	0.867	-0.34%	0.442	0.418	-5.42%
<i>High K</i>	7	-0.0963	0.0705	0.4979	0.0303	0.934	0.896	-4.11%	0.961	0.923	-3.95%	0.380	0.390	2.63%

Table 2.4: Social costs (lower numbers mean a mechanism is more efficient), procurement costs, and mean profits for an individual supplier relative to the engineer’s estimate under the FPAFE and an ERA with the number of entrants chosen to minimize procurement costs, for various parameters. The “% Diff” column corresponds to the percent change in the quantity when moving from an FPAFE to an ERA. Average outcomes are estimated using 500,000 simulations.

might decrease the expected cost of completing the project substantially.³⁴ However, most of the ERA's efficiency advantage comes from lower entry costs (for example, this accounts for 91% of the reduction in social costs in the base case). This fact helps to explain why the one case in Table 2.4 where the social costs are almost identical under the two mechanisms is where entry costs are very low (*Low K*, where K equals just 0.2% of the engineer's estimate). In this case, the additional entry into the FPAFE lowers the expected cost of completing the project while increasing entry costs by a relatively small amount. However, even in this *Low K* case, the ERA is still more efficient.

A natural question is how well the ERA that we consider does relative to a mechanism that was designed to maximize efficiency. As already noted, one weakness of our ERA is that the number of selected entrants does not depend on suppliers' signals. Another weakness, relative to a sequential search procedure, is that the procurer cannot get bids from more firms if the ones for whom it initially chooses to pay the entry cost turn out to have high costs. Based on our results so far, it is unclear whether these weaknesses are important or not. To investigate this further, we can compare social costs under our ERA and under two more efficient but hypothetical alternatives.³⁵ In the first alternative, which we can call the "signal-dependent ERA" (S-D ERA), we allow the procurer to choose the number of entrants with knowledge of suppliers' signals but without using any information on the costs of the selected entrants. In the second alternative, which we can call the "sequential search procedure" (SSP), we consider a procurer who knows the suppliers' signals and then approaches suppliers sequentially, in ascending order of their signals. The procurer asks an additional supplier to enter if, based on its signal and the lowest cost to complete the project of the suppliers that have entered already, the procurer expects that its entry will reduce the cost of completing the contract by enough to offset the additional entry cost.

For the baseline parameters, we find that our ERA performs very well compared to these hypothetical alternatives. Social costs for the FPAFE, ERA, S-D ERA, and SSP are 0.886, 0.865, 0.863 and 0.857, respectively. The ERA with fixed n therefore achieves 91% and 72% of the efficiency gains of switching from the FPAFE to the S-D ERA and SSP respectively, while maintaining the attractive elements of simultaneous bidding and being quite straightforward to implement. Qualitatively, these comparisons also hold for the other parameters in Table 2.4: with the exception of the *Low K* case, where all mechanisms have very similar social costs, the fixed n ERA achieves at least 78% of the gains that would result from switching from the FPAFE to the S-D ERA, and at least 60% of the gain from switching from the FPAFE to the SSP.

As can be seen in Table 2.4, the benefits of the increased efficiency of the ERA often accrue entirely to the procurer, with the profits of the average supplier tending to fall slightly (the percent changes in the supplier's profits are non-trivial because supplier profits are quite small under either

³⁴For the baseline parameters, there is no entrant into the FPAFE in 1% of simulated auctions, in which case the procurement cost is 1.5. This would only happen if both of the selected firms had costs above 1.5 in the ERA and this happens with almost zero probability, and never in our simulations. If we condition our comparison on cases where at least one firm enters the FPAFE, so that we ignore simulations where the outcome of the FPAFE is particularly inefficient, the expected winner's cost is very similar for the ERA and the FPAFE, but the overall differences in social costs, procurement costs and bidder profits are almost identical to those shown in Table 2.4.

³⁵The alternatives are hypothetical in the sense that we ignore the fact that it might be difficult or very expensive for the procurer to get suppliers to report their signals or costs truthfully.

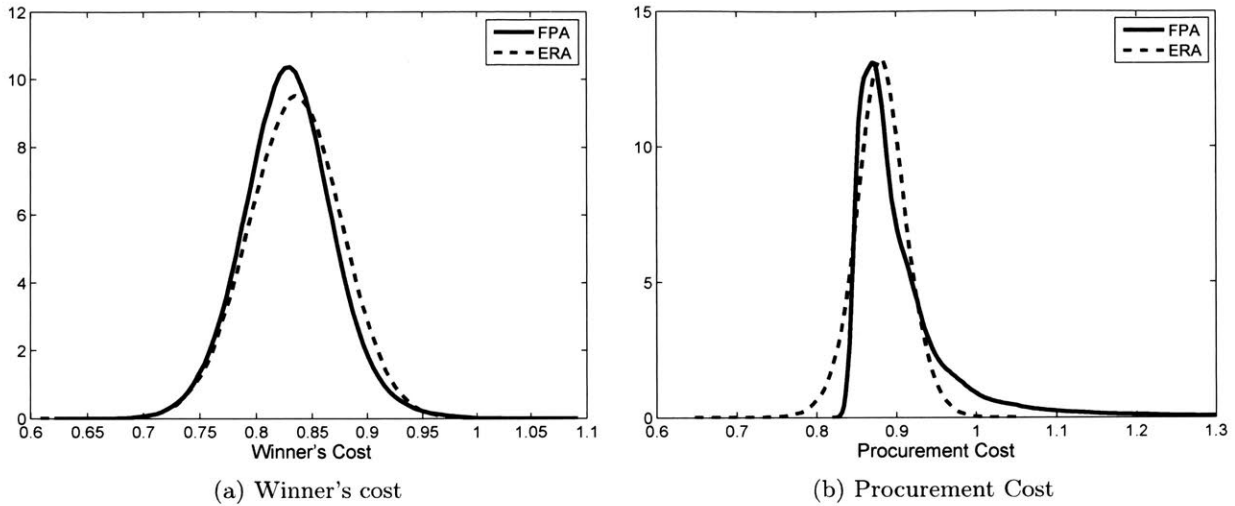


Figure 2-4: Kernel densities for (a) the winner's cost of completing the contract and (b) total procurement costs in both the FPAFE and the ERA, using the baseline parameters. The costs to the procurer in the ERA includes entry costs paid less the revenue from the first-stage bids. Densities are plotted based on 500,000 simulations. Simulations of the FPAFE where no suppliers entered are ignored.

mechanism), while procurement costs tend to fall by slightly more than the gain in efficiency. Aggregating over the whole sample, total supplier profits fall by \$242,000 (2.82%) under the ERA. An exception is when entry costs are high (*High K*), as the bidders also benefit under the ERA, but even in this case the procurer benefits more than the bidders do.

To understand why procurement costs fall, recall that in an FPAFE the procurer only pays the winning bid whereas in an ERA it pays the winning bid and the entry costs of all of the suppliers admitted to the second stage, but also receives the first-stage bids of all of the suppliers. Table 2.5 shows that in general first-stage revenues and the entry costs roughly offset from the procurer's perspective so that the fall in procurement costs comes primarily from the fact that winning bids tend to be significantly lower in the second stage of the ERA than in the FPAFE.

Winning bids are lower because the winner's markup tends to be lower in the ERA, not because the winner's cost is lower. Figure 2-4 plots the distribution of the winner's cost and the procurement cost for both mechanisms at the baseline parameters. The winner cost distributions are very similar in shape for the two mechanisms, but the distribution of procurement costs for the FPAFE has a more pronounced right tail which is driven by the fact that winners with relatively high costs have high markups in the FPAFE.³⁶

³⁶The left end of the procurement cost distributions in Figure 2-4(b) are also of note. In an FPAFE there is a single bid function for given parameter values (θ_a) because entrants do not know how many other planholders are bidding. This bid function maps costs to the interval $[\underline{b}, r]$, and the sharp cutoff in the kernel density for the FPAFE corresponds to \underline{b} . On the other hand, bid functions in the ERA depend on the realization of \bar{s} , and this causes \underline{b} to vary depending on signals even when the parameters of the auction are fixed. As a result, there is no sharp cutoff in the density for the ERA.

The difference in markups for the baseline parameters can be seen more clearly in Figure 2-5, which compares the FPAFE bid function and the average of the bid functions from the second stage of the ERA with $n_{\text{cost}}^* = 2$.³⁷ The diagrams include the pdfs of the cost distribution of a typical entrant into each mechanism to indicate which parts of the bid curve are the most empirically relevant (Figure 2-4 shows the pdf for the winning bidder). The markups in the ERA are slightly larger than in the FPAFE for low-cost bidders, but they are smaller for higher costs, which the majority of bidders, and even a large share of winning bidders, have. This leads to an average markup for the winning bidder in the FPAFE of 9.4% (ignoring cases where there is no sale), whereas the average markup for the winning bidder in the second stage of the ERA is only 6.8%.

Markups are lower in the ERA even though for the baseline parameters a bidder should expect to face 2.9 rival bidders in the FPAFE and only one rival bidder in the second stage of the ERA with $n_{\text{cost}}^* = 2$. This reflects how the lack of coordination in entry affects bidding incentives. In the FPAFE, a bidder with a high cost knows that he is most likely to win when no other firms enter, in which case he should be submitting a bid at the reserve price; even though it is unlikely that no other firms enter, the possibility leads him to submit a high bid in equilibrium.³⁸ On the other hand, in an ERA with $n_{\text{cost}}^* = 2$ a bidder knows that he can only win if he submits a lower bid than his competitor. Therefore, as a bidder's cost increases he will tend to bid more aggressively (i.e., with a lower markup), as can be seen in Figure 2-5 where the average ERA markup is monotonically decreasing in the bidder's cost for all costs below 1.2.³⁹

We can also compare the FPAFE and ERA assuming that the number of selected firms in the ERA is chosen to maximize efficiency (n_{eff}^*). Up to integer constraints, n_{eff}^* balances the cost of an additional entrant against the possible gain that will come from lowering the cost of completing the project. Given our estimates, it turns out that using n_{eff}^* or n_{cost}^* does not change the results very much because, given the integer constraint, n_{eff}^* and n_{cost}^* are usually the same. This is the case for all of the parameters in Tables 2.4 and 2.5. When we draw parameters for all of the auctions in our data, there are some cases with low entry costs where the optimal n s are higher (for example, in 12% of cases n_{cost}^* is greater than 2 and it takes on values as high as 6). In 17% of cases we find that $n_{\text{eff}}^* > n_{\text{cost}}^*$. However, although this is a significant proportion of cases, we find that using n_{eff}^* rather than n_{cost}^* only reduces social costs slightly, by less than \$0.1 million in aggregate (recall that the advantage of the ERA with n_{cost}^* over the FPAFE is \$7.8 million). In none of our experiments have we found a case where $n_{\text{eff}}^* < n_{\text{cost}}^*$. The intuition here is that

³⁷The bid function in the second stage of the ERA depends on the value of the highest losing bid in the first stage. Here the average bid function in the ERA is defined by $\bar{\beta}(c) \equiv \int \beta(c; \bar{s}) f_{\bar{s}}(\bar{s}) d\bar{s}$, where $f_{\bar{s}}(\cdot)$ is the density of the $(n + 1)^{\text{st}}$ lowest signal.

³⁸That high-cost bidders bid close to the reserve may suggest that the level of the reserve plays an important role in determining why average winning markups in the FPAFE are so large. However, Figure 2-4 shows that the distribution of the winner's cost is usually not in the region where bids are close to the reserve even if it is often in the range where markups are high. For the baseline parameters, reducing the reserve price to 1.1 only reduces the average winner's markup to 8.3% (compared to 9.4% when the reserve is 1.5), and increasing it to 4.75 (the maximum of the support of the cost distribution) only increases it to 10.9%.

³⁹Above 1.2, the markup increases reflecting the shape of the upper tail of the cost distribution. However, it is very rare for the winner in the second stage of an ERA to have a cost above 1.2, so this has little effect on procurement costs.




Figure 2-5: Comparing the bid function for a FPAFE (heavy solid line) with the average bid function in an ERA with $n = 2$ (heavy dashed line) for the baseline parameters ($\mu_C = -0.0963$, $\sigma_C = 0.0705$, $\alpha = 0.4979$, $K = 0.0147$, $N = 7$, and a reserve price $r = 1.5$). The dotted line is the 45° line, for comparison. The lighter lines show the densities for the cost of a typical entrant into each mechanism.

allowing an additional entrant reduces the bids that firms make in the first stage of the ERA, and a procurer that seeks to minimize procurement costs therefore has an incentive to exert market power by reducing the number of entry slots. This incentive appears to typically outweigh the incentive to reduce second-stage markups by increasing competition. This is sensible given that, even with $n = 2$, markups in the second stage of the ERA tend to be small.

Given the sizable gains of using an ERA that we estimate above, a potential puzzle is why they are not used more often for the sort of public procurement we study here. There are several possible explanations. First, it may be the case that the reduced supplier profits are actually harmful to the procurer in the long run, especially when it wants to let many contracts over a number of years. In this case, it may benefit when there are a large number of firms in the industry who are potentially interested in bidding, although, as noted above, the benefits from increasing the number of potential suppliers in an FPAFE are quite small in our setting. Kagel, Pevnitskaya, and Ye (2008) show that in a laboratory setting students tend to overbid in ERAs, leading to bankruptcies. If this is true in the field as well, this might provide an additional route by which ERAs might diminish long-run competition.

Second, in empirical settings like ours, many suppliers are relatively small companies that face liquidity constraints, and these may be tightened by having to make binding bids into ERAs.⁴⁰ A system that favors bigger companies may be undesirable, as one of the goals of public procurement is often to favor smaller or minority-owned businesses (e.g., Krasnokutskaya and Seim (2011),

⁴⁰Following the same logic, liquidity constraints or risk aversion might tend to reduce first-stage bids, increasing procurement costs.

Athey, Coey, and Levin (2013)). On the other hand, to the extent that suppliers are also hurt by uncoordinated entry when entry is costly, one could argue that these same concerns could also provide motivation for regulating entry.

Third, designing an ERA appropriately, in particular choosing n , may require the procurer to have more information, for example about the distribution of costs or the entry cost, than running a FPAFE. In our empirical setting, when we take a draw of the parameters for each of the auctions in our sample, $n = 2$ is the procurer-optimal choice for 88% of these draws, and it is always the best choice unless K is small.⁴¹ Therefore in general, adopting a policy that two suppliers will be selected should be effective.

Fourth, Quint and Hendricks (2013) suggest that ERA may provide sellers with perverse incentives to try to sell off contracts that, once they undertake research, bidders will find to be worthless. In some settings this may be an important consideration, but it may be less of a concern when a state agency is letting a large number of contracts and so should want to maintain a reputation for integrity.

Finally, it may of course be that some of the standard assumptions used to model procurement auctions, which we have maintained here, such independent private values, simultaneous entry decisions and independent signals (e.g., Athey, Levin, and Seira (2011), Krasnokutskaya and Seim (2011) or Li and Zheng (2011)) are incorrect and affect the comparison of the different designs. For example, if bidder signals are correlated or entry decisions are made sequentially, entry into the FPAFE may be less volatile, mitigating some of the gains to using the ERA. Alternatively, it may be that our assumption that firms act competitively in both mechanisms is not correct: if firms are more able to collude in an ERA, then this might offset the reductions in procurement costs that our model predicts.

2.6 Conclusion

In procurement settings, where it is expensive for potential suppliers to learn their costs of completing a project, unregulated entry can lead to volatile amounts of supplier participation. If suppliers have either very little information or almost perfect information about their costs prior to entering, then fixing the number of entrants, and choosing them through some type of entry rights auction procedure, will tend to raise efficiency. However, when suppliers are partially informed about their costs, fixing the number of entrants in advance can be inefficient because it ignores the fact that suppliers' private information can determine how valuable additional entry would be.

We compare the performance of a standard first-price auction with free entry and an entry rights mechanism in the empirical setting of highway construction projects involving bridge work in Oklahoma and Texas where we find suppliers' information to be partially informative and entry costs to be significant. For the vast majority of projects in our sample we predict that using an entry rights auction procedure would significantly increase efficiency and reduce procurement costs, and

⁴¹Exactly how small K needs to be to make having more entrants optimal depends on the other parameters, but particularly the level of α . As α increases, it becomes more desirable to select more entrants.

that using ERAs would have much larger effects on procurement costs than alternative procedures such as setting an optimal reserve price in a standard auction or even increasing the number of potential bidders. This reflects the fact that an entry rights procedure can both economize on entry costs by coordinating entry, but also encourage even high-cost suppliers to bid aggressively by ensuring that they will face competition.

Case	FPAFE			ERA				
	Comp. Cost	Win Bid	No. Entrants	Comp. Cost	Win Bid	n_{cost}^*/N	Total Entry Costs	Total 1 st Stage Bids
<i>Baseline</i>	0.837	0.914	3.381	0.835	0.891	2/7	0.029	0.028
<i>Low N</i>	0.845	0.923	3.125	0.844	0.902	2/5	0.029	0.025
<i>High N</i>	0.831	0.908	3.528	0.829	0.884	2/9	0.029	0.029
<i>Low μ_C</i>	0.821	0.899	0.888	0.820	0.875	2/7	0.029	0.027
<i>High μ_C</i>	0.852	0.930	0.921	0.851	0.907	2/7	0.029	0.029
<i>Low σ_C</i>	0.877	0.940	0.905	0.873	0.900	2/7	0.029	0.013
<i>High σ_C</i>	0.797	0.890	0.907	0.800	0.882	2/7	0.029	0.042
<i>Low α</i>	0.835	0.915	3.311	0.831	0.881	2/7	0.029	0.019
<i>High α</i>	0.838	0.914	3.479	0.841	0.902	2/7	0.029	0.038
<i>Low K</i>	0.828	0.870	4.592	0.830	0.843	3/7	0.007	0.013
<i>High K</i>	0.848	0.961	2.856	0.835	0.891	2/7	0.061	0.028

Table 2.5: Breakdown of the expected values of major quantities in an FPAFE and an ERA. The cases correspond to the parameters listed in Table 2.4. The “Comp. Cost” columns report the expected completion cost, which is the winner’s cost when the project is awarded to a bidder and $c_0 = 1.5$ when there is no winner. The “Win Bid” columns are the average bid of the winning bidder in the auction for the contract (this is identically the same as the procurement cost in the FPAFE). The “Total 1st Stage Bids” column lists the sum of the expected first-round bids in the ERA. The column n_{cost}^*/N lists the procurer-optimal number of entrants selected in the ERA. Average outcomes are estimated using 500,000 simulations.

Chapter 3

Imperfect Public Monitoring with a Fear of Signal Distortion

This chapter proposes a model of signal distortion in a two-player game with imperfect public monitoring. We construct a tractable theoretical framework where each player has the opportunity to distort the true public signal and each player is uncertain about the distortion technologies available to the other player. We show that when players evaluate strategies according to their worst-case guarantees—i.e., are ambiguity averse over certain distributions in the environment—perceived continuation payoffs endogenously lie on a positively sloped line. We then provide examples showing that, counterintuitively, identifying deviators can be harmful in enforcing a strategy profile; moreover, we illustrate how the presence of such signal distortion can sustain cooperation when it is impossible in standard settings. We finally extend our model to a repeated game where our concept is a natural generalization of strongly symmetric equilibria. In this setting, we prove an anti-folk theorem, showing that payoffs under our equilibrium concept are under general conditions bounded away from efficiency.

3.1 Introduction

Many real-world strategic interactions are mediated by public signals that are possibly random functions of the players' actions, and economists have applied the theory of imperfect public monitoring to study many such situations. Applications include oligopoly games where price is influenced by quantity as well as demand fluctuations (see, for instance, Green and Porter (1984)), trade agreements with volatile trade volume (Bagwell and Staiger (1990)), and incentive contracts where workers' actions are unobserved (e.g., Radner (1986) and Levin (2003)). In most of these settings

This chapter is joint with Lucas Manuelli and Ludwig Straub. Thanks to Drew Fudenberg for encouraging us to undertake this project and for his guidance along the way. We are also grateful for helpful discussions with Gabriel Carroll, Glenn Ellison, Jonathan Libgober, David Miller, and Juuso Toikka. Bhattacharya and Manuelli acknowledge financial support from the NSF Graduate Research Fellowship under Grant No. 1122374. Bhattacharya also acknowledges support from the John Krob Castle 1963 Fellowship. Straub acknowledges financial support from the ERP Fellowship. All errors are our own.

as well as in the theoretical work on games with imperfect public monitoring (e.g., Abreu, Pearce, and Stacchetti (1990) and Fudenberg, Levine, and Maskin (1994)), it is taken for granted that the signal structure—the map from the action played to the public signal generated—is fixed and commonly known among all players. Recent papers (e.g., Fudenberg and Yamamoto (2010, 2011)) have acknowledged that this assumption is often especially strong and have proposed methods to relax it.

This chapter proposes a new method to relax this assumption, based on the observation that in many of the applications above, players may fear that the signal that mediates their interaction can be *distorted* by their opponents. In a partnership game between two workers, say, compensation may be based on various dimensions of quality of an object that the workers produce jointly: workers may worry that their colleague may sabotage or otherwise alter the object after work on the project has concluded. In other settings, the signal is determined by a third party. Again in a worker-firm setting, promotions or bonuses may depend on performance evaluations conducted by a manager; a worker may worry about favoritism between his colleague and the manager that may cause the manager to doctor her evaluation in favor of the colleague. Cartel agreements are often based on measures like market share, which are computed by a consulting firm hired by the cartel.¹ Cartel members may worry that the consulting company is in the pocket of one of the firms and may be willing to alter these numbers in favor of this firm—perhaps in return for the promise of future business with this firm.

Signal distortion could directly be modeled as simply an extended game of imperfect public monitoring—in which players have a richer action space that allows them to affect signals without affecting per-period payoffs. Instead, we take a different and novel approach to modeling signal distortion in this chapter. It is natural to believe that in many settings there is a large amount of uncertainty in how one’s opponents can distort the signal as well as in how one will be able to distort the signal oneself. As such, in the model we present, we assume that players are unsure about the timing of the distortion as well as the distortion technology itself and are *ambiguity averse*, in the maxmin sense of Gilboa and Schmeidler (1989), over the possibilities. We can even imagine a situation in which there is no actual chance to distort the technology, but agents fear this possibility anyway. As a result, incentives are determined by “perceived” payoffs given by the preferences of ambiguity-averse agents.

This setting is best described by a simple story involving a partnership game. Suppose two workers are both working on a project, and they can choose to either work hard or shirk. Each worker does not see what his colleague is doing, but their manager does see their actions and writes down performance evaluations about them. The manager will show these performance evaluations to her boss the following day to determine compensation for the workers. So far, the setting has exactly been one of imperfect public monitoring: the decision of whether to work or shirk can be thought of as some game between the workers, the performance evaluation represents the signal

¹One problem such third parties solve in a cartel, for instance, is that firms may be unwilling to share their books with competitors but may be willing to do so with a third party. Section 6.6 of Marshall and Marx (2012) discusses consulting firms and trade associations as potential third-party facilitators and provides many examples.

generated by the actions, and the compensation represents the continuation payoffs. Now suppose that that evening, the workers can individually approach the manager and try to convince her to change the evaluation she will show her boss. When approaching the manager, each worker is unsure about (i) how the manager would be willing to modify the evaluation and about (ii) whether his colleague will be able to approach the manager later that evening. Given this uncertainty over the “distortion technology,” we can imagine that when a worker approaches her, the manager simply offers him a take-it-or-leave-it offer to change the evaluation to something else in particular: because of ambiguity aversion over this distortion technology, worker 1 (say) would fear that the manager’s offers to both workers would be especially undesirable to worker 1. We assume that the manager has no stake in this game and that there is no cost to convincing her to change the evaluation.²

The first main result of this chapter is that the combination of the possibility of one’s opponent distorting the signal and one being uncertain about and ambiguity averse over how one’s opponent may modify the signal forces perceived continuation payoffs to lie on a positively sloped line. That is, the incentives provided by continuation payoffs are *perfectly aligned* across players, even if the true continuation payoffs are misaligned. This linearity result is reminiscent of a number of results in the literature on principal-agent models, such as Holmström and Milgrom (1987), Edmans and Gabaix (2011), Chassang (2013), and Carroll (2015). The intuition behind the linearity result in our model is most similar to that in Carroll (2015). Just as the ambiguity-averse principal in Carroll’s model evaluates contracts according to worst-case guarantee over potential actions the agent can take, so do the players in our model evaluate continuation payoffs based on the worst-case over options that their opponents may be presented with. As such, they hesitate to distort the signal to values that may harm their opponents, since they fear that their opponents would then be more willing to distort the signal to potentially worse outcomes. Consequently, the perceived payoffs are endogenously aligned. To our knowledge, ours is the first analysis studying this sort of endogenous linearity in a game theory context rather than in contract design.

Our second main observation is that the equilibrium concept used in this chapter can indeed have starkly different consequences from standard concepts studied in the literature. We present examples that suggest that, contrary to the intuition from standard games of imperfect public monitoring, making deviations by players *less* distinguishable can actually aid cooperation. Moreover, we suggest a reinterpretation of standard normal-form games and show that introducing a fear of signal distortion can *help* support Pareto-efficient outcomes, which would not be sustainable without this possibility of distortion. Finally, we extend our equilibrium concept to an infinitely repeated setting, where we show that it naturally generalizes the concept of a strongly symmetric equilibrium. Here, we prove an anti-folk theorem, namely that the set of possible equilibrium payoffs under our equilibrium concept is bounded away from efficiency.

²Altering the evaluation would be costless in a setting where the manager may have personal relationships with the workers and the workers fear favoritism. We might also think of this setting as one in which the cost of the bribing the manager becomes negligible compared to the change in the continuation payoffs, although we do not formalize such an interpretation in this chapter.

The ingredient we use to generate the linearity result—that players are ambiguity averse over distributions in the environment—is nonstandard, but it is one in which there has been much recent interest in the context of mechanism design. Bose, Ozdenoren, and Pape (2006) study an auction in which there is uncertainty over the value distribution of the players, and Bodoh-Creed (2012) extends these results to more general ambiguity-averse preferences. Lopomo, Rigotti, and Shannon (2010, 2011) consider principal-agent models (and more general mechanisms) in which Knightian uncertainty is embedded through incomplete preferences. Bose and Renou (2014) propose a setting in which the mechanism designer is allowed to engineer ambiguity into an environment. Like these papers, we also modify a standard setting by introducing ambiguity to a particular feature of the environment and study differences induced by this ambiguity. We also relate to the literature on games with ambiguity-averse players, including Dow and da Costa Werlang (1994), Klibanoff (1996), and Lo (1996, 1999).

The remainder of the chapter proceeds as follows. In Section 3.2, we present a one-period model that formalizes our setup, explains the preferences we posit, and introduces the concept of *distortion equilibrium*—our key equilibrium concept in the chapter. Section 3.3 proves the key result of the linearity of incentives and provides further discussion of the equilibrium concept and its robustness. We then provide some examples to motivate the intuition presented above—and draw distinctions between distortion equilibria and standard Nash equilibria—in Section 3.4. Section 3.5 extends the model to an infinitely repeated game and formulates a variation of public perfect equilibrium that is consistent with the one-period model presented beforehand. Section 3.6 modifies the results from Abreu, Pearce, and Stacchetti (1990), Fudenberg and Levine (1994), and Fudenberg, Levine, and Maskin (1994) (henceforth APS, FL, and FLM, respectively) to provide a method to compute the set of perceived payoffs in the infinitely repeated game. We show that this set is closely connected with the set of payoffs of what we call *totally linear* public perfect equilibria—or PPEs in which all continuation payoffs must lie on a positively sloped line at all public histories—under standard preferences. We also present examples of repeated games to reinforce the intuition from the one-period model and also discuss the anti-folk theorem in this context. Section 3.7 concludes. The Appendix collects a formal derivation of our equilibrium concept (Appendix C.1), an extension of our model to N players (Appendix C.2), and all omitted proofs (Appendix C.3).

3.2 A One-Shot Game

3.2.1 Setup

Consider a two-player normal form game G where player i has a finite action set A_i . Payoffs are denoted by $g : A_1 \times A_2 \rightarrow \mathbb{R}^2$. An action profile $a = (a_1, a_2)$ also generates a public signal $y \in Y$ from a known distribution $\pi(a) \in \Delta Y$. We assume for simplicity that Y is finite. The realization y of the public signal will give player i an additional payoff $w_i(y)$. While we will assume that $w_i(y)$ is exogenously fixed, it can easily be thought of as the result of future strategic choices of the players; for instance, it can specify the equilibrium of a second-stage game that players coordinate on, or

it can specify the future path of play in a repeated game of imperfect public monitoring.³ We will thus refer to $w_i(y)$ as the *continuation payoff*. A (mixed) strategy for player i in this game is simply a distribution $\alpha_i \in \Delta A_i$, and the payoff to player i in this game from the mixed strategy profile $\alpha = (\alpha_1, \alpha_2)$ is

$$\mathbb{E}_{\alpha_1[a_1], \alpha_2[a_2]} [(1 - \delta)g_i(a_1, a_2) + \delta \cdot \mathbb{E}_{\pi(a_1, a_2)[y]} w_i(y)],^4$$

where $\delta \in [0, 1)$ is the discount factor. In what follows, we will use $g(\alpha)$ and $\pi(\alpha)$ as payoffs and distributions from mixed strategies in the obvious way.

In our setting, however, players will (possibly) be given an opportunity to modify the signal before it is publicly revealed, thereby giving them the ability to modify the realized continuation payoffs. However, there is uncertainty as to whether one's opponent will be able to modify the signal after a player does so himself. Figure 3-1 displays the extensive form game that models this uncertainty. Players play actions α , and a signal \hat{y} is drawn from $\pi(\alpha)$. However, the signal \hat{y} is not immediately shown to the players; rather, there is a possibility that this signal \hat{y} is distorted, and as such, we refer to this signal as a "temporary signal." After the temporary public signal \hat{y} is drawn, one of three things happens.

Case I: The temporary signal \hat{y} is shown to the players and becomes the signal y . This happens with probability $1 - \xi$.

Case II: Nature draws a signal distribution $\mu \sim F(\hat{y})$ for some distribution $F(\hat{y}) \in \Delta(\Delta Y)$. Note that this distribution can depend on the realization of the temporary signal \hat{y} . Player 1 is then given the choice between keeping \hat{y} or accepting the alternate signal distribution μ .

- If Player 1 decides to keep \hat{y} , then Nature draws another signal distribution $\mu' \sim F(a_2, \hat{y})$ (which can depend on the temporary signal \hat{y} and the action Player 2 played) and offers Player 2 the choice between \hat{y} and μ' . If Player 2 decides on \hat{y} , then the true signal y becomes \hat{y} ; otherwise, the true signal y is drawn from μ' .
- On the other hand, if Player 1 decides to change to μ , then a second temporary signal $\hat{\hat{y}}$ is drawn from μ , and a distribution μ' is drawn from $F(a_2, \hat{\hat{y}})$ and presented to Player 2.

This happens with probability $\gamma\xi$.

Case III: The same setup can happen as Case II, but the roles of Players 1 and 2 can be switched. This happens with probability $(1 - \gamma)\xi$.

Note that following the initial normal-form game G , nodes consist of tuples $(a_1, a_2, O, \hat{y}, \mu)$ in the middle row of Figure 3-1 and tuples $(a_1, a_2, O, \hat{y}, \mu, YN, \hat{\hat{y}}, \mu')$ in the bottom row of Figure 3-1.

³See Sections 3.5 and 3.6 for an application to repeated games.

⁴Throughout this chapter, we will use the notation $\mathbb{E}_{\pi[y]} f(y)$ to denote the expectation of $f(y)$ where y is distributed according to π . That is, the dummy variable will be listed in square brackets in the expectation to make expressions easier to follow.

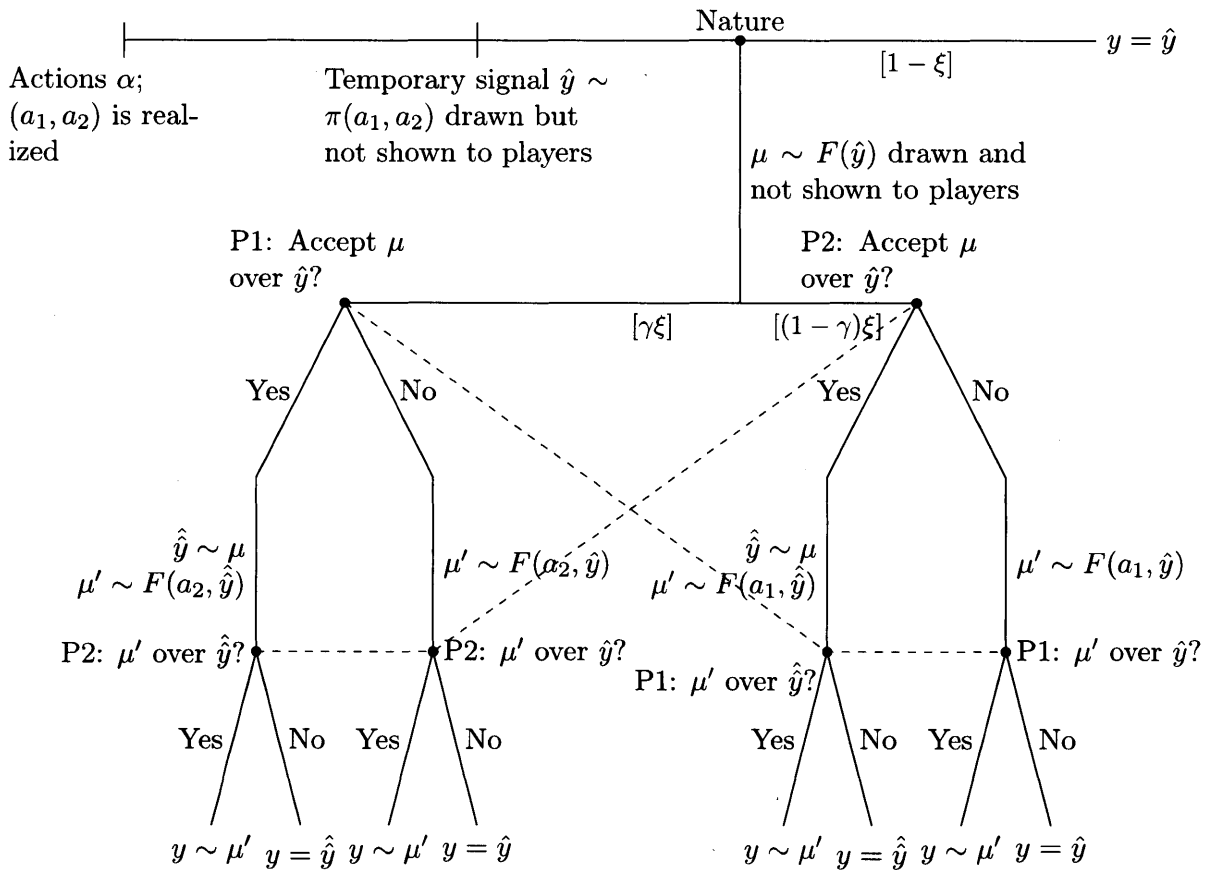


Figure 3-1: An extensive form representation of the first-stage game. The tree begins in the top left. Probabilities for the various moves of Nature are given in square brackets, and information nodes for the two players are connected by dashed lines. At the leaves of the tree, the signal y is announced as the public signal.

Here, (a_1, a_2) is the strategy realized in G , $O \in \{1, 2\}$ is the order of the players (i.e., which player modifies the signal first), \hat{y} is the temporary signal shown to the first-moving player, μ is the alternate signal distribution shown to the first-moving player, $YN \in \{\text{Yes}, \text{No}\}$ is the decision of the first-moving player, \hat{y}' is the temporary signal shown to the second-moving player, and μ' is the alternate signal distribution shown to the second-moving player. Player i 's information set, however, is simply $h_i = (a_i, \hat{y}, \mu)$, as he does not know the order in which the players can modify the signal (i.e., whether his opponent will have a chance to modify the signal from what he chooses), nor does he know the action his opponent played.

Of course, the nonstandard aspect of this setup is the distortion phase (Cases II and III above), and it is worth briefly discussing the interpretation of $F(\cdot)$ and the alternate signal distributions μ drawn from this distribution. To do so, it helps to revisit the story from the Introduction. In this situation, we can interpret the signal y as a subjective evaluation by the manager of the agents, which is dependent on the true actions taken by the agents. The distribution $F(\cdot)$ can be interpreted as the distortion technology available to each player—or perhaps jointly to the pair consisting of a particular player and the manager together. For instance, it may be easier to change the signal y to a signal y' that is “close” in some unmodeled sense: perhaps the manager will not wish to alter her subjective evaluation too drastically. The realization μ from this distribution is the actual distribution to which the manager (or the player) is able to change the signal when the player visits her.⁵

3.2.2 Distortion Equilibrium

We now define main equilibrium concept in our chapter under the assumptions alluded to in the Introduction. We assume that the economic agents are uncertain about *when* and *how* they will be able to distort the signal, and we model this uncertainty as ambiguity aversion over ξ , γ , and the F 's; that is, agents know that ξ , γ , and F each belong to some set, but they are not sure about the actual values. More specifically, agents are risk-neutral with respect to probability distributions they know (i.e., outcomes of mixed strategies and the signal distribution) but are ambiguity averse over distributions they do not know.⁶ We assume that agents have no information about ξ and γ , other than the trivial bounds that ξ and γ both lie in $[0, 1]$, and that agents believe each F is an element of $\Delta(\Delta(Y))^0$, or the set of distributions over ΔY that have full support.

Appendix C.1 derives the equilibrium concept in this section by starting with a sequential

⁵In our current setup, we have assumed that each player is offered a single alternative signal distribution to which he can choose to distort the temporary signal. We could imagine a more general model in which the players are allowed to choose from a *set* of possible distortion strategies instead of a single alternate signal. It is easy to see that such a model would not be operationally any different from the one proposed in this chapter. As long as there is maximal ambiguity aversion over this set—i.e., players worry that any possible set of distributions could be offered to their opponents—it would be without loss of generality for each player to worry about singleton sets. Any non-singleton set could be replaced by the singleton set containing a (possibly unique) choice that a player would be willing to make from the larger set.

⁶We can derive the same equilibrium concept under the assumption that agents are infinitely risk-averse *only* with respect to the outcomes of the distributions governed by γ and F . Furthermore, by interpreting the payoffs as von Neumann-Morgenstern utilities this “risk neutrality” over all other distributions includes the cases of risk-averse and risk-seeking behavior.

equilibrium of the game presented in Section 3.2.1 and formally introducing ambiguity aversion in ξ , γ , and each F . Since the notation in that appendix is involved, we focus on the intuition here; footnote 7 connects the discussion in this section to that in the Appendix. When given the opportunity to distort the signal, the agent is ambiguity averse over the order in which the players are approached (γ) as well as the technology that will be available to their opponents if their opponents do indeed have a chance to re-distort the signal ($F(a_2, y)$). Once faced with a temporary signal, the worst possible situation for player 1 is that his opponent will have a chance to alter the decision, and that the technology available to player 2 will be such that he chooses the worst possible outcome for player 1. Note, however, that we do *not* exogenously assume that player 2's incentives are inherently misaligned from those of player 1: player 1 still only believes that player 2 will choose distortions that make player 2 better off (in player 2's own eyes, keeping in mind that player 2 also believes that he is distorting first). Therefore, when choosing actions in the first stage, player 1 effectively believes that (i) there will be a distortion stage, (ii) the alternate distribution he is offered will not allow him to improve the signal, and (iii) his opponent will have the final say in distorting the signal. It is important to note that, unlike beliefs about γ , both players believe that $\xi = 1$. There would be no difference in the model if ξ were therefore known to be 1. However, allowing for ambiguity aversion in ξ allows for the interpretation that the distortion phase is a fictitious construct that nevertheless affects incentives and alters the actions that can be supported in equilibrium.

We will describe the strategies in our equilibrium concept with ambiguity aversion with two objects: (i) an action in the game G followed by (ii) a *distortion strategy* for each player i , which we will denote $D_i : Y \rightrightarrows \Delta Y$. This distortion strategy must be a compact-valued correspondence with $\delta_y \in D_i(y)$ for each $y \in Y$. $D_i(y)$ represents the set of distributions that player i is willing to accept in lieu of the temporary signal when the temporary signal is y .

Let \mathcal{D} denote the set of all possible distortion strategies for either player, i.e., the set of all compact-valued correspondences D_i with $\delta_y \in D_i(y)$. Define the *perceived continuation value* from the signal y as

$$\tilde{w}_i(y) \equiv \min_{\mu \in D_{-i}(y)} \mathbb{E}_{\mu[y']} w_i(y'). \quad (3.1)$$

The interpretation is that player i perceives the continuation payoff from the signal y to be the one generated from the *worst possible distribution* that his opponent would prefer over y . Adopting this definition of perceived continuation payoffs, the time-zero utility of player i is defined to be

$$v_i(\alpha) \equiv (1 - \delta)g_i(\alpha) + \delta \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_i(y), \quad (3.2)$$

with $\delta \in (0, 1)$ the discount factor. To introduce the equilibrium concept, we define a notion of *consistency* between the perceived continuation payoffs in (3.1) and the distortion strategy. This condition essentially requires that the distortion strategy be optimal given the perceived continuation payoffs, i.e., that $D_i(y)$ is the set of distributions to which player i prefers to distort y . It is the analogue of a sequential rationality condition that would be imposed in a sequential

equilibrium of the extensive-form game described in Section 3.2.1.⁷

Definition 3.1 (Consistency). *A triple (w, \tilde{w}, D) is said to be consistent if \tilde{w} satisfies (3.1) and*

$$D_i(y) = \{\mu \in \Delta(Y) : \mathbb{E}_{\mu|y'}[\tilde{w}_i(y')] \geq \tilde{w}_i(y)\}. \quad (3.3)$$

Note that if (w, \tilde{w}, D) is consistent, then $D_i(y)$ is clearly compact and contains the degenerate distribution δ_y , so the restriction that $D_i \in \mathcal{D}$ is implied by consistency. Implicit in (3.3) is a tie-breaking assumption that player i is willing to distort to any signal distribution that leaves him *weakly* better off.

Finally we define our equilibrium concept as follows.

Definition 3.2 (Distortion Equilibrium). *A strategy profile (α, D) is a distortion equilibrium given continuation payoffs w if*

- (i) *defining $\tilde{w}(y)$ via (3.1), the triple (w, \tilde{w}, D) is consistent as in Definition 3.1; and*
- (ii) *for each player, α_i is optimal given α_{-i} and \tilde{w}_i , meaning for all $a_i \in \alpha_i$,*

$$a_i \in \arg \max_{a'_i \in A_i} (1 - \delta)g_i(a'_i, \alpha_{-i}) + \delta \cdot \mathbb{E}_{\pi_y(a'_i, \alpha_{-i})|y} \tilde{w}_i(y).$$

The main operational difference between the standard imperfect public monitoring setup and this alternate setup is that incentives in this setup are given by the *perceived* continuation values $\tilde{w}_i(y)$ instead of the standard continuation values $w_i(y)$.

We note that our notion of distortion equilibrium is closely related to that of a *multiple priors equilibrium* from Lo (1999). This equilibrium concept consists of a *set* of beliefs for each player, which collects all the beliefs over opponents' actions that the player "can imagine happening" and over which he is ambiguity averse. It can be shown that any distortion equilibrium is equivalent to a multiple priors equilibrium applied to the extensive game described by Figure 3-1 in which players are not ambiguity averse about each others' actions but only about some of nature's draws.

3.3 Properties of Distortion Equilibria

3.3.1 Linearity

Fix continuation payoffs $w(y)$. The consistency requirement for the D can be viewed as a fixed point problem between the distortion strategies and the perceived continuation payoffs $\tilde{w}(y)$. Given the $w(y)$ and a D , we can compute $\tilde{w}(y)$ as the solution to the minimization problem given in (3.1); given the $\tilde{w}(y)$, the $D(y)$ are uniquely defined from (3.3) in the definition of consistency. This fixed point procedure gives us the following key result.

⁷ See Appendix C.1 for details. Note that (3.1) is exactly analogous to $\tilde{U}_i(\delta_y|h)$ in (C.11) in Appendix C.1, and (3.2) is the analogue of (C.12). Consistency is exactly the formal counterpart of the way we construct the $D_i(y)$ sets in the sequential equilibrium, in (C.6).

Theorem 3.1. *Consistency of the triple (w, \tilde{w}, D) requires that $\tilde{w}(y)$ lie on a line of slope in $(0, \infty)$, when plotting the pairs $\{(\tilde{w}_1(y), \tilde{w}_2(y))\}_{y \in Y}$.*

This proof is a consequence of two simple lemmas.

Lemma 3.1. *Suppose $\tilde{w}(y)$ and D satisfy consistency. Then, the $\tilde{w}(y)$ are strongly Pareto-ranked in that $\tilde{w}_1(y') \geq \tilde{w}_1(y'')$ if and only if $\tilde{w}_2(y') \geq \tilde{w}_2(y'')$.*

Proof. Suppose y and y' are such that $\tilde{w}_1(y) \leq \tilde{w}_1(y')$. Then, $D_1(y') \subseteq D_1(y)$. This implies that

$$\tilde{w}_2(y) = \min_{\mu \in D_1(y)} \mathbb{E}_{\mu[y'']} w_2(y'') \leq \min_{\mu \in D_1(y')} \mathbb{E}_{\mu[y'']} w_2(y'') = \tilde{w}_2(y'),$$

as needed. \square

Lemma 3.2. *Suppose $\tilde{w}(y)$ and D satisfy consistency and, among the $\tilde{w}(y)$, we have a unique Pareto-best point \tilde{w}_B such that $\tilde{w}_B \geq \tilde{w}(y)$ for all y and a unique Pareto-worst point \tilde{w}_W such that $\tilde{w}_W \leq \tilde{w}(y)$ for all y .⁸ Then, the $\tilde{w}(y)$ lie on a line with slope in $(0, \infty)$.*

Note, of course, that the fact that there is a unique Pareto-best point and a unique Pareto-worst point is implied by Lemma 3.1. Moreover, there may be multiple signals that give rise to these Pareto-extremal perceived continuation payoffs.

Proof of Lemma 3.2. First, if $\tilde{w}_{B,2} = \tilde{w}_{W,2}$ then Lemma 3.1 ensures that $\tilde{w}_B = \tilde{w}_W$, meaning $\tilde{w}_B = \tilde{w}_W = \tilde{w}(y)$ for all y . Trivially, all the $\tilde{w}(y)$ lie on a positively sloped line since they all coincide.

It remains to consider the case where $\tilde{w}_W < \tilde{w}_B$. Let $Y_B \equiv \{y \in Y : \tilde{w}(y) = \tilde{w}_B\}$ and $Y_W \equiv \{y \in Y : \tilde{w}(y) = \tilde{w}_W\}$. Let line ℓ connect \tilde{w}_B to \tilde{w}_W , and suppose for contradiction that there exists \hat{y} such that $\tilde{w}(\hat{y})$ lies above line ℓ .⁹ First note that there exist y_B^* and y_W^* such that $w_2(y_B^*) = \tilde{w}_{B,2}$ and $w_2(y_W^*) = \tilde{w}_{W,2}$.¹⁰ Note that $\tilde{w}_1(\hat{y}) = \alpha \tilde{w}_1(y_B^*) + (1 - \alpha) \tilde{w}_1(y_W^*)$ for some $\alpha \in [0, 1]$. It follows that $\alpha \delta_{y_B^*} + (1 - \alpha) \delta_{y_W^*} \in D_1(\hat{y})$. But, $\tilde{w}_2(\hat{y}) > \alpha \tilde{w}_2(y_B^*) + (1 - \alpha) \tilde{w}_2(y_W^*) = \alpha w_2(y_B^*) + (1 - \alpha) w_2(y_W^*)$, which is a contradiction to the definition of $\tilde{w}_2(\hat{y})$ as $\min_{\mu \in D_1(\hat{y})} \mathbb{E}_{\mu[y]} w_2(y)$. \square

It is important to stress that Theorem 3.1 is purely a result of the consistency requirement and is entirely independent of incentive compatibility in the actions chosen in G . Nevertheless, this theorem has strong implications for which strategies can be supported in an equilibrium. To explain this statement, note that the two ways to provide incentives that most theories offer are (i) value-burning and (ii) orthogonal enforcement.¹¹ The folk theorem in FLM relies crucially on the second method of providing incentives, as transfers of continuation values along the tangent hyperplanes

⁸When comparing vectors, we use $u \geq v$ to mean that $u_i \geq v_i$ for all components i .

⁹We are viewing the $\tilde{w}(y)$ as plotted on the $(\tilde{w}_1, \tilde{w}_2)$ plane. Also, the case where $\tilde{w}(\hat{y})$ lies below ℓ is symmetric; simply exchange the roles of players 1 and 2.

¹⁰This is because $D_1(y_B) = \Delta Y_B$ for all $y_B \in Y_B$ and $D_1(y_W) = \Delta Y$ for all $y_W \in Y$, which implies that $\tilde{w}_{2,B} = \min\{w_2(y_B) : y_B \in Y_B\}$ and $\tilde{w}_{2,W} = \min\{w_2(y) : y \in Y\}$. Take the arg mins to find the desired signals y_B^* and y_W^* .

¹¹See Sannikov and Skrzypacz (2010) for a discussion.

are used to punish players if the public history is suggestive of a deviation from the prescribed strategy. In contrast, since the $\tilde{w}(y)$ are the values that give incentives, Theorem 3.1 implies that the incentives of the two players are perfectly aligned; that is, the only incentive provision mechanism is value burning. We can view the incentive provision mechanism in this setting as a generalization of the strongly symmetric equilibria of Green and Porter (1984) or Abreu, Pearce, and Stacchetti (1986): conditional on a choice of slope for the perceived continuation values, the incentives of the two players cannot be misaligned. However, unlike in strongly symmetric equilibria, the slope and the intercept between the incentives is a choice variable; this flexibility is relevant in Section 3.5 when extending this setting to an infinitely repeated game.

3.3.2 Existence and Multiplicity

It is trivial to see that distortion equilibria always exist. Simply let α be a Nash equilibrium of the original game G , and let $D_i(y) = \Delta Y$ for all i and $y \in Y$. Then, $\tilde{w}(y) = \bigwedge \{w(y')\}_{y' \in Y}$ for all y ,¹² and the continuation payoffs simply serve to shift the payoffs of the original game by a constant.

In general, however, distortion equilibria are not unique. Indeed, there are often infinitely many distortion equilibria for a particular game, as there are infinitely many consistent triples for a particular set of $w(y)$. Figure 3-2 displays this multiplicity. We let Y be a three-element set and fix $w(y)$ exogenously so that $w(y_1) = (0, 0)$, $w(y_2) = (2, 1)$, and $w(y_3) = (1, 2)$. The points $(w_1(y), w_2(y))$ are plotted as blue circles. Each panel of Figure 3-2 then shows a different set $\{(\tilde{w}_1(y), \tilde{w}_2(y))\}_y$, which can be rationalized with a different $D(\cdot)$. The first panel shows the case that we used to discuss existence: all $D_i(y) = \Delta Y$ and the perceived continuation payoffs collapse. The second panel illustrates the case where $D_i(y_1) = \Delta Y$ for both i , and $D_i(y) = \Delta(\{y_2, y_3\})$ for $y \in \{y_2, y_3\}$; this causes the “top two” continuation payoffs to collapse while there remains some distinction between different signals. In the third panel, $D_i(y_1) = \Delta Y$, $D_i(y_2) = \delta_{y_2}$, and $D_i(y_3) = \{(p_1, p_2, p_3) : 2p_2 + p_3 \geq 1\}$. This is the set of distributions such that the expected value of \tilde{w}_i under that distribution is larger than $\tilde{w}_i(y_3)$. Finally, the fourth panel illustrates a consistent triple such that $D_i(y_1)$ and $D_i(y_2)$ are as before but the distortion set for y_3 is larger: in this case, $\tilde{w}(y_3) = (1/2, 1/4)$ and thus $D_i(y_3) = \{(p_1, p_2, p_3) : 2p_2 + p_3/2 \geq 1/2\}$.¹³ It can easily be checked that these choices of D are also compatible with the definition of \tilde{w} in (3.1).

The third and fourth panels of Figure 3-2 show that once the positively sloped line on which the \tilde{w} lie is chosen to be the one connecting $w(y_1)$ to $w(y_2)$, the location of $\tilde{w}(y_3)$ on this line is still indeterminate and can be picked as any point such that $\tilde{w}_i(y_3) \leq w_i(y_3)$ for both i . Note that if we were to posit $\tilde{w}(y_3) = (3/2, 3/4)$ so that $\tilde{w}_1(y_3) \geq w_1(y_3)$, then we would still be able to find D

¹²The meet of two ordered pairs x and y , denoted $x \wedge y$, is the componentwise minimum. Denote the meet of a set S as $S \equiv \{x \wedge y : x, y \in S\}$. Finally, denote by $\bigwedge S$ the element x where $x_i = \min\{s_i : s \in S\}$; that is, $\bigwedge S$ is a tuple whose i^{th} coordinate is the minimum of the i^{th} coordinates of all elements in S .

¹³Note that in any distortion equilibrium $D_1(y) = D_2(y)$ for all y . This is because we can write $\tilde{w}_1(y) = c\tilde{w}_2(y) + d$ for some $c \in (0, \infty)$ and d . Then,

$$D_1(y) = \{\mu : \mathbb{E}_\mu[\tilde{w}_1(y')] \geq \tilde{w}_1(y)\} = \{\mu : \mathbb{E}_\mu[c\tilde{w}_2(y')] + d \geq c\tilde{w}_2(y) + d\} = \{\mu : \mathbb{E}_\mu[\tilde{w}_2(y')] \geq \tilde{w}_2(y)\} = D_2(y).$$




Figure 3-2: Illustration of consistent triples. The blue circles represent $w(y)$, and the red \times represent corresponding $\tilde{w}(y)$ for various choices of $D_i(y)$. The $w(y)$ are mapped to the $\tilde{w}(y)$ via the dashed arrows.

such that the pair (\tilde{w}, D) still satisfies (3.3) from the definition of consistency. However, the triple (w, \tilde{w}, D) would not be consistent, as \tilde{w} would not be defined from (3.1); indeed, since $\delta_y \in D_{-i}(y)$ for all i , (3.1) requires that $\tilde{w}_i(y) \leq w_i(y)$. Perceived continuation values are necessarily more “pessimistic” than the actual continuation values.

The discussion suggests a very simple method for constructing consistent triples from exogenous $w(y)$. First, it is clear that the Pareto-worst \tilde{w}_W will be $\bigwedge \{w(y)\}_{y \in Y}$. Next, we choose the Pareto-best \tilde{w} . To do so, we find a set $Y' \subseteq Y$ such that no element of $\{w(y)\}_{y \in Y'}$ is Pareto-dominated by some element of $\{w(y)\}_{y \in Y \setminus Y'}$. By setting $D_i(y') = \Delta Y'$ for each i and $y' \in Y'$, we choose the Pareto-best \tilde{w}_B as $\bigwedge \{w(y')\}_{y' \in Y'}$. In the second panel of Figure 3-2, $Y' = \{y_2, y_3\}$ while in the third and fourth panels we have $Y' = \{y_2\}$. The consistency requirement affords a lot of flexibility in choosing $\tilde{w}(y)$ for $y \notin Y'$ such that $w(y) \geq \tilde{w}_W$ (for both components). *Any* point on the line connecting \tilde{w}_W to \tilde{w}_B such that $\tilde{w}_i(y) \leq w_i(y)$ for both i is a valid choice if the line on which the \tilde{w} lie has slope in $(0, \infty)$.

The last sentence in the previous paragraph is not immediately obvious, and we now offer an explanation. Suppose we have a set of $\{\tilde{w}(y')\}$ that satisfy $\tilde{w}_i(y') \leq w_i(y')$ for each i and y' and also all lie on the line connecting \tilde{w}_W to \tilde{w}_B . The choice of $\{\tilde{w}(y')\}$ pins down the $D_i(y')$ for each i and y' , and thus we are interested in showing that the conjectured (w, \tilde{w}, D) triple is indeed consistent, i.e., $\tilde{w}_i(y') = \min_{\mu \in D_{-i}(y')} w_i(\mu)$ for each i and y' . Since we assumed that the line connecting \tilde{w}_W and \tilde{w}_B has slope in $(0, \infty)$, we have $D_i = D_{-i}$ by Footnote 13. Fix a signal y . Thus, for any $\mu \in D_{-i}(y) = D_i(y)$, we will have

$$\mathbb{E}_{\mu[y']} w_i(y') \geq \mathbb{E}_{\mu[y']} \tilde{w}_i(y') \geq \tilde{w}_i(y),$$

since $D_i(y)$ is defined to be the set of μ such that the second inequality holds. Thus, it suffices to find a $\mu \in D_i(y)$ such that $\mathbb{E}_{\mu[y']} w_i(y') = \tilde{w}_i(y)$. As in Lemma 3.2, there must exist y_W and y_B such that $w_i(y_W) = \tilde{w}_{W,i}$ and $w_i(y_B) = \tilde{w}_{B,i}$. Furthermore, it must be that $\tilde{w}_i(y_W) \leq \tilde{w}_i(y) < \tilde{w}_i(y_B)$. Find the distribution μ^* that takes mass only on y_W and y_B such that $\mathbb{E}_{\mu^*[y']} \tilde{w}_i(y') = \tilde{w}_i(y)$. Then,

$\mu^* \in D_i(y)$. But then $\mathbb{E}_{\mu^*}[y']w_i(y') = \tilde{w}_i(y)$ as well (since \tilde{w} and w agree on the support of μ^*). Thus, the constructed (w, \tilde{w}, D) satisfy (3.1) and are a consistent triple.

3.3.3 Discussion

Theorem 3.1 shows that in this model of signal distortion with ambiguity aversion, perceived continuation payoffs—and thus incentives—are perfectly aligned between two players in a game. This result is reminiscent of a literature in contract theory that strives to explain linear incentive contracts designed by a principal for an agent. One early such paper is Holmström and Milgrom (1987), who present a model in which a contract that is linear in observable outcomes is optimal. The intuition, unlike in our model, is most closely related to the fact that such a scheme does not allow the agent to arbitrage nonlinearities in the incentive provision mechanism; it is in this sense that Holmström and Milgrom (1987) suggest that linear contracts are robust to a large strategy space. More recently, Edmans and Gabaix (2011) provides a separate set of sufficient conditions for linear contracts by considering a situation in which an agent decides on how much effort to exert after observing the noise in the system. Chassang (2013) considers a dynamic contracting environment and shows that linear contracts satisfy attractive properties, including performance bounds over many environments.

Carroll (2015) provides an explanation for linear contracts that contains many of the same ingredients that our model does. In his paper, an agent can choose a distribution of output at some cost, but the principal is ambiguity-averse over the set of distribution-cost pairs available to the agent. Linear contracts, in which the principal's payoff is tied directly to the agent's, then essentially guarantee that the principal also benefits from any self-interested action the agent takes. A similar setup and reasoning underlies our model as well. We can interpret the choice of signal distribution in our model as the analogue of the distribution of output available to the agent in Carroll (2015). There is an analogue of the cost of choosing such distributions as well: distributions that the player will not be able to choose (e.g., alternate distributions that will not be offered by the manager to the workers) may be prohibitively costly, and the distribution that player will be able to select can be thought of as costless. Each player, however, is uncertain over the distortions and costs available to his opponent. In equilibrium, even though there is no actual contract being designed by either player, payoffs are endogenously aligned. The intuition is that each player is cognizant of the fact that distorting the signal in a manner that harms his opponent will only incentivize the opponent to alter the signal further (if he has the opportunity to do so). Given that each player is uncertain about what his opponent can do, increasing the set of possibilities that the opponent is open to can only hurt the player. In this manner, each player chooses signals to effectively tie his (perceived) payoffs to his opponent's.

While the extensive form game presented in Section 3.2 is seemingly complicated, it simply encapsulates three main components: (i) ambiguity aversion about the signal distortion technology of both oneself and one's opponent, (ii) ambiguity aversion about the timing of distortion, and (iii) uncertainty about the realized order at the time of distortion. We believe that these elements

together capture a setting in which the players fear signal distortion at many stages of the game. In the first stage, while taking actions in G , (i) and (ii) together encapsulate that both players fear that their opponents will change the signal to something that happens to be unfavorable to them. The role of (iii) is to say that this fear of signal distortion exists *even at the time of distortion*, which is especially reasonable in many of the settings one may imagine. If a worker is worrying about possible favoritism between his colleague and the manager, it seems reasonable that a meeting with the manager still would not allay his worries.

One may wonder why we have restricted the model to situations where both players necessarily have an opportunity to distort. The specifics of the potential orders of distortions presented in Figure 3-1 are not especially important. It is of course possible to extend the extensive form so that there are other distortion arrangements; for example, we can also imagine that there is a possibility that only player 1 gets to distort the signal. It is straightforward to check that as long as at each distortion node the players believe there is a chance that their opponent will distort the signal after them—i.e., the fear of future signal distortion persists throughout the extensive-form game—, then the definition of consistency still remains relevant. The linearity result in Theorem 3.1 would still hold as well, as long as the fear of signal distortion persists at the time of signal distortion.

Modeling the distortion phase becomes more involved when considering extensions of this game to $N > 2$ players. Players may worry about which opponent will distort in the future, or perhaps about the order in which his opponents will distort, or even about whether his opponents will jointly distort the signal. Appendix C.2 extends the concept of distortion equilibrium to games with N players and considers such issues. We show an analogue of Theorem 3.1, directly extending the arguments in Section 3.3.1 under many reasonable models for the distortion phase in a game with three players.

Other assumptions are deliberately extreme as well but are immaterial to the intuition of the result. For instance, we have assumed that the players fear distortions to *any* signal distribution in ΔY . We can imagine instead a model in which players only fear distortions to signal distributions in some compact subset $S \subseteq \Delta Y$, which does not depend on the particular signal y being distorted. An example would be the case where S is the set of Dirac delta distributions on Y , so that agents only fear distortions to other signals (rather than signal distributions). In such a situation, perceived continuation values would still be strongly Pareto ranked, as Lemma 3.1 would still hold, but they would not necessarily lie on a positively sloped line. As the set S grows to approach ΔY , the perceived continuation values would become closer linear ones. We have also assumed maxmin ambiguity aversion, following the literature (especially Carroll (2015)). In addition, we have assumed no direct costs of distorting the signal, which may be appropriate for situations like favoritism but inapplicable to settings where significant bribes are required to persuade a third party to alter the signal. While explicit formalizations of general forms of ambiguity aversion or of costs of distortion are beyond the scope of this chapter, we do not expect the Pareto-ranked property to exist in such models. However, as these approach the limit cases of full ambiguity aversion and zero costs, it is natural that perceived continuation values would approach linear ones too.

We finally conclude with a short description of a tension in this model: there is uncertainty about the distortion technology available to the players, but players are expected utility maximizers conditional on a particular signal distribution. This tension is present in Carroll (2015) as well, and it has an especially clear interpretation in our model. In the settings we envision, there is usually an understood map between the signal—be it the evaluations of a manager, output from a team, or assessments from a third-party arbitrator—and the payoffs. The uncertainty we are modeling in signal distortion, however, may come from less quantifiable sources: a fear of favoritism, say, or the potential for sabotage.

3.4 Examples

In this section, we present examples of games of imperfect public monitoring and compute both the standard Nash equilibrium in the game as well as the distortion equilibria. The purpose of these examples is to highlight differences between these two forms of equilibria. Our first example underscores that improving the monitoring technology—so that deviations by individual players from a prescribed strategy are distinguishable—does *not* necessarily make it easier to sustain a certain first-stage action in equilibrium. Our second example notes that it is possible to sustain certain actions in a distortion equilibrium that are *not* sustainable in a Nash equilibrium.

In both examples, we investigate whether a certain action profile α is “sustainable” in the sense that it is possible to find suitable continuation values $w(y)$ in a certain (given) set W , and distortion strategies D such that (α, D) constitutes a distortion equilibrium.

3.4.1 Identifying Deviators Can Be *Harmful*

Consider the situation where G is a prisoner’s dilemma, given as

G	C	D
C	1, 1	−1, 2
D	2, −1	0, 0

There are three public signals, denoted y' , y'' , and y''' , and we are choosing w from the set W to try to support the outcome (C, C) as an equilibrium. Suppose for concreteness that we are picking w from the convex hull of feasible outcomes in the prisoner’s dilemma itself.¹⁴ The distribution $\pi(\alpha)$ can be either

¹⁴Choosing w from the set of feasible outcomes is reminiscent of w representing a true continuation payoff in a repeated game, which we explore in Section 3.5. In a different interpretation, we can imagine that the players play a second-stage game G' in which the Nash equilibria have payoffs $(1, 1)$, $(-1, 2)$, $(2, -1)$, and $(0, 0)$, and a distortion equilibrium involves coordinating on a Nash equilibrium of G' after each signal.

S1	y'	y''	y'''		S2	y'	y''	y'''
(C, C)	1/3	1/3	1/3		(C, C)	1/3	1/3	1/3
(D, C)	0	1/3	2/3	or	(D, C)	2/3	0	1/3
(C, D)	0	1/3	2/3		(C, D)	0	2/3	1/3
(D, D)	0	0	1		(D, D)	0	0	1

Under signal structure S1, cooperating yields an equal probability of each of the three signals. If either player deviates, then y''' becomes more likely. In this sense, it is possible to (statistically) see whether *some* player deviated but impossible to tell who it is. Under signal structure S2, the distribution induced by (C, C) remains the same as in S1. The difference between S1 and S2 is that in S2, it is possible to identify who deviated: if player 1 deviates, the y' is more likely, while y'' is more likely if player 2 deviates. Note that in both structures, if both players deviate, then the signal will be y''' with certainty; however, since in this example we are trying to enforce (C, C) and are only concerned with unilateral deviations, this distribution is irrelevant.

First consider whether it is possible to sustain (C, C) as the outcome of standard Nash equilibria under signal structures S1 and S2. Under S1, the incentive compatibility condition for (C, C) for player 1 is that

$$(1 - \delta) \cdot 1 + \delta \cdot \left(\frac{1}{3}w_1(y') + \frac{1}{3}w_1(y'') + \frac{1}{3}w_1(y''') \right) \geq (1 - \delta) \cdot 2 + \delta \cdot \left(\frac{1}{3}w_1(y'') + \frac{2}{3}w_1(y''') \right),$$

which simplifies to $w_1(y') - w_1(y''') \geq 3(1 - \delta)/\delta$. Similarly, we find that the incentive compatibility condition for player 2 is that $w_2(y') - w_2(y''') \geq 3(1 - \delta)/\delta$. Thus, for $\delta \geq 3/4$, one way to enforce (C, C) in the first stage of a distortion equilibrium is to set $w(y') = w(y'') = (1, 1)$ and $w(y''') = (0, 0)$. Under S2, the incentive compatibility conditions for players 1 and 2 are $w_1(y'') - w_1(y') \geq 3(1 - \delta)/\delta$ and $w_2(y') - w_2(y'') \geq 3(1 - \delta)/\delta$, respectively. In this case, with $w(y') = (-1, 2)$, $w(y'') = (2, -1)$, and $w(y''') = (1, 1)$, setting $\alpha = (C, C)$ is a Nash equilibrium with standard preferences as long as $\delta \geq 1/6$.

To compute distortion equilibria, we must specify a value $w(y)$ for each public signal y , compute a consistent triple (w, \tilde{w}, D) , and show that (C, C) can be sustained as a first-stage action in equilibrium given the perceived continuation values. Under S1, we have found $w_i(y)$ that already lie on a positively sloped line that sustain (C, C) as a Nash equilibrium. We can set $\tilde{w}_i(y) = w_i(y)$ and achieve consistency by setting $D_i(y') = D_i(y'') = \Delta\{y', y''\}$ and $D_i(y''') = \Delta Y$. However, note that the incentive compatibility conditions for (C, C) under S2 require that $\tilde{w}_1(y'') - \tilde{w}_1(y') \geq 3(1 - \delta)/\delta$ and $\tilde{w}_2(y') - \tilde{w}_2(y'') \geq 3(1 - \delta)/\delta$; these are the same as the Nash equilibrium, with w replaced by \tilde{w} . However, for any $\delta < 1$, these conditions imply that $\tilde{w}(y')$ and $\tilde{w}(y'')$ line on a *negatively* sloped line, which contradicts Theorem 3.1. Indeed, under signal structure S2, it is impossible to sustain (C, C) in the first stage of a distortion equilibrium regardless of the continuation payoffs (i.e., even if they were not restricted to be in the convex hull of the prisoner's dilemma's payoffs).

In the above examples, moving from signal structure S1 to S2 facilitates cooperation in the

standard case;¹⁵ however, it actually hinders cooperation when studying distortion equilibria in that cooperation is *impossible* under S2. The difference between the two signal structures is that S1 is such that deviations by players are not distinguishable and result in the same signal distribution: whenever *either* player deviates from the prescribed strategy, there is a high likelihood of y''' . This does not allow for continuation payoffs where specifically the deviating player is punished. On the other hand, S2 does distinguish between a deviation by player 1 and one by player 2. A standard Nash equilibrium can leverage this distinction by using continuation payoffs of the “I win/you lose” form, thereby making both players unwilling to deviate and stomach these personal losses. Indeed, to satisfy the incentive compatibility conditions for (C, C) using S2, continuation payoffs *must* be in this anti-aligned form. This is not possible for perceived continuation payoffs, as shown by Theorem 3.1.

Informally, the fear of signal distortion induces player 1 to worry that player 2 will be able to “point the finger” at him by switching the signal to one that suggests that player 1 deviated. Under S1, such finger-pointing is not possible, since there is no signal that suggests that a deviation by player 1 is any more likely than one by player 2. This is a key intuitive difference between this setup and the one in FLM, where distinguishing deviators is quite helpful in achieving enforceability. We should note, however, that the notion of “distinguishing deviators” used in this section is different from the pairwise full rank condition in FLM. Neither S1 nor S2 satisfies pairwise full rank for (C, C) since neither satisfies the pairwise-identifiability condition from FLM. However, S1 and S2 fail pairwise-identifiability for different reasons. Signal structure S1 fails since a convex combination of deviations from player 1 (i.e., playing D with probability 1 to generate the profile (D, C)) yields the same signal distribution as a convex combination of deviations from player 2 (also playing D with probability 1 to generate the profile (C, D)). Signal structure S2 does not fail because of this property: any deviation by player 1 can be statistically differentiated from any deviation by player 2.¹⁶ Rather, it fails since a linear combination of the distributions from (C, C) , (D, C) , and (C, D) equals zero. Thus, in our example, our notion of being unable to distinguish deviators entails more than simply a failure of pairwise-identifiability.¹⁷

3.4.2 Cooperating Can Be *Easier*

In this section, we again study the prisoner’s dilemma, but instead of changing the signal structure between two games, we change G itself. To motivate this section, consider a variation of the story of the two workers given in the Introduction. Two prisoners are asked to play the prisoner’s dilemma,

¹⁵We can see this informally by comparing the discount factors needed to support (C, C) in a Nash equilibrium. An alternate route would be compute something like the score in FL.

¹⁶It is impossible to find a convex combination of the distributions for (C, C) and (D, C) that equals a convex combination of the distributions from (C, C) and (C, D) . We of course ignore the one that places mass 1 on (C, C) , in which case neither player deviated.

¹⁷We should note, however, that failure of pairwise full rank for (C, C) is necessary for it to not be enforceable in a distortion equilibrium. Indeed, if a strategy profile α has pairwise full rank as in FLM, then it is always enforceable in a distortion equilibrium. This is a consequence of Lemma 5.4 in FLM: such a strategy profile is strongly enforceable (à la FLM) on all pairwise hyperplanes, including those that ensure that the continuation payoffs $w_1(y)$ and $w_2(y)$ lie on a line of strictly positive slope.




Figure 3-3: Using a proxy in the prisoner's dilemma.

but instead of reporting their decisions to a judge who immediately carries out the punishment, they report their decisions to a bailiff. The bailiff then relays the prisoners' decisions to the judge. However, the prisoners worry that there is a chance that the bailiff stops by one or both of their cells and gives them a chance to change the report he will present the judge.

To formalize this story, let G be a game with two actions for each player, C for cooperate and D for defect. Suppose that $g_i(a) = 0$ for all $a \in \{C, D\}^2 \equiv A^2$ and all i . There are four signals, denoted y_a for each $a \in A^2$. The signal structure is such that $\pi(a) = \delta_a$ for all $a \in A^2$. Moreover, $w_i(y_a) = \tilde{g}_i(a)$, where \tilde{g} represents the payoffs from the standard prisoner's dilemma (e.g., the one from Section 3.4.1).

It is easy to see that the Nash equilibrium of this game is exactly the one in the prisoner's dilemma: both players play D . Indeed, the effective game in the first stage (taking into account payoffs from both G and the continuation payoffs w) is simply the same prisoner's dilemma, rescaled by δ . However, (C, C) can be sustained as a (nontrivial) distortion equilibrium.¹⁸ This construction is illustrated in Figure 3-3. We have $D_i(y_{(C,C)}) = \delta_{y_{(C,C)}}$, $D_i(y) = \Delta Y$ for $y \in \{y_{(C,D)}, y_{(D,C)}\}$, and $D_i(y_{(D,D)}) = \{(p_{(C,C)}, p_{(C,D)}, p_{(D,C)}, p_{(D,D)}) : p_{(C,C)} - p_{(C,D)} - p_{(D,D)} \geq 0\}$, where p_a is the probability the signal distribution μ assigns to signal y_a . This induces a game with payoffs

\tilde{w}	C	D
C	1, 1	-1, -1
D	-1, -1	0, 0

so that the payoffs of the entire game is simply δ times the numbers in the above matrix. Thus, (C, C) can also be sustained as a distortion equilibrium. Note, however, that (D, D) can also be sustained.

¹⁸Every strategy profile in the first-stage game can be rationalized as a distortion equilibrium if $D_i(y) = \Delta Y$ for all i and y , which would make the continuation payoffs trivial.

The classic intuition behind the prisoner’s dilemma is of course that if player 1 knows his opponent is cooperating, then it is a best response to defect. (That defecting is always a best response does not matter for our current purposes.) Suppose, however, the player 1 fears the bailiff will allow player 2 to alter the signal $y_{(D,C)}$, were player 1 to actually defect. The discussion in Section 3.2.2 suggests that player 1 would worry that player 2 would get the last word in distorting the signal. If player 2 were offered the temporary signal $y_{(D,C)}$, then he would be in an especially desperate situation and would be willing to alter the signal to anything—including $y_{(C,D)}$. Player 1 thus wants to avoid this and does not deviate to D in a distortion equilibrium. Note that player 2 would *not* change the signal $y_{(C,C)}$ to $y_{(C,D)}$ —and player 1 realizes this—since he fears that player 1 will then be able to change the signal to something else.

The observation that introducing a “proxy”—an agent who, like the bailiff above, executes the action and can potentially be (costlessly) bribed by the players—into a game can help sustain efficient outcomes in a distortion equilibrium applies generally. In particular, consider a normal form game \tilde{G} with action spaces A_i for $i = \{1, 2\}$ and payoffs $\tilde{g}_i(a) \in \mathbb{R}$ for $i \in \{1, 2\}$ and $a \in A_1 \times A_2$. The “proxied” version of this game is a game G with actions spaces A_i such that $g_i(a) = 0$ for all i and a , a signal set $Y = A_1 \times A_2$, and a signal structure $\pi(a) = \delta_a$. The Nash equilibria of the proxied game coincide with the Nash equilibria of the original game, but the distortion equilibria are substantially different. As noted in Footnote 18, any first-stage action profile can be trivially sustained as part of a distortion equilibrium of the proxied game if the agents expect an extreme level of distortion in the second stage ($D_i(y) = \Delta Y$). However, it is possible to devise reasonable refinements of distortion equilibria under which the Pareto-efficient action profiles are the unique action profiles which can be sustained as outcomes in a distortion equilibrium. We do not pursue a rigorous formalization in the chapter, as the economic intuition is captured well by the example in this subsection.

3.5 An Infinitely Repeated Game

Thus far, we have assumed that the continuation payoffs are exogenously fixed, and this simplification allowed us to cleanly describe much of the relevant intuition, along with the main incentive alignment result in this chapter (Theorem 3.1). It is nevertheless interesting to study methods to endogenize the continuation payoffs, and embedding the one-period setup in a supergame is one such method. In this section, we will define a *recursive distortion equilibrium*, a natural generalization of the equilibrium concept presented in the one-period model to a repeated game setting by first defining the concept recursively. We will then define *public perfect equilibria with distortion*, which both bears more resemblance to public perfect equilibria from FLM and also can be interpreted as the natural equilibrium concept when players are ambiguity-averse over the possibility of signal distortion in all future periods. We conclude by proving that the two concepts are equivalent.

We use the same setup as in Section 3.2. Two players $i \in \{1, 2\}$ are playing a stage game G with action set A_i for player i and payoff matrix given by g . Moreover, there is a finite set Y of signals and a known map $\pi : \prod_i \Delta A_i \rightarrow \Delta Y$ that gives the signal structure as a function of actions

taken in a given period. Players have a common discount factor δ . A *public history* at time t is a sequence of all observed signals until time t ; we let the set of all public histories be \mathcal{Y} . We begin with a recursive definition of equilibrium that is the natural analogue to the distortion equilibrium defined in the previous section.

Definition 3.3 (RDE). *A recursive distortion equilibrium (RDE) is a triple $(\alpha(y^{t-1}), D(y^{t-1}), w(y^{t-1}))$ of a (mixed) strategy $\alpha(y^{t-1}) \in \prod_i \Delta A_i$, a distortion strategy $D(y^{t-1}) \in \mathcal{D}^2$, and continuation payoffs $w(y^{t-1})(\cdot) : Y \rightarrow T$, for some bounded set $T \subseteq \mathbb{R}^2$, such that for each public history $y^{t-1} \in \mathcal{Y}$*

(i) *if we define*

$$\tilde{w}_i(y^{t-1})(y_t) \equiv \min_{\mu \in D_{-i}(y^{t-1})(y_t)} \mathbb{E}_{\mu[y]} [w_i(y^{t-1})(y)],$$

then the triple $(w(y^{t-1}), \tilde{w}(y^{t-1}), D(y^{t-1}))$ is consistent;

(ii) *the Bellman equation*

$$w(y^{t-1})(y_t) = \max_{a_i} \{ (1 - \delta)g_i(a_i, \alpha_{-i}(y^t)) + \delta \mathbb{E}_{\pi(a_i, \alpha_{-i}(y^t))|y} [\tilde{w}_i(y^t)(y)] \},$$

where $y^t = (y^{t-1}, y_t)$, is satisfied; and

(iii) *$\alpha_i(y^{t-1})$ solves the previous maximization for player i .*

In the one-period model, note that we essentially specified an action, a distortion, and a vector of continuation payoffs at time 0 for each player. In this infinitely repeated model, we are specifying an action, a distortion, and a vector of continuation payoffs for each player *and each public history*; it is in this sense that recursive distortion equilibrium is a natural generalization of distortion equilibrium to an infinitely repeated game. The only additional constraint, of course, is that the continuation payoffs we specify at a particular history have to be compatible with those from different histories via the Bellman equation. Moreover, it is easy to see that as a result of the consistency condition, $\tilde{w}_i(y^{t-1})(\cdot)$ lie on a positively sloped line at all public histories. As such, incentives are again provided by (perceived) value-burning in this infinitely repeated game. An important distinction, however, between this refinement and equilibria in strongly symmetric strategies is that the slope and intercept of the perceived continuation payoffs are history-dependent choice variables, and this provides a method of differentially incentivizing different players.

For the purposes of comparison to FL, a more relevant solution concept would be one in which strategies bear direct resemblance to strategies in FL and FLM. We now formulate such a concept and show that it is indeed equivalent to recursive distortion equilibria. As in FL, we focus on public strategies.

Definition 3.4 (Public Strategy). *A public strategy for player i is a map $h_i : \mathcal{Y} \rightarrow \Delta A_i \times \mathcal{D}$. That is, a public strategy specifies a (possibly mixed) action $\alpha_i(y^{t-1})$ and a distortion $D_i(y^{t-1})(\cdot) \in \mathcal{D}$ for each public history y^{t-1} .*

Now we want to specify payoffs in this infinitely repeated game directly from public strategies. As illustrated by our one period example, and the recursive definition above, the “worst case” for player i occurs when player $-i$ modifies the signal *after* him. Using this insight, we can define continuation payoffs: we will prove later that the continuation payoffs so defined are compatible with the continuation payoffs defined as part of the recursive distortion equilibrium. Fix a public strategy profile as above and suppose we enter period $t + 1$ with public history y^t . Define

$$\mathcal{C}(y^t) \equiv \left\{ (\mu_1(y^{t+s})(y), \mu_2(y^{t+s})(y))_{s \geq 0} : \mu_i(y^{t+s})(y) \in D_i(y^{t+s})(y) \right\}.$$

That is, $\mathcal{C}(y^t)$ specifies a realization of an alternate distribution μ for each public history y^{t+s} and each temporary signal y . Note that if we specify that the distortion is always such that player 1 distorts first and then player 2 does, then an element $c \in \mathcal{C}$ induces a probability distribution over public histories y^s via the sequence

$$y^t \xrightarrow{\pi(\alpha(y^t))} \hat{y} \xrightarrow{\mu_1(y^t)(\hat{y})} \hat{\hat{y}} \xrightarrow{\mu_2(y^t)(\hat{\hat{y}})} y \xrightarrow{\text{current public history}} (y^t, y) \equiv y^{t+1},$$

where each arrow represents a draw from the distribution underneath the arrow. Let $\sigma_1(y^t) : \mathcal{C}(y^t) \rightarrow \Delta(\{y^{t+s}\}_{s \geq 0})$ be this map. Then let $\mathcal{S}_1(y^t) \equiv \sigma_1(y^t)$ be its image. Then, the perceived payoff for player 1 of the public strategy profile, starting from history y^t , is

$$v_1(y^t) \equiv \min_{\nu \in \mathcal{S}_1(y^t)} \left\{ (1 - \delta) \cdot g_1(\alpha(y^t)) + (1 - \delta) \cdot \mathbb{E}_\nu \left[\sum_{s=1}^{\infty} \delta^s g_1(\alpha(y^{t+s})) \right] \right\}. \quad (3.4)$$

For player 2, simply interchange the roles of players 1 and 2 in the above discussion.¹⁹ To preserve the notation from the definition of RDE, given these payoffs $v(y^t)$ from histories $y^t \equiv (y^{t-1}, y_t)$ onward, define continuation payoffs at y^{t-1} as $w(y^{t-1})(y_t) \equiv v(y^t)$.²⁰ Now, define $\tilde{w}(y^t)(\cdot)$ as in (i) in the definition of recursive distortion equilibrium (Definition 3.3). We have that the natural consistency condition is that

$$D_i(y^{t-1})(y) \equiv \left\{ \mu \in \Delta Y : \mathbb{E}_{\mu[y']} [\tilde{w}_i(y^{t-1})(y')] \geq \tilde{w}_i(y^{t-1})(y) \right\}. \quad (3.5)$$

With this notion, we can define an equilibrium.

Definition 3.5 (PPED). *A public perfect equilibrium with distortion (PPED) is a public strategy profile (α, D) such that starting at each public history y^{t-1} ,*

¹⁹While the continuation payoffs defined as above are technically treated as definitions of the preferences given public strategies, we should note that they are reasonable given the extensive form game defined in Section 3.2. There, we showed that when evaluating payoffs, players fear that γ is such that their opponent always modifies the signal after they do. This is why when evaluating the payoffs of player 1, we only consider distributions generated when player 2 modifies the signal after player 1 does in the minimization problem. Moreover, note that it is without loss of generality to only consider $\mu \in D_i$ as part of \mathcal{C} since $\delta_y \in D_i(y^{t-1})(y)$ for all public histories; as such, the possibility of the opponent rejecting the proposed signal change is built into the set.

²⁰The distinction between the $v(y^{t-1})$ and $w(y^{t-1})(\cdot)$ is pedantic, but it serves to highlight the fact that in an RDE, we view continuation payoffs as vectors (for each player) at each node, and in a PPED, we view payoffs as a single value (for each player) at each node. These notions can of course be interchanged freely.

- (i) α_i is a best response to α_{-i} and D_{-i} in the subgame following y^{t-1} , using utility $v_i(y^{t-1})$ defined in (3.4), and
- (ii) $D(y^{t-1})(\cdot)$ is consistent with the strategy profile as in (3.5).

The connection between a PPED and an RDE is relatively intuitive, as the best response and consistency conditions in the definition of PPED are exactly as in an RDE. It remains to show that the minimization problem posited in (3.4) is compatible with the payoffs defined as part of the RDE. The following result shows that it is.

Theorem 3.2. *Given an RDE $(\alpha(y^{t-1}), D(y^{t-1}), w(y^{t-1}))$, the public strategy profile $(\alpha(y^{t-1}), D(y^{t-1}))$ is a PPED. Conversely, if $(\alpha(y^{t-1}), D(y^{t-1}))$ is a PPED, then $w(y^{t-1})(y_t) \equiv v(y^{t-1}, y_t)$, defined in (3.4), satisfies the Bellman equation.*

This equivalence helps us develop the APS-style characterization in the next section.

3.6 APS-Style Characterization of Perceived Payoffs

For a given discount factor δ , let $E_P(\delta)$ be the set of perceived payoffs—i.e., defined in (3.4)—in any PPED. Our goal is to provide an APS-style characterization of $E_P(\delta)$. In this section, we show that all the main results of APS still hold in our setting, with the appropriate modifications. We also define the concept of a *totally linear PPE*, which is a generalization of strongly symmetric PPE to arbitrary supergames, and then provide an APS-style characterization of the set $E_L(\delta)$ of payoffs sustainable in totally linear PPEs. We then show our key results that the sets $E_P(\delta)$ and $E_L(\delta)$ —along with the limits of these sets as $\delta \rightarrow 1$ —are closely connected. We then use this relationship to prove an anti-folk theorem in our setting, showing that perceived payoffs in PPEDs are bounded away from efficiency as $\delta \rightarrow 1$. This relationship also allows us to compute the limit set of perceived payoffs numerically for certain games, and we present some examples that mirror the ones in Section 3.4.1. Proofs are omitted from this section and are instead relegated to Appendix C.3.

3.6.1 Recursive Characterization of Perceived Payoffs

As in APS, we view a PPED as a pair (α, w_P) that specifies an action profile α in period 0 along with perceived continuation payoffs $w_P(y)$ for all signals $y \in Y$. The enforceability condition from APS is modified so that incentives are given by an appropriate within-period minimization.

Definition 3.6 (P-enforceable). *The pair (α, v) of an action profile and a perceived payoff is P-enforceable with respect to δ and a set W of allowed continuation payoffs if there exists $D \in \mathcal{D}^2$ and a $w(y) \in W$ for each $y \in Y$ such that*

$$(i) \quad v_i = (1 - \delta)g_i(\alpha) + \delta \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_i(y),$$

$$(ii) \quad \alpha_i \in \arg \max_{\alpha'_i} (1 - \delta)g_i(\alpha'_i, \alpha_{-i}) + \delta \mathbb{E}_{\pi(\alpha'_i, \alpha_{-i})[y]} \tilde{w}_i(y), \text{ and}$$

(iii) the triple (w, \tilde{w}, D) is consistent.

Definition 3.7. For any set $W \subseteq \mathbb{R}^2$,

$$B_P(\delta, W) \equiv \{v = (1 - \delta)g(\alpha) + \delta \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_P(y) : (\alpha, v) \text{ is } P\text{-enforceable with respect to } W\}.$$

We say a set W is *P-self generating (P-SG)* if $W \subseteq B_P(\delta, W)$.

Lemma 3.3. If W is bounded and P-SG, then $W \subseteq E_P(\delta)$.

As in APS, we see that $E_P(\delta)$ is the largest fixed point of the $B_P(\delta, \cdot)$ operator.

Lemma 3.4. $E_P(\delta) = B_P(\delta, E_P(\delta))$.

Now we want to characterize the set $E_P(\delta)$. We use an approach very similar to the linear programming strategy of FL. We set up the analogous planning problem and define a set Q_P , which will be an intersection of half-planes. As in FL, under a full dimensionality²¹ condition we will have $E_P(\delta) \rightarrow Q_P$ as $\delta \rightarrow 1$. First, we introduce the relevant programming problem.

Definition 3.8. For a strategy profile α and a direction $\lambda \in \mathbb{R}^2$, define $k_P(\alpha, \lambda, \delta)$ as the value of the program

$$\begin{aligned} & \sup_{v, w(y), D} \lambda \cdot v \\ \text{s.t. } & v = (1 - \delta)g(\alpha) + \delta \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}(y) \\ & v_i = (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_i(y) \quad \forall a_i \in \alpha_i \\ & v_i \geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_i(y) \quad \forall a_i \in A_i \\ & \lambda \cdot v \geq \lambda \cdot w(y) \quad \forall y \in Y \\ & D \in \mathcal{D}^2 \text{ is such that } (w, \tilde{w}, D) \text{ is consistent.} \end{aligned} \tag{3.6}$$

Define $k_P^*(\lambda, \delta) \equiv \sup_{\alpha} k_P(\alpha, \lambda, \delta)$.

It is easy to see that $k_P(\alpha, \lambda, \delta)$ is independent of δ by the same scaling argument as in FL. We thus write $k_P^*(\lambda)$ instead of $k_P^*(\lambda, \delta)$. We define Q_P as

$$Q_P \equiv \bigcap_{\lambda \in \mathbb{R}^2} \{v \in \mathbb{R}^2 : \lambda \cdot v \leq k_P^*(\lambda)\}. \tag{3.7}$$

Theorem 3.3. For all δ , $E_P(\delta) \subseteq Q_P$, where Q_P is as in (3.7). Moreover, if Q_P has full dimension (in \mathbb{R}^2), then $E_P(\delta) \rightarrow Q_P$ as $\delta \rightarrow 1$.

3.6.2 Totally Linear PPEs

The result of Theorem 3.1 suggests that there may be a connection between PPEDs and standard PPEs with continuation payoffs that lie on a positively sloped lines. In this section we study a

²¹Throughout this section, whenever we refer to full dimension of a set like Q_P , we view it as a subset of \mathbb{R}^2 .

subset of PPEs satisfying the additional restriction that after every public history the continuation values lie on a positively sloped line. We show in the next subsection that the set of PPE payoffs of these *totally linear* PPEs is a subset of the payoffs of PPEs.

Definition 3.9 (Totally Linear PPE). *A PPE is said to be totally linear if after each public history y^t the continuation values $w(y^t, y)$ lie on a nonnegatively sloped line. $E_L(\delta)$ is the set of payoffs of totally linear PPEs, given a discount factor δ .*

Note that a symmetric totally linear PPE in a symmetric game is simply a strongly symmetric equilibrium, and in this sense, we can view totally linear PPEs as a generalization of strongly symmetric ones. A recursive characterization for totally linear PPEs follows very similarly to that in APS.

Definition 3.10 (Linearly Enforceable). *A pair (α, v) is said to be linearly enforceable with respect to (δ, W) if there exists $w(y) \in W$ lying on a positively sloped line such that*

- (i) $v = (1 - \delta)g(\alpha) + \delta \mathbb{E}_{\pi(\alpha)[y]}[w(y)]$ and
- (ii) $\alpha_i \in \arg \max_{\alpha'_i} \left\{ (1 - \delta)g(\alpha'_i, \alpha_{-i}[y]) + \delta \mathbb{E}_{\pi(\alpha'_i, \alpha_{-i})}[w(y)] \right\}$.

For any set $W \subseteq \mathbb{R}^2$, let

$$B_L(\delta, W) \equiv \{v = (1 - \delta)g(\alpha) + \delta \mathbb{E}_{\pi(\alpha)[y]}w(y) : (\alpha, v) \text{ is linearly enforceable with respect to } W\}.$$

A set $W \subseteq \mathbb{R}^2$ is said to be *linearly self-generating* if $W \subseteq B_L(\delta, W)$. Let $E_L(\delta)$ be the set of payoffs of totally linear PPEs. Then, analogously to Lemma 3.3, we have

Lemma 3.5. *If W is bounded and linearly self-generating then $W \subseteq E_L(\delta)$.*

As before, we also have $E_L(\delta) = B_L(\delta, E_L(\delta))$ so that $E_L(\delta)$ is the largest fixed point of the B_L operator. Using this APS characterization, we get a corresponding programming problem. For a strategy profile α and a direction $\lambda \in \mathbb{R}^2$, define $k_L(\alpha, \lambda, \delta)$ as the value of the program

$$\begin{aligned} & \sup_{v, w(y)} \lambda \cdot v \\ \text{s.t. } & v = (1 - \delta)g(\alpha) + \delta \mathbb{E}_{\pi(y)}w(y) \\ & v_i = (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \mathbb{E}_{\pi(y)}w_i(y) \quad \forall a_i \in \alpha_i \\ & v_i \geq (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta \mathbb{E}_{\pi(y)}w_i(y) \quad \forall a_i \in A_i \\ & \lambda \cdot v \geq \lambda \cdot w(y) \quad \forall y \in Y \\ & w(y) \text{ lies on a line with slope in } (0, \infty). \end{aligned} \tag{3.8}$$

Define $k_L^*(\lambda, \delta) \equiv \sup_{\alpha} k_L(\alpha, \lambda, \delta)$.²² As before $k_L(\alpha, \lambda, \delta)$ is independent of δ by a standard scaling

²²Programs (3.6) and (3.8) may bear resemblance to the program suggested in Fudenberg, Levine, and Takahashi (2007) (FLT), since we have an additional constraint on the continuation payoffs. Note, however, that the results in

argument. We can show as before that if Q_L is defined analogously to Q_P in Theorem 3.3, i.e.,

$$Q_L \equiv \bigcap_{\lambda \in \mathbb{R}^2} \{v \in \mathbb{R}^2 : \lambda \cdot v \leq k_L^*(\lambda)\}, \quad (3.9)$$

then the analogous result to Theorem 3.3 holds as well.

3.6.3 Connecting Payoffs of PPEDs and Totally Linear PPEs

We now show the main results of this section. Note that given the structure of program for $k_P^*(\lambda)$, it is difficult to solve explicitly for the scores k_P^* and consequently compute the set Q_P . On the other hand, the program for computing k_L is much more tractable.²³ It turns out that Q_L and Q_P are in fact connected, as are $E_L(\delta)$ and $E_P(\delta)$.

Theorem 3.4. *For any δ , $E_L(\delta) \subseteq E_P(\delta)$. Moreover, $Q_L \subseteq Q_P$.*

The proof of the inclusion $E_L(\delta) \subseteq E_P(\delta)$ proceeds by showing that any pair (α, v) that is linearly enforceable is also P-enforceable with respect to the same δ and set W . The inclusion $Q_L \subseteq Q_P$ follows from comparing the programs (3.6) and (3.8).

While one might hope that $Q_L = Q_P$, this is not true in general. Since points in Q_P must only be enforceable by \tilde{w} that are consistent with w , it is possible to punish players more harshly in a PPED than in a totally linear PPE, and as such more actions are sustainable. However, all points in Q_P which Pareto-dominate a Nash equilibrium payoff v_{NE} are contained in Q_L . This is captured in the next theorem.

Theorem 3.5. *Let $\mathcal{L}(v^*) \equiv \{v \in Q_P : v \geq v^* \text{ elementwise}\}$. If v_{NE} is the payoff of some Nash equilibrium of the stage game, then $\mathcal{L}(v_{NE}) \subseteq Q_L$.*

The proof of Theorem 3.5 is given in Appendix C.3.4 and simply involves comparing Programs (3.6) and (3.8). The usefulness of Theorem 3.5 is that it allows us to compute Q_P exactly for certain classes of games. For instance, suppose that the minmax payoff pair \underline{v} can be supported by a stage game Nash equilibrium. Then, we know that $\underline{v} \in E_P(\delta)$ for all δ and $\underline{v} \in E_L(\delta)$ for all δ . It follows that it is in both Q_P and Q_L . Now, Q_P cannot contain any payoff pair where one agent receives less than his minmax payoff, so we know that $\bigwedge Q_P = \underline{v}$, which means that $\mathcal{L}(\underline{v}) = Q_P$, and Theorem 3.5 implies that $Q_P \subseteq Q_L$. Together with Theorem 3.4, we have If the minmax payoff pair can be supported as a stage game Nash equilibrium, then $Q_P = Q_L$.

this section are not simply applications of FLT. FLT computes the limit set of payoffs from \mathcal{A}^0 -PPEs, which place constraints on the actions allowed in equilibrium after any public history. The programs in our setup place restrictions on the *continuation payoffs* that do not map to simple restrictions on actions. A difference between totally linear PPEs and strongly symmetric PPEs is important here. The method in FLT applies to strongly symmetric PPEs since the same restriction on actions can be placed after each public history; for totally linear PPEs, since the slope and intercept of the continuation values can be different at each history, there is no such restriction on actions that applies globally after each public history.

²³On a computational note, we should point out that the program for $k_L(\alpha, \lambda)$ is not linear. Instead, due to the quadratic equality constraint, it is a nonlinear, nonconvex problem. However, upon fixing the slope of the line on which the $w(y)$ lie, it becomes linear. One simple computational implementation of this program will then boil down to a one-dimensional optimization of the score as a function of the slope, where the value of the function is the solution to a linear program.

3.6.4 An Anti-Folk Theorem

One might wonder how close Q_P can be to the Pareto frontier of the set of feasible and individually rational payoffs V^* . We next show that, in fact, the Pareto frontier of Q_P is strictly inferior to the one of V^* . In other words, our PPED solution concept gives rise to an *anti-folk* theorem.

For this result, we restrict our attention to the situation where the minmax payoff pair can be supported by a stage game Nash equilibrium (see Corollary 3.6.3). This assumption is made for simplicity—the general case follows if Q_P and V^* are restricted to their respective subsets that Pareto dominate a given stage game Nash equilibrium, as in Theorem 3.5. Further, we assume that $\pi(a)$ has full support, for all actions a . While not necessary, this is important as we would like to rule out cases where an action a is enforceable with only off-equilibrium punishments. The following result holds.

Theorem 3.6 (Anti-Folk Theorem). *Assume the minmax payoff pair can be supported as a stage game Nash equilibrium, and $\pi(a)$ has full support. Then, for every point $v \in Q_P$ that cannot be supported by a stage game Nash equilibrium, there is a point $v' \in V^*$ which Pareto dominates v .*

This anti-folk theorem illustrates that our PPED solution concept leads to fundamentally different limit sets Q_P than the limit PPE set Q . Even in situations which satisfy the PPE folk theorem assumptions in FLM, the PPED limit set Q_P is bounded away from the Pareto frontier. The key intuition for this result is as follows. In order to enforce an action that is very close to the Pareto frontier, it is necessary that punishments are almost parallel to the Pareto frontier (and hence barely waste any resources on Pareto inefficient continuation values). In a PPED, however, continuation values must be aligned and lie on a *positively* sloped line, preventing them to line up with the *negatively* sloped Pareto frontier.

3.6.5 Examples of Q_P

In this section, we present examples of Q_P . We will restrict our attention to prisoner's dilemmas. Since these games have minmax payoffs that are stage Nash equilibria, we can apply Corollary 3.6.3 and simply compute Q_P by computing Q_L numerically. The main observations from this section are that (i) Q_P can consist of more than simply the Nash equilibrium payoffs, (ii) Q_P is bounded away from the efficient frontier of V^* , and (iii) changes in the signal structure that may expand the set of equilibrium payoffs under standard imperfect public monitoring may actually *decrease* the set of perceived payoffs in our setting. The second observation corresponds to Theorem 3.6, and the final one mirrors the ones presented in Section 3.4.

Consider two separate cases of the the Prisoner's Dilemma with three public signals. The payoff matrix is the same in both games, but the signal structure differs. The tables below list payoffs and $\pi(\alpha)$.

	C	D
C	1, 1	-1, 1.1
D	1.1, -1	0, 0

S3	y'	y''
(C, C)	1/2	1/2
(D, C)	1/4	3/4
(C, D)	1/4	3/4
(D, D)	0	1

S4	y'	y''
(C, C)	1/2	1/2
(D, C)	3/4	1/4
(C, D)	1/4	3/4
(D, D)	0	1

The relation between S3 and S4 mirrors that between S1 and S2. In S3, y'' is more likely as long as either player defects. In S4, y' and y'' are equally likely if both players cooperate, but if Player 1 defects, then y' is more likely, while if Player 2 defects, y'' is more likely.

Consider Game S4 first. Without ambiguity aversion, it can easily be checked that the score in the direction $\lambda = (1, 1)$ is 2, as the strategy profile (C, C) can be supported by the continuation payoffs $w(y') = (0.8, 1.2)$ and $w(y'') = (1.2, 0.8)$ for $\delta = 1/2$. In fact, despite that the full rank conditions are not satisfied, it turns out that the entire set of feasible, individually rational payoffs can be supported by a PPE. Now consider the same signal structure, but consider PPEDs. Suppose the w were in a *negatively* sloped line at the optimum in the linear program. In this case, it is easy to see that there would be no incentives for at least one player: either both \tilde{w} would collapse to a point, or the \tilde{w} would lie on a horizontal (or vertical) line. Thus, the only action profile that could be supported with w in a negatively sloped line is (D, D) . Now suppose instead that the w were on a *positively* sloped line, and suppose without loss that $w(y') \geq w(y'')$. Then, we can choose the distortions D such that $\tilde{w} = w$, so we can just work with these continuation payoffs. Suppose we are trying to support a strategy where player i plays C with probability $\alpha_i > 0$. Then, we must have

$$\begin{aligned} \alpha_2 \cdot 1 + (1 - \alpha_2) \cdot (-1) + \alpha_2 \left(\frac{1}{2}w_1(y') + \frac{1}{2}w_1(y'') \right) + (1 - \alpha_2) \left(\frac{1}{4}w_1(y') + \frac{3}{4}w_1(y'') \right) \\ = \alpha_2 \cdot 1.1 + (1 - \alpha_2) \cdot 0 + \alpha_2 \left(\frac{3}{4}w_1(y') + \frac{1}{4}w_1(y'') \right) + (1 - \alpha_2)w_1(y''), \end{aligned}$$

which means that $0.9\alpha_2 - \left(\frac{3}{4} + \frac{\alpha_2}{2}\right)(w_1(y') - w_1(y'')) = 1$, and the equal sign is replaced by \geq if $\alpha_1 = 1$. But, $w_1(y') - w_1(y'') \geq 0$ by assumption, and $0.9\alpha_2 \leq 0.9$, so this equation can never hold.²⁴ Thus, the score in the direction $\lambda = (1, 1)$ is 0, and $\{(0, 0)\}$ is the set of perceived PPE payoffs.

Now suppose that the signal structure is “aligned” as in Game S3. It can be checked that under standard preferences, the maximum score in the $\lambda = (1, 1)$ direction is 1.6, given by the profile (C, C) and the continuation payoffs $w(y') = (0.8, 0.8)$ and $w(y'') = (0.4, 0.4)$ for $\delta = 1/2$. That is, the score *decreases* relative to Game S3, as does the set of PPE payoff that can be supported as $\delta \rightarrow 1$. Next consider PPEDs. By the same argument as before, continuation payoffs (w) that lie on a negatively sloped line give no incentives (\tilde{w}), so we must restrict ourselves to choosing continuation payoffs on a positively sloped line if we want any hope of supporting anything more than the stage Nash equilibrium. But, note that the score then coincides with that in the case of

²⁴Obviously, if we assumed that $w(y'') \geq w(y')$ instead, then we would look at the decision rule for Player 2.




Figure 3-4: The sets Q (single-hatched) and Q_P (double-hatched) for a prisoner's dilemma with the aligned signal structure S3 given in the text.

regular PPEs (since we can choose the distortion strategies to set $w = \tilde{w}$) meaning that it *increases* relative to Game S4. Figure 3-4 plots Q , Q_P , and V^* for Game S4: the dotted area is the convex hull of the payoff set, the dashed area is V^* , the single-hatched area is Q , and the criss-crossed area (also included in the single-hatched area) is Q_P . Note that Q_P is also bounded away from efficiency.

3.7 Conclusion

In this chapter, we have proposed an explicit method of signal distortion that captures uncertainty over the timing and the technology of distortion. We show that if players are ambiguity averse over this timing and the distortion technology available to their opponents when they do so, then the “perceived incentives” (i.e., the perceived continuation values) lie on a positively sloped line. The intuition is relatively simple: given that a player will only change the signal to something that will benefit himself, the other player can ensure that she is made no worse off by this possibility of distortion if continuation payoffs are perfectly aligned. If players are unaware of the true probability of distortion and are ambiguity averse over this possibility, then the incentive alignment persists even as the probability becomes arbitrarily small.

We also present some examples that highlight differences between our setup and standard ones without this fear of signal distortion. First, signal structures that make deviators less distinguishable may in fact be beneficial for cooperation. Secondly, the possibility of distorting the signal can itself sustain cooperation when it is not possible in a standard setting. We use the prisoner's dilemma to illustrate this intuition, and we leave pursuing a generalization of these observations for future work.

Given that the proposed method of signal distortion is novel to the game theory literature, we see many other possibilities for future research as well. First, our emphasis in this chapter has

been on the “perceived payoffs.” While this focus can be justified by either imagining a world in which $\xi = 1$ or noting that the perceived payoffs are the ones with which the players in our model are concerned, studying properties of the case where $\xi = 0$ (but players are unaware of this fact and still ambiguity averse over ξ) would be an interesting avenue to push this model in the future. Of particular interest would be to develop a method for characterizing the equilibrium set of “real payoffs,” i.e. the discounted payoffs of the stage game for any path of play that can be sustained as a PPED. An APS-style recursive characterization would likely still apply in this case, although pursuing such a characterization is beyond the scope of this chapter. Second, we endogenize the continuation payoffs in our model through a repeated game, but there are other appealing methods to endogenize these payoffs. For instance, the $w(y)$ in the one-period model can be interpreted as payoffs set by a principal who wishes to motivate a team of agents who have a fear of signal distortion as presented in this chapter. We view this line of research, connecting back to contract design, as especially promising.

Appendix A

Appendix to Chapter 1

A.1 Further Details on DOD Profits Under Changes to the Contest Design

In this Appendix, I repeat the exercises of Sections 1.7 through 1.9 but report the effects on DOD profits (as well as Phase III DOD profits, as defined in Section 1.10). I discuss how these design counterfactuals affect DOD profits differently from social surplus. I also quantify the “Laffer” curve for R&D efforts in Appendix A.1.2.

A.1.1 Changing the Number of Competitors

I first repeat the analysis in Section 1.7 by computing the total effect of changing N_1 and \bar{N}_2 on DOD profits and Phase III DOD profits, which ignore research costs. Table A.1 supplements Table 1.7 by focusing on these measures of DOD profits. At the baseline contest, the DOD run an expected loss of about \$100,000 per contest.¹ About \$200,000 of this is due to R&D cost reimbursements (see Table 1.7), so the DOD earns about \$90,000 per contest in expected profits when ignoring research costs.

Panel (a) reports the effect of changing competition on DOD profits. The key difference between the DOD profits and social surplus is that the DOD only captures a relatively small portion of the surplus, and it thus places much more weight on saving on research efforts. For these parameters, therefore, increasing *either* measure of competition reduces the DOD’s profits, as unlike the social planner, the DOD is unable to recover the cost of research from the larger generated surplus. The patterns are analogous to the ones with social surplus: for instance, due to the low substitutability between contestants, increasing both N_1 and \bar{N}_2 *decreases* DOD surplus almost linearly. Panel (b) shows the effects on profits, ignoring research costs. Adding more competitors to Phase I only has

¹For these parameters, the expected value of a successful acquisition is about \$15 million, and the expected cost is about \$8.5 million. Thus, the surplus to the DOD for a successful acquisition is about \$1.75 million, ignoring research expenditures. Only about 5.2% of contests result in a successful acquisition, so we can reinterpret this number as saying that the DOD spends a total of about \$2 million per successful acquisition but only recovers \$1.75 million in the delivery process.

	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	-0.023	-0.094			$N_1 = 2$	-0.001	0.094		
$N_1 = 3$	-0.024	-0.134	-0.180		$N_1 = 3$	0.001	0.098	0.189	
$N_1 = 4$	-0.026	-0.135	-0.222	-0.258	$N_1 = 4$	0.003	0.100	0.195	0.287

(a) DOD profits (Baseline is -0.103 \$M)

(b) Phase III DOD profits (Baseline is 0.091 \$M)

Table A.1: Total effects of moving from a baseline of $N_1 = \bar{N}_2 = 1$ to various values of (N_1, \bar{N}_2) on (a) DOD profits and (b) Phase III DOD profits (i.e., ignoring research costs). Each entry in the table lists the *change* from the baseline value, and the baseline values are listed in the respective captions. All values are in millions of dollars.

a negligible effect, because the benefit of an additional draw is small and the reduction in research effort can be harmful for the chances of entering Phase II. Increasing both N_1 and \bar{N}_2 improves outcomes; each additional competitor adds about \$90,000 in expected surplus.

Table A.2 decomposes the total effect on DOD profits into direct and incentive effects; for comparison, it includes direct and incentive effects on social surplus, given in Table 1.8. Fixing $\bar{N}_2 = 1$, the direct effect in Phase I of increasing N_1 is negative and approximately \$100,000 per Phase I competitor. This number is slightly larger than the \$83,000 for social surplus. The direct effect of Phase II (shown in the second panel) is basically zero for DOD profits, compared to a large and positive number for social surplus. The net direct effect is negative for the DOD whereas it is (often) positive for the social planner. These differences can be traced back to the DOD capturing a small fraction of the benefit from more draws, both in Phase I (higher values of v throughout the contest) and in Phase II (a larger chance that someone succeeds and that $\max(v_i - c_i) > 0$). The incentive effects (third and fourth panels) are comparable between the DOD and the social planner; they are slightly larger for the DOD once again because it places more of a premium on the benefit of reducing research costs, which is part of what the incentive effects cover.

This decomposition suggests that the difference between the DOD-optimal level of competition and the socially optimal level of competition can be traced back primarily to differences in the *direct* effect of Phase II. Increasing entry into late stage competition has a strong effect on the final failure rate and thus a large effect on surplus, but it also comes at a large cost. The incentive effect in Phase II does little to mitigate this cost because the firms do not change their Phase II effort appreciably in response to (the effectively low level of) competition.

A.1.2 Changing η : The Procurer's Laffer Curve

As discussed at the end of Section 1.8, there is a natural Laffer curve associated with R&D outcomes from the perspective of the procurer. Low values of η yield low profits for the DOD because firms have little incentive to exert effort. High values of η also yield low profits for the DOD because the DOD is unable to capture much of the surplus from delivery but still refunds costs.

If the firm is not promised any part of the surplus (i.e., if $\eta = 0$), then the firm has no incentives

Baseline ($N_1 = 1 = \bar{N}_2 = 1$)	Social Surplus				DOD Profits			
	0.144				-0.103			
	\bar{N}_2							
Direct (Phase I)	1	2	3	4	1	2	3	4
$N_1 = 2$	-0.083	-0.083			-0.099	-0.099		
$N_1 = 3$	-0.175	-0.175	-0.175		-0.198	-0.198	-0.198	
$N_1 = 4$	-0.270	-0.270	-0.270	-0.270	-0.297	-0.297	-0.297	-0.297
Direct (Phase II)	1	2	3	4	1	2	3	4
$N_1 = 2$	–	0.215			–	0.001		
$N_1 = 3$	–	0.232	0.429		–	0.002	0.006	
$N_1 = 4$	–	0.239	0.453	0.636	–	0.003	0.009	0.018
Incentive (Phase I)	1	2	3	4	1	2	3	4
$N_1 = 2$	0.059	0.000			0.076	0.000		
$N_1 = 3$	0.153	0.045	0.000		0.174	0.058	0.000	
$N_1 = 4$	0.250	0.135	0.042	0.000	0.271	0.155	0.054	0.000
Incentive (Phase II)	1	2	3	4	1	2	3	4
$N_1 = 2$	0.000	-0.002			0.000	0.004		
$N_1 = 3$	0.000	-0.002	-0.007		0.000	0.004	0.012	
$N_1 = 4$	0.000	-0.002	-0.007	-0.012	0.000	0.004	0.012	0.022

Table A.2: Decomposition of the total effect of changing the number of competitors in Phase I (N_1) and the limit on the number of competitors allowed to enter Phase II (\bar{N}_2), following (1.12).

N_1	\bar{N}_2	η^*			SS		DOD		DOD (NR)	
		SS	DOD	DOD (NR)	Base	Opt	Base	Opt	Base	Opt
1	1	0.56	0.38	0.67	0.012	0.013	-0.021	0.005	0.012	0.012
2	1	0.50	0.32	0.63	0.028	0.029	-0.044	0.013	0.027	0.029
3	2	0.57	0.33	0.65	0.111	0.122	-0.139	0.045	0.096	0.102
4	2	0.59	0.29	0.52	0.246	0.271	-0.238	0.125	0.191	0.248

Table A.3: Optimal values of η from the perspective of social surplus (SS), DOD profits (DOD), and Phase III DOD profits ignoring the reimbursements for research costs (DOD (NR)). The values of these outcomes, in millions of dollars, are given in the columns marked “Opt.” The columns marked “Base” are the ones at the estimated value of $\eta = 0.73$.

to exert effort, and there will be very little surplus generated in the R&D contest.² On the other hand, if the firm is promised the entire surplus (i.e., if $\eta = 1$), then even though there may be a large amount of social surplus generated, the DOD will capture none of it; indeed, the DOD will run a negative profit in this setting if it accounts for refunding the firms’ research efforts out of pocket as well. For intermediate values of η , we would expect an inverted-U curve.³ Bargaining powers on the right side of this curve guarantee the procurer the same surplus while giving the firm a larger share of the social surplus, and they thus Pareto-dominate bargaining powers on the left side of the curve.

Figure A-1 plots these Laffer curves for $(N_1, \bar{N}_2) = (1, 1)$ and $(4, 2)$, at the parameters estimated for the respective values of N_1 . The figures for other levels of competition are qualitatively similar, and Table A.3 shows summary statistics for all four values of N_1 . At the estimated parameters, the DOD runs a small loss when taking into account research efforts. However, we do see the characteristic Laffer curve, and the DOD can increase profits by about \$26,000 in the case of $N_1 = 1$ and about \$350,000 per contest in the case of $N_1 = 4$ by giving the firms about a third of the surplus rather than three-fourths of it. Most of this savings comes from savings in research costs, however. If the DOD’s objective does not penalize research costs, the DOD would still prefer to reduce the bargaining parameter, but it would only do so to about two-thirds for $N_1 \leq 3$ and about one-half for $N_1 = 4$. The gains in this objective are considerably more modest as well. Note, however, that none of these changes would be Pareto improvements, as it would cost the firm surplus (but could in principle increase aggregate social surplus). Note further that the socially optimal level of η (discussed in Section 1.8 and included in Table A.3 for convenience) are larger than the DOD-optimal levels; this is expected, since the socially optimal point corresponds to a particular point on the frontier that weights DOD profits and firm profits equally. The conclusion from this section, therefore, is that I do not find evidence of a clear Pareto inefficiency in this setting: the estimated bargaining parameter suggests that the program lies on the Pareto frontier between the firms and the DOD.

²As long as exerting no effort corresponds to failure with probability 1, there will indeed be zero surplus generated through the R&D contest.

³In principle, there could be other nonmonotonicities in the intermediate region.

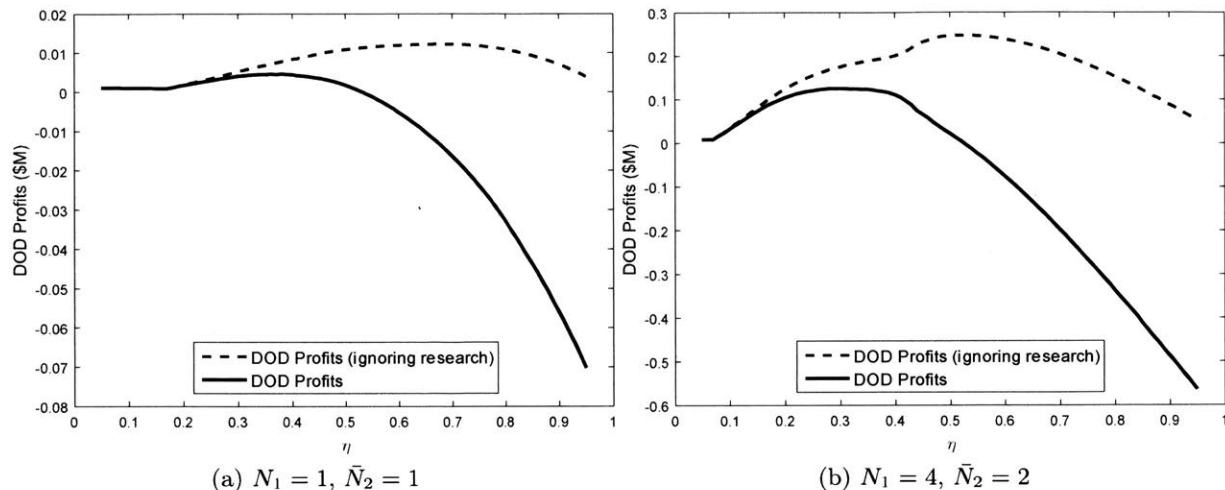


Figure A-1: Profit of the DOD (both ignoring and incorporating reimbursing the cost of effort) as a function of η for two different levels of competition.

Following Section 1.7.2, one could in principle define direct and incentive effects on the DOD surplus from an increase in η . The direct effect is unambiguously negative, as it simply reduces the surplus that the DOD captures. When firms readjust their research efforts from an increase in η , the DOD pays a cost for increased research but also enjoys a larger generated surplus. When η is high enough for the contest to lie on the Pareto frontier between firms and the DOD, the net effect is negative: the DOD loses a constant proportion of a larger expected surplus.⁴

A.1.3 Mandating IP Sharing

Finally, I study whether the DOD would benefit from mandating IP sharing. Figure A-2 shows a similar plot to Figure 1-4 for DOD profits, taking into account that the DOD pays both research costs and the Phase I prizes. Here, adding prizes without IP sharing only decreases profits: effort is slightly low relative to the social optimum but high relative to the DOD's ideal level. Introducing IP sharing (without prizes) cuts efforts drastically, but it does increase the DOD profits to around zero. Adding prizes in addition drops DOD profits further since it only increases Phase I efforts further. Thus, for these parameters, there is an inherent tension between the DOD and the social planner: the social planner would want to introduce IP sharing along with a nontrivial prize, whereas the DOD—which internalizes a larger share of research costs (relative to surplus) and the entire share of the prizes—prefers information sharing with no prize.

Table A.4, like Table 1.10, conducts the analysis for other values of N_1 and \bar{N}_2 , reporting

⁴To fix ideas, we can write the surplus when $N_1 = 1$ as $(1 - \eta)s(\eta) - t(\eta)$, where $s(\eta)$ is the equilibrium value of $\mathbb{E}[(v - c)^+]$ and $t(\eta)$ denotes expected research costs. The total effect is the derivative $-s(\eta) + s'(\eta)(1 - \eta) - t'(\eta)$, and the first term is the direct effect. This direct effect is clearly negative and increasing in magnitude with η . The first term of the incentive effect decreases to 0 for large enough η , and the exact values of slopes of $s(\cdot)$ and $t(\cdot)$ depend on the difference between values and costs as well as the estimated elasticity of costs. The incentive effect can also become negative for large η .




Figure A-2: DOD profits as a function of the Phase I prize, with and without mandatory information sharing. I use the estimated parameters for $N_1 = 2$ and use $N_1 = \bar{N}_2 = 2$.

N_1	N_2	$K = 0$		$\mathbb{E}[K(\cdot)]$	K_{SS}^*	
		SS	DOD		SS	DOD
2	2	0.054	-0.046	0.026	0.055	-0.083
3	2	0.111	-0.135	0.000	0.111	-0.135
3	3	0.139	-0.160	0.000	0.139	-0.160
4	2	0.246	-0.238	0.000	0.246	-0.238
4	3	0.362	-0.325	0.000	0.362	-0.325
4	4	0.498	-0.361	0.000	0.498	-0.361

(a) Phase I prizes with no IP sharing

N_1	N_2	$K = 0$		$\mathbb{E}[K(\cdot)]$	K_{SS}^*		K_{IC}		
		SS	DOD		SS	DOD	SS	DOD	
2	2	0.004	-0.000	0.062	0.093	-0.103	0.131	-0.001	-0.410
3	2	0.162	-0.130	0.000	0.162	-0.130	0.099	0.047	-0.469
3	3	0.013	-0.000	0.073	0.264	-0.103	0.083	0.185	-0.445
4	2	0.268	-0.245	0.000	0.268	-0.245	0.041	0.212	-0.407
4	3	0.400	-0.342	0.000	0.400	-0.342	0.002	0.396	-0.353
4	4	0.647	-0.294	0.001	0.650	-0.296	0.001	0.602	-0.345

(b) Phase I prizes with mandatory IP sharing

Table A.4: Expected prizes ($\mathbb{E}[K(\cdot)]$), social surplus, and DOD surplus for various levels of prizes both (a) without and (b) with IP sharing. I report outcomes for (i) no prizes, (ii) the socially optimal prize, and, in the case of IP sharing, (iii) the minimum prize to make IP sharing incentive-compatible. The DOD-optimal level of the prize is always 0. All values are in millions of dollars.

outcomes for DOD profits as well as social surplus for comparison. Panel (a) studies prizes without IP sharing. The DOD-optimal prize is always zero, as we would expect given the results on the socially optimal prize: the DOD-optimal level of Phase I research is lower than the socially optimal

one, and research is already overprovided in Phase I (see the discussion in Section 1.6). Note also that in the one situation where the socially optimal prize is positive ($N_1 = \bar{N}_2 = 2$), DOD profits decrease when the socially optimal prize is set. This result mirrors the one in Section 1.10.

Panel (b) considers mandatory IP sharing. IP sharing by itself ($K = 0$) can sometimes reduce losses for the DOD, usually by discouraging Phase I effort. Note that DOD-optimal level of the prize is always zero as well for these parameters. The socially optimal level of prize is often zero as well, and there are a few cases when social surplus can be improved and DOD losses reduced by a combination of IP sharing and prizes. In all but the case of $N_1 = \bar{N}_2 = 4$, the incentive-compatible prize reduces DOD profits—often by a large margin. These incentive-compatible prizes are often large, and the DOD has to pay them out of pocket.

A.2 Additional Empirical Results

In this Appendix, I provide further details about the descriptive statistics presented in Section 1.2.3. I also present and discuss structural estimates conditional on fixed values of η , to apply the baseline model to this setting.

A.2.1 Details on the Descriptive Statistics

Table A.5 reports OLS regressions of contest-level “success” rates, from Phase I to Phase II and from Phase II to Phase III, and columns (1) and (3) corresponds to the first two columns of Table 1.4. Columns (2) and (4) replace N_1 and N_2 by dummy variables, with $N_1 = 1$ and $N_2 = 1$ being the omitted categories. Column (2) shows that of the increase in Phase I success of 6.6 percentage points per Phase I competitor, the jump is especially prominent when moving from a single Phase I competitor to two Phase I competitors (about 12.7 percentage points), and then the marginal effect of adding competitors tapers off slightly. Column (4) shows that much of the increase in contest-level success in Phase II comes from moving from 2 competitors to 3 competitors in Phase II.

Table A.6 investigates the probability that an *individual* competitor generates successful research, and it extends columns (3) and (4) of Table 1.4. Note that individual-level success is unobserved, so I first provide details of the models estimated in this table. Column (2) suggests that the decrease in the probability of individual success (shown in column (1)) is especially prominent when moving from 2 to 3 Phase I competitors. Column (3) indicates that contestants in contests with one additional Phase II competitor have a higher probability of success, by about 2.8 percentage points, but column (4) qualifies this result by noting that the success rate drops when going from 3 to 4 competitors in Phase II.

Column (2) of Table A.7 extends the results shown in column (5) of Table 1.4 on funding amounts. Recall that contests with one more Phase I competitor have on average 1.6% more funding, and that more Phase II competitors has no impact on average funding. However, column (2) qualifies this result and indicates that there is a noticeable negative effect for contests with

	Pr(Contest Enters Phase II)		Pr(Contest Enters Phase III)	
	(1)	(2)	(3)	(4)
# Phase I Comp	0.066*** (0.009)		-0.018** (0.008)	
$N_1 = 2$		0.127*** (0.027)		-0.034 (0.023)
$N_1 = 3$		0.187*** (0.028)		-0.039 (0.025)
$N_1 = 4$		0.200*** (0.032)		-0.073** (0.029)
# Phase II Comp			0.076*** (0.016)	
$N_2 = 2$				0.066*** (0.017)
$N_2 = 3$				0.188*** (0.063)
$N_2 = 4$				0.301 (0.226)
Log(Avg Phase II Amt)			0.157*** (0.018)	0.157*** (0.018)
R^2	0.083	0.842	0.128	0.218
N	2773	2773	2292	2292

Table A.5: Regressions of a dummy of whether the contest enters Phase II (columns (1) and (2)) or Phase III (columns (3) and (4)) on the number of competitors in Phases I and II, controlling for year fixed effects, SYSCOM fixed effects, and topic covariates. I restrict the sample to contests with no more than 4 Phase I competitors. Columns (3) and (4) restrict to the set of contests that enter Phase II.

larger numbers of Phase II competitors: contests with four Phase II competitors have more than 20% less funding than contests with just one Phase II competitor.

Figure A-3 illustrates that projects with larger Phase II contracts tend to be more likely to succeed and enter Phase III. Panel (a) plots a local linear regression of contest-level success rate on average Phase II funding, and the estimated function is monotonically increasing both when uncontrolled and when controlling for covariates in a partially linear model.⁵ Panel (b) shows that this pattern holds *within contest* as well: I regress an indicator for winning a Phase III contract on (the log of) the ratio of the individual competitor's funding and the lowest funding awarded to a firm in Phase II of the same contest. Once again, better-funded projects have a higher probability of transitioning to Phase III even within contest.

⁵The regression that does not control for covariates simply runs a kernel regression of a dummy for success (either whether the contest succeeded, as in (a), or whether the individual firm was awarded a Phase III contract, as in (b)) on the log of the Phase II award amount, using the asymptotically optimal bandwidth. The regression that controls for covariates uses the semiparametric estimator proposed in Robinson (1988) to estimate the model $y = g(t) + X\beta + \epsilon$, where $g(\cdot)$ is nonparametric, and X is a vector of year fixed effects, SYSCOM fixed effects, and topic information.

	Phase I		Phase II	
	(1)	(2)	(3)	(4)
# Phase I Comp	-0.128*** (0.008)		-0.023*** (0.008)	
$N_1 = 2$		-0.027 (0.028)		-0.036 (0.029)
$N_1 = 3$		-0.264*** (0.028)		-0.043 (0.033)
$N_1 = 4$		-0.291*** (0.030)		-0.098** (0.037)
# Phase II Comp			0.028*** (0.010)	
$N_2 = 2$				0.010 (0.018)
$N_2 = 3$				0.097*** (0.039)
$N_2 = 4$				0.062 (0.370)
Log(Phase II Amount)			0.250*** (0.031)	0.271*** (0.058)
N	2773	2773	2292	2292

Table A.6: MLE estimates of the probability that an individual firm generates a successful innovation in Phase I (columns (1) and (2)) or Phase II (columns (3) and (4)), correcting for unobserved successes in a model-based manner, as described in the text. Regressions control for year fixed effects, SYSCOM fixed effects, and topic covariates. I restrict the sample to contests with no more than 4 Phase I competitors. Columns (3) and (4) restrict to the set of contests that enter Phase II.

Finally, column (6) of Table A.7 extends the results of column (6) of Table 1.4 by adding dummies for the number of competitors, and the results are not substantively different. Columns (3) and (4) of Table A.7 run the same regressions but include Phase III contracts with values less than \$1 million.⁶ The qualitative results do not change, and the quantitative results are similar for Phase I competition. However, the effect of Phase II competition is reduced by about 33% (see column (3)), and the dummies for $N_2 = 2$ and $N_2 = 3$ are smaller as well. Much of the effect comes from the large effect of $N_2 = 4$ on Phase III amount. Moreover, the correlation between Phase II amount and Phase III amount is larger. While it is reassuring that the qualitative results do not depend on arbitrary sample selections, I choose to drop especially low Phase III amounts throughout the analysis in the chapter. They are overwhelmingly near the beginning of my sample, when the Phase III contract in the data did not always correspond to delivery. They are sometimes placeholder contracts as well, to give the firm some interim funding for “transitioning” the product further into something the Navy plans to use, and I am unable to identify the actual delivery contract in the FPDS. I thus treat these low values as data that is missing at random.

⁶Some of these contracts are actually unreasonably small and sometimes amount to \$50,000.

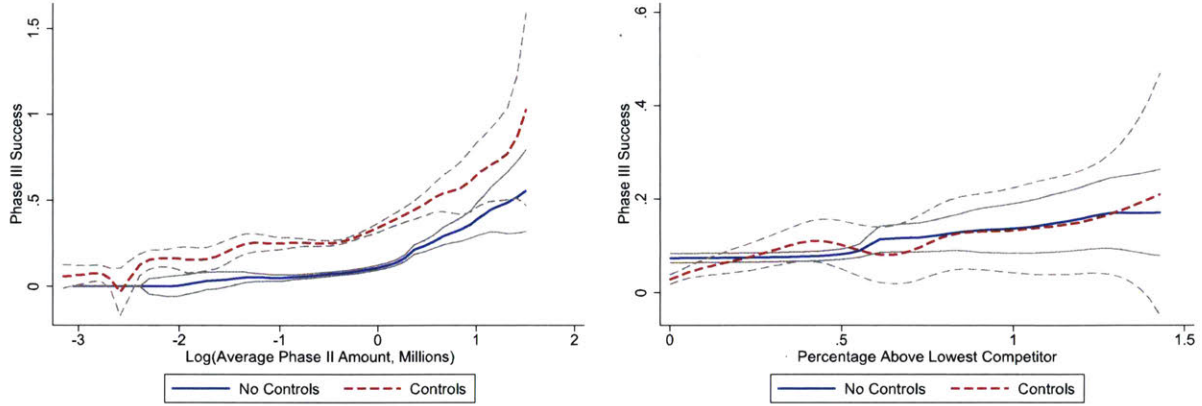


Figure A-3: (a) Local linear regression of contest-level success from Phase II to Phase III on the log of average Phase II funding, both controlling and not controlling for covariates. (b) Local linear regression of individual success on the ratio between the particular individual's Phase II funding amount and the lowest funding amount within that contest.

A.2.2 Structural Estimates without Assumption O

In this section, I provide estimates using Assumption M and an analogue of the scaling property of Proposition 1.4. In particular, I make the following assumption, which amounts to assuming the result of Proposition 1.4 without imposing Assumption O. The following assumption is also natural and, like Proposition 1.4, can be thought of as simply a change of units.

Assumption S. Fix N_1 and \bar{N}_2 . Suppose that \hat{p} and $\hat{t}_{N_2}(v)$ are the effort rules in Phases I and II (if N_2 firms enter Phase II) for a contest with primitives $\psi(\cdot)$, V , $C(t)$, and η . Consider a “scaled” contest with primitive $\tilde{\psi}(\cdot) = \theta \cdot \psi(\cdot)$, $\tilde{V} = \theta \cdot V$, $\tilde{C}(t) = \theta \cdot C(t/\theta)$ (i.e., so that $\tilde{H}(c, t) = H(c/\theta, t/\theta)$), and $\tilde{\eta} = \eta$. Then, \hat{p} and $\theta \cdot \hat{t}_{N_2}(v)$ are the effort rules in the scaled contest.

Under Assumptions M and S(caling), Steps 1–3 of the estimation procedure outlined in Section 1.4.3 are still valid. I run these steps, fixing three different values of η . Doing so will help give a sense of how η affects the estimates and how alternate assumptions about the bargaining procedure would affect the results. Table A.8 presents these new structural results for the value and delivery cost distribution. Since the distribution of unobserved heterogeneity is independent of η , I do not report those results again and instead refer to Table 1.5. Furthermore, I do not estimate the Phase I cost function, because doing so requires me to either make assumptions about the optimality of Phase I research effort or use the observed values for Phase I contracts as informative of $\psi(p^*)$.

Table A.8(a) reports moments value distributions for three different values of η . For the low ($\eta = 0.20$) and medium ($\eta = 0.50$) values of the bargaining parameters, competitions with a larger number of Phase I competitors are not necessarily associated with more valuable projects; the means of the value distributions are relatively similar across different values of N_1 . Indeed, the value estimates for these two values of η are similar to each other as well. The DOD values projects at around \$23–\$28 million, and the 95% range within a contest is about 12–13% of the mean.

	Log(Avg Phase II Amt)		Log(Phase III Amt)			
	(1)	(2)	(3)	(4)	(5)	(6)
# Phase I Comp	0.016 (0.012)		0.283* (0.154)		0.234** (0.110)	
$N_1 = 2$		-0.003 (0.038)		0.063 (0.372)		0.141 (0.310)
$N_1 = 3$		0.017 (0.039)		0.602 (0.418)		0.502 (0.322)
$N_1 = 4$		0.041 (0.045)		0.541 (0.465)		0.572 (0.381)
# Phase II Comp	-0.002 (0.016)		-0.289 (0.217)		-0.429** (0.176)	
$N_2 = 2$		-0.000 (0.020)		-0.221 (0.296)		-0.391* (0.205)
$N_2 = 3$		0.030 (0.052)		-0.280 (0.498)		-0.946* (0.488)
$N_2 = 4$		-0.254* (0.133)		-1.908** (0.872)		-1.044 (0.634)
Log(Avg Phase II Amt)			0.665** (0.272)	0.668** (0.276)	0.330* (0.195)	0.361* (0.204)
Min Phase III Amt			None	None	\$1 Million	\$1 Million
Adj R^2	0.133	0.448	0.326	0.465	0.422	0.881
N	2292	2292	233	233	151	151

Table A.7: Regressions of Phase II and Phase III award amounts on the number of competitors in other phases, controlling for year fixed effects, SYSCOM fixed effects, and topic covariates. Columns (3) and (4) restrict to the set of contests with a Phase III contract, and columns (5) and (6) restrict to contests with Phase III contracts of at least \$1 million. All columns restrict to contests with no more than 4 Phase I competitors.

The noticeable difference happens when η increases to 0.80. Here, the mean values are between \$6–\$17 million, and values do seem to increase with N_1 . The variation in value is similar to the other estimates, however. Note that the estimates with $\eta = 0.73$ presented in Table 1.5 roughly lie between the estimates for $\eta = 0.50$ and $\eta = 0.80$.

The identification argument presented in Section 1.4.1 would seem to suggest that the estimate of the value distribution should be independent of η . The dependence on η comes from two sources. First, identification in the model with unobserved heterogeneity does not rely on a clean upper bound, and the value estimates could interact with η . Second, parametric restrictions come into play. The most important such effect is that since the lower bound on costs is zero, the lowest possible value of a transfer for a particular value v is ηv . Thus, when η is low, low Phase III contracts can be explained by just a combination of values and costs even when values are estimated to be high; for higher values of η , low Phase III contracts must be explained by unobserved heterogeneity as well—or by low values. Thus, for a fixed data generating process, the observed values are

η	Values (\$M)	$N_1 = 1$	$N_1 = 2$	$N_1 = 3$	$N_1 = 4$
0.20	Mean	28.14 (3.67)	27.24 (3.25)	26.37 (3.50)	23.25 (4.96)
	Standard Deviation	0.90 (0.12)	0.86 (0.14)	0.83 (0.12)	0.74 (0.16)
	95% Range	3.51 (0.46)	3.38 (0.55)	3.26 (0.46)	2.91 (0.61)
0.50	Mean	26.91 (5.74)	22.93 (4.42)	24.50 (4.24)	25.97 (5.23)
	Standard Deviation	0.85 (0.19)	0.72 (0.13)	0.77 (0.14)	0.83 (0.17)
	95% Range	3.34 (0.73)	2.82 (0.52)	3.02 (0.54)	3.25 (0.65)
0.80	Mean	6.06 (1.36)	8.35 (1.65)	6.70 (1.55)	17.76 (3.06)
	Standard Deviation	0.20 (0.05)	0.25 (0.05)	0.19 (0.05)	0.53 (0.10)
	95% Range	0.77 (0.21)	0.97 (0.21)	0.76 (0.18)	2.09 (0.39)

(a) Value distributions

η	Pr($c < v$)		$\mathbb{E}[c c < v]$		Quantiles (\$M)			
	Value	Semi-Elasticity	Value	Elasticity	1%	5%	10%	Elasticity
0.20	0.068	0.013	14.83	-0.017	6.42	20.90	39.21	-0.172
	(0.010)	(0.005)	(1.19)	(0.005)	(0.77)	(2.53)	(4.76)	(0.047)
0.50	0.062	0.012	13.45	-0.017	6.26	20.37	38.22	-0.182
	(0.011)	(0.005)	(1.52)	(0.005)	(0.83)	(2.75)	(5.20)	(0.041)
0.80	0.064	0.012	4.57	-0.017	2.08	6.78	12.72	-0.176
	(0.010)	(0.002)	(0.51)	(0.002)	(0.29)	(0.95)	(1.79)	(0.024)

(b) Delivery cost distributions

Table A.8: Structural estimates for the baseline model, for three different values of η . The parameters presented are the same ones as in Table 1.5.

depressed when η is higher.

Table A.8(b) shows the moments of the cost distribution. The main observation is that the cost distributions are decreasing with η (both the mean and all quantiles). This dependence is the outcome of two main forces. First, the failure rate depends primarily on the difference between values and costs, and fitting the failure rate well when values decrease (due to higher η) requires costs to decrease as well. Indeed, the proxy for failure (the probability that the cost draw is larger than the mean value) is relatively fixed across η . A second, counteracting force is that the Phase III contract amount is roughly $\eta v + (1 - \eta)c$. Thus, to match the same transfer distribution with a higher η , the costs must be slightly *larger*, ignoring any change in v . The MLE procedure balances these

two effects. Finally, note that the elasticity is roughly independent of η ; this elasticity is mainly a function of the dependence of the failure rate and transfer distribution on Phase II research, and we would thus not expect it to vary with η other than for reasons due to heterogeneity or parametric assumptions.

A.3 Extensions to the Identification Result

In this Appendix, I discuss a number of extensions to the identification results in Section 1.4.1. The first result shows that Assumption M by itself provides some information about η . The other results consider generalizations of the model with Assumptions M and O and study identification of these more general models. First, I briefly note that the baseline argument can be applied to firms with asymmetric cost functions with almost no modifications. I then show a more involved argument that all primitives can be identified in the model with multiplicative unobserved heterogeneity, as in the empirical model in Section 1.4.2. Finally, I consider models in which there is an unobserved benefit to Phase II research not captured by the Phase III contract.

I also wish to briefly note that I conjecture that the primitives of the model are still identified when we only observe data in Phase II with multiple firms. While such a setting is irrelevant in the context of the model presented in this chapter, as pure randomness will ensure that there are at least some contests with one only competitor in Phase II, it could be relevant in a related but different model—such as one that begins in Phase II. Furthermore, I conjecture that the primitives are identified in certain generalizations of the model in which Phase II research efforts are based on the *expected* value of the project to the DOD but the actual value is realized only after Phase II research is completed. Future versions of this appendix will contain the technical conditions that yield identification in these settings.

A.3.1 Bounding η Using Assumption M

In this section, I show that while we cannot identify the bargaining parameter η exactly purely from Assumption M, we can identify a lower bound on this parameter. Fix values t' and $t'' > t'$ for the Phase II research effort, and suppose these research efforts correspond to $v(t')$ and $v(t'')$, respectively. Let the success rate at a research effort of t be denoted $g(t)$. Consider the $g(t')$ th quantile of the distribution of Phase III research efforts conditional on t'' , denoted $T_3(g(t'), t'')$.⁷ Note that $T_3(g(t'), t'') = \eta v(t'') + (1 - \eta)C(g(t'), t'')$, where $C(q, t'')$ is the q th quantile of the distribution of costs when research efforts are t'' . By stochastic dominance, $C(g(t'), t'') \leq C(g(t'), t') = v(t')$. Then,

$$T_3(g(t'), t'') = \eta v(t'') + (1 - \eta)C(g(t'), t'') \leq \eta v(t'') + (1 - \eta)v(t').$$

Rearranging, we have

$$\eta \geq \frac{T_3(g(t'), t'') - v(t')}{v(t'') - v(t')}. \quad (\text{A.1})$$

⁷Do not normalize the distribution of T_3 by the failure rate; this quantile would correspond to the $[g(t')/g(t'')]^{\text{th}}$ quantile of the distribution of Phase III transfers conditional on success.

Since (A.1) has to apply for all t' and t'' , we have the lower bound

$$\eta \geq \max_{t', t'' > t'} \frac{T_3(g(t'), t'') - v(t')}{v(t'') - v(t')}. \quad (\text{A.2})$$

Thus, we have the following proposition.

Proposition A.1. *Suppose we have data on the distributions of Phase III transfers and Phase II research efforts. If Assumption M holds, then a lower bound on η is identified from (A.2).*

A.3.2 Asymmetric Firms

Suppose that there are multiple types of firms, indexed by k , whose types are known by the researcher ex-ante. They differ in their cost distributions as well as the distributions from which their values are drawn. That is, a type k firm draws a value $v \sim V_k$ upon entering Phase II and a cost $c \sim H_k(\cdot; t)$ if spends t on Phase II research effort. The Phase III allocation rule is the same as in Section 1.3. Entry into Phase II is determined so that in equilibrium the number of Phase II competitors of type k is drawn from N_{1k} Phase I competitors, each of whom succeed in Phase I with probability p_k . As long as $p_k \in (0, 1)$, there is a positive probability that the *only* entrant into Phase II is an entrant of type k . We can then focus only on these contests and apply the argument in Section 1.4.1 directly to identify η , V_k , and $H_k(\cdot; \cdot)$;

A.3.3 Unobserved Heterogeneity

Suppose that, as in the empirical specification in Section 1.4.2, each contest is associated with a multiplicative error term θ_j so that for contests with multiple Phase II competitors, we observe $\theta_j \tilde{t}_i$ for each firm i and a Phase III transfer $\theta_j T$, if there is a winner. Suppose further that the distribution of θ_j is common across *all* contests. I will show that this distribution is identified using data from contests with $N_2 \geq 2$. The value distribution \tilde{V} is identified as well, and the cost distribution $\tilde{H}(\cdot; \tilde{t})$ is identified as a function of η from contests with $N_2 = 1$. The identification results also require some technical conditions on the (endogenous) failure rate and unobserved heterogeneity.⁸

Identification of Θ . Consider two firms with Phase II efforts $\theta_j \tilde{t}_1$ and $\theta_j \tilde{t}_2$. From the joint distribution of the logs of these efforts, we can recover the distribution of θ_j , and its associated density $f_\Theta(\cdot)$, as well as that of \tilde{t} via appealing to Kotlarski (1967), as long as these distributions have nonvanishing characteristic functions.⁹ Assume the distribution of unobserved heterogeneity does not have especially large mass around 0, so that $1/\theta$ has a mean.

Identification of the Failure Rate. For the remainder of the argument, focus on contests with $N_2 = 1$. Define the failure rate conditional on $\theta = 1$ as $g(\tilde{t}) \equiv 1 - \tilde{H}(\tilde{v}(\tilde{t}); \tilde{t})$ (and zero if \tilde{t} is outside

⁸While I have not shown that these technical conditions are necessary, they are needed for the proof I provide.

⁹Evdokimov and White (2012) provide alternate conditions under which this identification result remains true.

the support of its distribution). In the data, note that we observe the failure rate conditional on a particular value of t (rather than a value of \tilde{t}); that is, we observe

$$k(t) \equiv \int_0^\infty f_\Theta(\theta) \cdot g\left(\frac{t}{\theta}\right) d\theta, \quad (\text{A.3})$$

for each t that can be expressed as some $\theta\tilde{t}$, where θ and \tilde{t} are both in the supports of the distributions of unobserved heterogeneity and efforts. The basic idea is that since we know $k(t)$ and $f_\Theta(\theta)$, we can apply a deconvolution argument to find $g(\cdot)$. For convenience, we consider Fourier transforms with respect to the group $G \equiv (\mathbb{R}^+, \cdot)$ and the associated Haar measure $d\theta/\theta$.¹⁰

Suppose that $g(\cdot)$ is such that $\int_0^\infty g(\theta) d\theta < \infty$.¹¹ Then, $\theta \mapsto \theta g(\theta)$ is in $L^1(G)$. Furthermore, $f_\Theta \in L^1(G)$ as well since $1/\theta$ is assumed to have a finite mean. Express (A.3), after regrouping terms and reparameterizing, as

$$k(t) \equiv \int_0^\infty \left[f_\Theta\left(\frac{t}{u}\right) \right] \cdot [u \cdot g(u)] \frac{du}{u}.$$

Both terms in brackets are L^1 and thus so is k by Hölder's inequality. Thus, taking the Fourier transform, we have that $\hat{k}(s) = \hat{f}_\Theta(s) \cdot \widehat{[\theta g(\theta)]}(s)$. As long as the transform of f_Θ is nonvanishing, we can recover the transform of $\theta \cdot g(\theta)$ and thus the failure rate $g(\cdot)$ itself.

Identification of the Cost Distribution. Consider the conditional distribution $\ell(T|t)$ of Phase III transfers conditional on a particular value of the Phase 2 transfer t , normalized so that $\int \ell(T|t) = 1 - k(t)$. This quantity is observed in the data, and we can express it as

$$\begin{aligned} \ell(T|t) &= \int_0^\infty f_\Theta(\theta) \cdot \Pr\left(\eta\tilde{v}\left(\frac{t}{\theta}\right) + (1-\eta)c = T; \frac{t}{\theta}\right) d\theta \\ &= \int_0^\infty f_\Theta(\theta) \cdot \frac{1}{1-\eta} \tilde{h}\left(\frac{T/\theta - \eta\tilde{v}(t/\theta)}{1-\eta}; \frac{t}{\theta}\right) \frac{d\theta}{\theta} \\ &\equiv \int_0^\infty f_\Theta(\theta) \cdot q\left(\frac{T}{\theta}; \frac{t}{\theta}\right) \frac{d\theta}{\theta}, \end{aligned}$$

¹⁰Fourier transforms are defined for L^1 functions for any locally compact commutative group, including the group of positive numbers endowed with multiplication as long as the measure with respect to which we are integrating is translation-invariant (i.e., is the *Haar measure*). In this case, $f \in L^1$ if $\int_0^\infty f(\theta)/\theta d\theta < \infty$. Convolution is defined as $(f * g)(s) \equiv \int_0^\infty f(\theta)g(s/\theta)/\theta d\theta$, and Fourier transforms are such that $\widehat{(f * g)}(s) = \hat{f}(s) \cdot \hat{g}(s)$. See Theorems 1.2.4(b) and 1.1.6(e) in Rudin (1962) for the result that multiplication is the dual of convolution and Hölder's inequality, respectively, in this setting. For the purposes of this proof, it would suffice to reparameterize functions and the transform, but doing so at each step would be unnecessarily cumbersome.

¹¹Since $g(\cdot)$ is endogenous, this may not be a natural condition. But, note that $g(t) = 1 - H(v(t); t)$, and $v(t)$ is increasing in t but $1 - H(v; t)$ is decreasing in t . Thus, $g(t) = 1 - H(v(t); t)$ is integrable in the standard sense if $H(\cdot; t)$ is integrable in the standard sense for all v . That is, if the cost distribution decays quickly enough, then we can guarantee that the failure rate satisfies the technical condition.

where the final line defines $q(T; t)$. But, note that we can redefine $\ell(T|t)$ as $\ell^*(T/t; t)$ and $q(T; t)$ as $q^*(T/t; t)$. Thus, fixing a t , we can use the Fourier transform to recover $q^*(\cdot; t)$ for all t .¹² This in turn recovers $q(T; t) = \tilde{h}(T/(1 - \eta) - \eta\tilde{v}(t)/(1 - \eta); t) \cdot (1/(1 - \eta))$, which means that the cost distribution is recovered as a function of η and $\tilde{v}(t)$.

Identification of V . To identify $\tilde{v}(t)$, note that

$$\int_{-\infty}^{\tilde{v}(\tilde{t})} \tilde{h}(c; \tilde{t}) dc = g(\tilde{t}).$$

Substituting $c = T/(1 - \eta) - \eta\tilde{v}(\tilde{t})/(1 - \eta)$, we have

$$\int_{-\infty}^{\tilde{v}(\tilde{t})} q(T; \tilde{t}) dT = g(\tilde{t}).$$

Since $q(T; \tilde{t}) \geq 0$, this equation has a solution, and $\tilde{v}(\tilde{t})$ is identified from matching the failure rate. Transforming the distribution of \tilde{t} by $\tilde{v}(\cdot)$ recovers the distribution of \tilde{V} .

Identification of η . Since we have recovered $\tilde{v}(\cdot)$ and $\tilde{H}(\cdot; \cdot)$ as a function of η , we can apply the same argument as in Section 1.4.1 to recover η , if we make the assumption that the transfer is the firm-optimal one.

I summarize the results of this section in the following proposition.

Proposition A.2. *Consider the equilibrium model with multiplicative unobserved heterogeneity. Suppose that the distribution of unobserved heterogeneity is such that its inverse has a mean and that the distribution of costs is such that $\int_0^\infty [1 - H(c; t)] dt < \infty$ for all c . Then, as long as we observe contests in which at least two competitors enter Phase II, the distribution of Θ is nonparametrically identified, as are V , $H(\cdot; t)$, and η .*

I briefly comment on the strategy of using different sets of contests to identify these distributions, as a concern may be that the distribution of unobserved heterogeneity may itself differ across these contests. However, note that a natural reason for this distribution to differ is that the DOD may choose a different number of Phase I competitors for contests with different distributions of competitors; however, once this choice is made, the remainder of the contest follows in a somewhat mechanical manner. In particular, we can let Θ depend on N_1 , but as long as $N_1 > 1$ (and $\bar{N}_2 \geq 2$), there will be contests that enter Phase II with $N_2 \geq 2$ and $N_2 = 1$ due to pure randomness. Of course, the arguments in this section would not apply if there were a special set of auctions where only one competitor were allowed to enter into Phase II (i.e., if $N_1 = 1$ or $\bar{N}_2 = 1$). There is not much we can do in this situation, as unobserved heterogeneity can only be identified from information on correlation between actions of the firms within a particular contest.

¹²It can be checked that $1/\theta$ have a mean is sufficient for this transform to apply. Note that $q^*(\cdot; t)$ is a density and is thus integrable in the standard sense.

A.3.4 External Benefits of Research

In the model presented in Section 1.3.2, the only benefit of conducting research comes from the possibility of winning a Phase III contract. Since the firms involved in an SBIR contest retain intellectual property rights over their innovations, one may speculate that there could be an additional benefit of doing research. Extend the model in Section 1.3.2 to one in which a firm that exerts effort t also gets a benefit $b(t)$ in addition to the benefit from the Phase III contract, with $b'(\cdot) \geq 0$ and $b''(\cdot) < 0$. It is easy to check from a monotone comparative statics argument that firms with higher values of v will still exert more effort. Furthermore, in the case of $N_2 = 1$, the observed Phase III contract amounts will still be $\eta v + (1 - \eta)c$, so Proposition 1.2 applies immediately to this model.

However, the first order condition of the firm changes from (1.5) to

$$b'(t_2) + \eta \int_{\underline{c}}^{v(t_2)} (v(t_2) - c) \frac{dh}{dt}(c; \eta, t_2) dc = 1, \quad (\text{A.4})$$

so applying Proposition 1.3 requires more conditions. Suppose that $b'(t_2)$ is known for some value $t_2 = t_2^*$. Then, for t_2^* , we can apply the argument in Proposition 1.3 to identify η . Then, for all other t_2 , (A.4) identifies $b'(t_2)$. We codify this in the following proposition.

Proposition A.3. *Consider the model in Section 1.3.2 but suppose firms get a benefit $b(t)$ from exerting effort t . Suppose that the value of $b'(\cdot)$ is known at some point and that the value of $b(\cdot)$ is also known at some (possibly different) point. Then, $b(\cdot)$ is identified as well over the range over which firms exert effort.*

The summary of this extension is that the overidentifying restrictions embedded in the fact that the firms' first-order condition must hold at all points can be used to identify any external marginal benefit of research. The two caveats are that we need some external information about both this marginal benefit as well as some information about the level of the benefit itself. Without an external information about the marginal benefit, we cannot disentangle it from the impact of η . Without some information about the benefit itself, we have no hope of identifying it, since the data contain absolutely no information about the level of this benefit. However, natural conditions exist for both the benefit and the marginal benefit. We may expect $b'(t) = 0$ for a sufficiently large value of t , as the marginal benefit may decrease; we may also expect $b(0) = 0$.

A.4 Incentive Compatibility in the Model of IP Sharing

In this Appendix, I discuss the issue of incentivizing firms to share their Phase I breakthroughs with its competitors. To do so, I explicitly model the subgame following a deviation in which a firm with value v chooses not to share its IP. This allows for the computation of the value of deviating, and I can then compute the minimum prize schedule necessary to preclude this deviation. Note that I do not ask whether this is the *optimal* schedule for the DOD in the setting in which IP sharing is not

mandatory; the DOD may well choose a prize schedule that induces sufficiently high-value firms to keep their IP private. Rather, this incentive-compatible schedule simply serves as a benchmark for comparison.

To compute this incentive-compatible prize schedule $K(v)$, I first need to compute the profits under a “deviation” in which a firm with value v refuses to share information. I consider the following setting: at the end of Phase I, the DOD offers the prize $K(v)$ to all firms with successful innovations. However, unlike before, the DOD allows any firm to forego the prize $K(v)$ in exchange for keeping the invention secret; the firm is still allowed to enter Phase II if its draw of v is high enough to merit entry into Phase II. The DOD does not reveal whether firms shared their information or not, and it still shares the plans of the highest-value project from the other firms with the holdout. Moreover, it does not reveal whether or not each firm accepted the prize.¹³

The deviation I consider, therefore, consists of the following steps.

- (i) A firm with value v gives up the prize $K(v)$ but enters Phase II if its draw of v is in the top \bar{N}_2 of the draws. It must decide whether to accept the Phase I prize before learning how many other firms succeeded.
- (ii) In Phase II, it gets access to the highest draw v' of all other firms and chooses which project it wishes to develop. If no other firm succeeded in Phase I, the deviator is the only firm in Phase II and exerts effort according to the equilibrium of the model in Section 1.3.2, with $N_2 = 1$.
- (iii) Beliefs of all other firms are passive, so all other firms in Phase II (if any) exert the equilibrium effort $t^*(v')$ on the project v' .

These criteria together let us derive the equilibrium effort exerted by a firm with value v that deviates, if all other firms are using the technology with value v' . Denote this profit by $\hat{\pi}_{\text{success}}(v, v'; p^*, K(\cdot))$. The incentive compatibility condition is that $K(v) \geq \mathbb{E}[\hat{\pi}_{\text{success}}(v, v'; p^*, K(\cdot))] - \pi_{\text{success}}(v; p^*, K(\cdot))$, where the first expectation is taken over the realization of successes as well as the best value of the opponents. Note that if this IC constraint holds with equality, $K(\cdot)$ must be increasing, as a firm with a higher-value project will pay a large cost in terms of expected forgone profits if it shares its breakthrough with its competitors.

A.5 Omitted Proofs

I collect the proofs of Propositions 1.2, 1.3, and 1.5 in this appendix.

¹³Specifying this deviation involves a number of assumptions on the details of the information sharing mechanism. One could imagine other mechanisms that differ in some respects; for instance, the DOD could refuse to share other firms' plans with a firm that does not accept the prize $K(v)$. One justification for the willingness to share plans is from social surplus reasons: while the DOD could in principle improve its profits by committing to not share plans (and reduce the prize it has to pay), the social planner would always want a firm to have access to a project that could be potentially better. The DOD can also choose to announce which firms were willing to share their plans. However, I avoid this possibility out of convenience: if deviations were public, I would have to be explicit about off-path beliefs, which in turn would affect the incentives to deviate.

A.5.1 Proof of Propositions 1.2 and 1.3

The one remaining step to prove Proposition 1.2 is to consider the case in which there is selection into Phase II (i.e., when $\bar{N}_2 = 1$ and $N_1 > 1$). In this case, the argument in Section 1.4.1 shows that the *selected* distribution is identified; denote this V_S . However, note that this selected distribution is a known mixture of order statistics of the unselected distribution V . In particular, V_S is the maximum of N_S draws from V if N_S firms succeed in Phase I. Since the probability p^* of any individual firm succeeding in Phase I is identified directly from the data, we can express the cdf $F_{V_S}(\cdot)$ of the V_S in terms of the cdf $F_V(\cdot)$ of V as

$$F_{V_S}(v) = \frac{1}{(1-p^*)^{N_1}} \sum_{N_S=1}^{N_1} \binom{N_1}{N_S} (p^*)^{N_S} (1-p^*)^{N_1-N_S} F_V(v)^{N_S}. \quad (\text{A.5})$$

The right-hand side of (A.5) is a convex combination of increasing functions of $F_V(v)$. Since $F_{V_S}(v)$ is identified, we can invert (A.5) to identify the cdf of V .

The missing step in Proposition 1.3 is to show that (1.5) has a unique solution for η . Recall that the firm sets its research effort in response to the first-order condition in Section 1.4.1, given by

$$\eta \int_{\underline{c}}^{v(t_2)} (v(t_2) - c) \frac{dh}{dt}(c; \eta, t_2) dc = 1 \quad (1.5)$$

Integrating (1.5) by parts, we have

$$\eta \int_{\underline{c}}^{v(t_2)} \frac{dH}{dt}(c; \eta, t_2) dc = 1. \quad (\text{A.6})$$

However note that

$$\begin{aligned} H(c, t_2; \eta) &= \Pr(C(t_2) \leq c | t_2, \eta) = \Pr(\eta v(t_2) + (1-\eta)C(t_2) \leq \eta v(t_2) + (1-\eta)c) \\ &\equiv \hat{F}(\eta v(t_2) + (1-\eta)c; t_2), \end{aligned}$$

which is the cdf of the transfer evaluated at $\eta v(t_2) + (1-\eta)c$, an observed quantity (as a function of η). Substituting into (A.6), we have

$$\eta \int_{\underline{c}}^{v(t_2)} \left(\frac{d\hat{F}}{dt}(\eta v(t_2) + (1-\eta)c; t_2) + \eta v'(t_2) \hat{f}(\eta v(t_2) + (1-\eta)c; t_2) \right) dc = 1.$$

Setting $u = \eta v(t_2) + (1-\eta)c$, we have

$$\frac{\eta}{1-\eta} \int_{\underline{T}}^{v(t_2)} \frac{d\hat{F}}{dt}(u; t_2) du + \frac{\eta^2}{1-\eta} \int_{\underline{T}}^{v(t_2)} v'(t_2) \cdot \hat{f}(u; t_2) du = 1, \quad (\text{A.7})$$

where \underline{T} is the minimum transfer observed. Given that $v(\cdot)$ and thus $v'(t_2)$ are both identified already from the support of the transfer distribution, the integrals are identified directly from the

data. Thus (A.7) can be rearranged to a quadratic in η and has at most two solutions, only one of which corresponds to the actual optimum (as the other violates the second order condition). Thus, η is identified, which in turn identifies the cdf $H(c; t_2)$ of $C(t_2)$ nonparameterically for all $c \leq v(t_2)$.

A.5.2 Proof of Proposition 1.5

Suppose the social planner can pick a schedule $t_i(v)$ of effort for each firm i as a function of the firm's realized value v . This choice induces a random variable $S_i(v, t_i(v))$ of the surplus each firm i . Fix a distinguished firm i . Then, the social planner's problem can be written as

$$\max_{t_i, t_{-i}} \left\{ \mathbb{E} \left[\max \{ S_i(v_i, t_i(v_i)), \max_{-i} S_{-i}(v_{-i}, t_{-i}(v_{-i})) \}^+ \right] - \mathbb{E} [t_i(v_i)] - \sum_{-i} \mathbb{E} [t_{-i}(v_{-i})] \right\},$$

where the expectations are taken over realization of v . If we denote the social planner's optimum as $t_{-i}^*(\cdot)$, to determine $t_i^*(v)$, the planner will be optimizing

$$\begin{aligned} & \max_t \left\{ \mathbb{E} \left[\max \{ S_i(v, t), \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i})) \}^+ \right] - t \right\} \\ &= \max_t \left\{ \mathbb{E} \left[\{ S_i(v, t) - \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i})) \}^+ + \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i})) \right] - t \right\} \\ &= \max_t \left\{ \mathbb{E} \left[\{ S_i(v, t) - \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i})) \}^+ \right] - t \right\}, \end{aligned} \tag{A.8}$$

where from the second to the third line, I drop $\max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i}))^+$ since it is independent of t . (Note that expectations are taken only over realization of v_{-i} in this sequence.) But, (A.8) is identically the expression for firm i 's problem when $\eta = 1$. Thus, the social planner's optimum corresponds to a Nash equilibrium of the game.

To show that the social surplus is monotone in η , we consider a different problem for notational convenience. Consider the problem $\max_t [\eta f(t; v) - t]$ where f is increasing in t . Denote the solution to this problem as $t^*(\eta; v)$ and note that this solution is increasing in η due to the fact that the maximand has increasing differences in η and t . Consider the function $g(v; \eta) \equiv f(t^*(\eta; v); v) - t^*(\eta; v)$. The derivative with respect to η is

$$\frac{dt^*(\eta; v)}{d\eta} [f'(t^*(\eta; v); v) - 1].$$

But, $dt^*(\eta; v)/d\eta \geq 0$. Moreover, we know that $\eta f'(t^*(\eta; v); v) = 1$ at an interior solution, so $f'(t^*(\eta; v); v) \geq 1$. Thus, $g(v; \eta)$ is increasing in η for all v . Since the social surplus is simply $\mathbb{E}[g(v; \eta)]$, where the expectation is taken over η , we have that the social surplus is increasing in η .

A.6 Data Appendix

In this appendix, I provide further details about the data collection and cleaning procedure, as well as how datasets from different sources are cross-checked and merged together.

A.6.1 SBIR Data from the Office of Naval Research

The website www.navysbirsearch.com has information about all SBIR contracts let by the Navy. Each entry contains the SBIR topic number, company information (name, address, DUNS number, and information about the PI in charge of the project), the SBIR Phase the contract is associated with, the federal contract number associated with the award, the SYSCOM in charge of letting the project, an award amount (which I clean later using the Federal Procurement Data System), and the start and end dates of the contract. It also includes the title of the proposal along with the full text of the abstract. I first scraped the data from the website and corrected obvious mistakes in the dataset, including fixing invalid contract numbers (where the correct numbers are clear) and dropping duplicate observations. I define a contest to consist of all Phase I, II, and III awards given under a particular topic number, and I can track a firm through the three phases using its unique DUNS number.

There are two minor considerations at this step. First, in a small number of cases, two different Phase III SBIR awards (given to two different companies) were listed with the same contract number but belonging to two separate contests. I treat these joint awards simply as separate awards for each contest. Second, there are a small number of contests in which the number of Phase II competitors is larger than the number of Phase I competitors. Since the Navy does not award direct-to-Phase II awards (i.e., every firm that wishes to compete in Phase II must also have competed in Phase I), I assume that these are data errors and that the competitor who appears first in Phase II actually was awarded a Phase I contract as well that was not in my dataset; however, I do not see the abstract and title for this project, and I assume that the Phase I contract amount (which I do not use in the analysis) is the standard amount without an option.

A.6.2 Federal Procurement Data System

From the Federal Procurement Data System (FPDS) via www.usaspending.gov, I downloaded contract data for all contracts from the Department of Defense from 2000 onwards. I use this dataset as the source for contract values: data from the ONR sometimes simply lists a standard SYSCOM-specific award amount. From this dataset, I extract all contracts where the contract id number matches one from the ONR dataset. I then check that the DUNS number and the firm name match for the merged contracts. For cases that are not exact matches, I verify through online searches that the difference can be attributed to a name change or an acquisition. I am unable to verify whether the datasets are merged properly for a small handful of contracts.

For each contract number, the FPDS contains an entry for the base contract, which contains information for the total contract value and the total value of all options to the contract. The FPDS

also contains an entry for each contract modification (e.g., remitting payment, change of scope, or exercising an option) which lists the changes to the total contract value. From this data, I can compute the total funding provided through the contract by summing across the dollars obligated in the base contract as well as all contract modifications, and for the majority of contracts, I use this measure as the contract value. For the vast majority of contracts, this amount agrees with the ONR data to within \$1, and the amounts differ by less than 5% for the majority of the remainder. Many of the remaining discrepancies can be explained by a single contract modification (exercise of an option, change of scope, or dollars de-obligated) recorded in the FPDS data that is not reflected in the ONR dataset. I use the data from the ONR as the measure of contract values if (i) I am unable to verify whether the merge is correct, (ii) the FPDS yields a contract amount that is less than 25% of the ONR data, or (iii) the base contract is missing in the FPDS.

A.6.3 SBIR Solicitations

I copied the full text of the Navy SBIR solicitations from the DOD archive of solicitations, from 1999 onward. For each topic number, I created a document containing all abstracts from all winning firms and all phases related to the topic as well as the full solicitation. This set of documents comprises the “corpus” that I fed into MALLET (McCallum, 2002) to generate the technology topics.

I train topics using the entire set of contests available to me, including those that I exclude from the final sample. When using MALLET, I treat each document as a sequence of word features (rather than merely a vector), remove stopwords such as “the” or “and”, and keep punctuation as part of the words. I also allow for hyperparameter optimization every 20 iterations so that MALLET optimizes over the distribution of topics and allows some topics to be more prominent than others. I let the sampling run for 5000 iterations; note that there is no upper limit on this, and I have noticed that running it for much fewer iterations would yield essentially indistinguishable results.

I set the number of topics to 20, but I have done robustness checks on the descriptive regressions using between 10 and 100 topics. MALLET outputs a set of topics, each of which is described by a list of words that categorize these topics. Of the 20 topics, 19 of them correspond to technology areas; the most popular one, however, consists of generic terms such as “system”, “phase”, “technology”, “design”, and “navy”. I drop this topic from the list and use the remaining 19 topics. MALLET also assigns each document d a weight p_{dt} for each topic t such that $\sum_t p_{dt} = 1$ for all d . I renormalize these proportions after eliminating the generic topic and am left with a set of 19 “fixed effects” for each contest.¹⁴ When grouping words into more topics, a larger number of the generated topics correspond to generic words instead of technology areas. For instance, when generating 100 topics, I categorize 5 of them as generic and ignore them when computing the proportions for each contest.

¹⁴Note that these variables are not fixed effects because they are proportions rather than binary variables. However, they do still sum to 1.

A.7 Computational Methods

A.7.1 Computing and Optimizing the Likelihood Function

Computing the Likelihood Function

I make three comments about computing the likelihood function. First, I parameterize $\mu(t)$ to be a decreasing quadratic in $\log t$ on the interval $-2.0 \leq \log t \leq 2.5$, which encompasses almost all the data. (Note that t is measured in terms of multiples of the mean Phase II amount, so this range is rather large.) I then parameterize this function by three values: (1) $\mu(t)$ at $\log t = -2.0$, (2) $\mu'(t)$ at $\log t = -2.0$, and (3) $\mu'(t)$ at $\log t = 2.5$. I constrain the parameters in (2) and (3) to be negative. To avoid numerical issues related to the quadratic function becoming increasing outside this range (which may be encountered for especially small or large values of θ), I let $\mu(t)$ be linear in $\log t$ for values of t outside this range, and I ensure that $\mu(t)$ is differentiable everywhere by setting the semi-elasticity of $\mu(t)$ (i.e., $d\mu(t)/d\log t$) outside this range equal to the value at the closest endpoint. Extending the range over which $\mu(t)$ is quadratic does not seem to change the results appreciably.

Second, I evaluate all integrals numerically on a fixed set of grid points (although the method and number of grid points varies by the particular integral). The likelihood seems to be robust to the number of grid points I use.

Finally, for certain parameters, specific observations are computed to have a likelihood of zero. Instead of letting the log likelihood function be $-\infty$ at these parameters, I replace the zeros with a penalty term π_{penalty} . One can imagine this procedure as an ad-hoc analogue of the “robust likelihood” for discrete distributions, suggested by Owen (2001). For the results in this chapter, I use $\log \pi_{\text{penalty}} = -100$. None of the data points in the main estimates of the chapter are affected by this penalty term, and at most three of the data points in each of the estimates reported in this chapter have a likelihood given by this penalty term. Indeed, estimates do not change for penalties in an neighborhood of this value, but if the penalty is taken to be too low, then a large portion of points are simply rationalized by this penalty.

Optimizing the Likelihood Function

The likelihood may have multiple local optima, and the discrete penalties to deal with data points that have zero probability (see below) for certain candidate parameter values introduce discrete jumps. Thus, I use a derivative-free global optimizer first, to narrow my search to a region of the parameter space where most of the data is rationalized well by the model. I then use a derivative-free local optimizer to polish the solution within this parameter region. Both algorithms I am using are implemented in the software package NLOpt by Johnson (2010).

I use the DIRECT-L (locally-biased dividing rectangles) algorithm of Gablonsky and Kelley (2001) to perform the global optimization. This algorithm employs a branch-and-bound method that progressively subdivides the parameter space into regions in which it suspects the optimum lies based on computations of the function at various points within each rectangle. It retains

information about multiple subdivided rectangles throughout the computational process and does not immediately discard rectangles that seem suboptimal early in the algorithm—a process that can guard against settling into a local optimum. While there is no guarantee that this algorithm will return the global optimum for arbitrary functions (as is the case with all global optimizers), I have found it to work efficiently and return results comparable to those given by much more computationally intensive genetic algorithm. I terminate this global optimizer after 2,500 function evaluations: I have found this number of evaluations to be sufficient to tune the parameters to a region that does not leave most of the data unrationalized, and a local optimizer can move more efficiently to the optimum. Indeed, Johnson (2010) recommends that the termination condition for a global optimizer should be either a limit on runtime or on the number of function evaluations, and using more functional evaluations does not change the output of the final local optimizer appreciably (and indeed often does not change the output of the global optimizer to more than 10^{-6} either).

I then use a local optimizer to polish this solution. I start with the BOBYQA (bound optimization by quadratic approximation) algorithm by Powell (2009) as a derivative-free local optimizer, starting at the optimum found by DIRECT-L. This algorithm utilizes a trust-region method that constructs a quadratic approximation of the objective at each iteration and updates the candidate optimum using this approximation.

A.7.2 Inverting the Characteristic Function

In practice, computing (1.10) is numerically challenging, and I follow suggestions outlined by Krasnokutskaya (2011). Applying (1.10) directly leads to densities that oscillate at the tail, so in practice (i) the integrals are evaluated on a compact interval $[-T, T]$ and (ii) I multiply the integrand by a damping function $d(t) = 1 - \exp(-|t|/T)$. I choose T in a data-driven fashion that is somewhat similar to the method that Krasnokutskaya (2011) proposes, which involves matching moments of the recovered density function with moments computed from the data.

1. I estimate the mean and the standard deviation of the distributions of $\log \Theta$ and \tilde{t} from the data. The mean of $\log \Theta$ is 0, so the mean of the \tilde{t} for a particular (N_1, N_2) is simply the mean of the residuals ν_{ij} for that parameter space. The standard deviation of $\log \Theta$ can be estimated as the standard deviation of the difference $\nu_{i_1j} - \nu_{i_2j}$, divided by $\sqrt{2}$. The standard deviation of \tilde{t} can then be estimated from this estimate together with an estimate for the standard deviation of ν_{ij} .
2. I compute the density for a particular value of T , replace negative values by zeros, and then renormalize the density.
3. I compute the mean and standard deviation of the generated distribution via numerical integration and choose T to minimize the squared deviations from the estimated mean and standard deviations.¹⁵

¹⁵I have noticed through visual inspection that this procedure sometimes still leads to distributions that oscillate

A.7.3 Solving for the Equilibrium Effort Function

When solving the equilibrium model in Section 1.3.2 numerically, I utilize a three-step procedure, fixing a set of parameters for the value and cost distributions as well as the bargaining parameter and the cost of research effort.

1. I solve for the equilibrium effort function in Phase II when $N_2 = 1$ on a fine grid of values. For most instances, I use a grid ranging between 1/3 of the 0.1th percentile of the value distribution to the 3 times the 99.9th percentile of the distribution. I usually use an equally spaced grid with 1000 grid points. For each value v on the grid, I compute the optimal effort using a single-variable optimization routine, such as `fminbnd` in Matlab.
2. For each $1 < N_2 < \bar{N}_2$ (unless $\bar{N}_2 = N_1$, in which case this step applies to $N_2 = \bar{N}_2$ as well), I use an iterated best response procedure to compute the equilibrium Phase II effort function $t_{N_2}^*$. I do so by iterating on (1.2), with $t_{N_2}^*(\cdot)$ in (1.2) replaced by the candidate effort function from the previous iteration of the algorithm. Note that for these values of N_2 , there is no selection, so the choice of p^* is irrelevant, and all types v have the same beliefs over their opponents' surplus. On each step of the iteration, I solve a one-dimensional optimization problem at each grid point using `fminbnd`. I iterate until the maximum change in the effort function is less than 10^{-6} .¹⁶
3. When $N_2 = \bar{N}_2$ and $\bar{N}_2 > 1$ and $\bar{N}_2 < N_1$, the solution method has to account for selection. For a fixed probability p of success in Phase I, I can compute beliefs for all types v from (1.1). I then use the same best response iteration to compute the equilibrium. I do this computation for p on a grid from 0.01 to 0.99.¹⁷
4. In the final step, I compute p^* from the first-order condition associated with (1.4), using a one-dimensional solver such as `fzero` in Matlab. I compute the profits from p not in the grid used in Step 3 by linearly interpolating the computed effort functions.

wildly near the tails, an issue that may be especially relevant using bootstraps. I thus reduce the T estimated for Θ by 10% and the T estimated for \bar{t} by 50% for some additional smoothing.

¹⁶I have noticed that for some parameters, this algorithm tends to fluctuate between two functions instead of converging to one. However, these two functions are always within 10^{-5} of each other, so I simply truncate the algorithm after 30 iterations if it still has not converged. I have found that if instead of updating the effort function to the solution of (1.2) I update it to a convex combination of the solution and the previous iteration, the algorithm is more likely to converge.

¹⁷In practice, I have found it is sufficient to compute the equilibria for $p = \{0.01, 0.1, 0.2, \dots, 0.9, 0.99\}$.

Appendix B

Appendix to Chapter 2

B.1 Numerical Methods to Solve the First-Price Auction

In this section we provide some additional details on how we solve the FPAFE model developed in Section 2.2.

The choice of grid $\{x_i\}_{i=1}^N$ has a minor effect on the solution to the programming problem. Hubbard and Paarsch (2009) use an N -point Gauss-Lobatto grid on $[b, r]$,¹ but we have found in our experiments that an evenly spaced grid is often more efficient. Throughout the article, we use $N = 500$ grid points and $P = 25$ polynomials in the expansion; these choices solve the optimization problem reliably and efficiently.

Bid functions are monotonic in the cost of the bidder, so we follow Hubbard and Paarsch (2009) and impose that $\beta^{*-1}(x_i) \geq \beta^{*-1}(x_{i-1})$ for $2 \leq i \leq N$ when solving the problem in (2.5). Furthermore, we impose the rationality constraint that an agent never bids less than his cost: $\beta^{*-1}(x_i) \leq x_i$ for all i . Finally, we replace the constraint in equation (2.4) by the condition that the two sides cannot differ by more than 0.1% of K . The integral in equation (2.4) is replaced by an approximation using the trapezoidal rule on the N -point grid defined previously.

While we need to calculate bid functions to use in our estimation procedure, we are able to verify the accuracy of our solutions by exploiting the fact that a second price auction, which is easier to solve (Roberts and Sweeting (2013b) use this format to analyze timber auction data), should give the same equilibrium entry thresholds and the same expected procurement cost as a first-price auction with the same parameters. In the many examples that we have looked at, both the entry thresholds and the expected revenues typically agree to several decimal places.

B.2 Details of Sample Construction

De Silva, Dunne, Kankanamge, and Kosmopoulou (2008) study a range of different types of highway procurement auctions from January 1998 through August 2003. However, we only focus on auctions

¹A Gauss-Lobatto grid on $[-1, 1]$ is defined to be the points $y_k = \cos[k\pi/(N-1)]$ for $k \in \{0, 1, \dots, N-1\}$. To define a grid on other intervals, we simply scale these points linearly.

after March 2000 due to an important change in policy in Oklahoma: prior to March 2000 they did not disclose the engineer’s estimate to bidders. Throughout this article, we restrict our attention to auctions for new projects (not re-auctions of failed sales) involving bridge-work with between 4 and 11 planholders. From the remaining set of bridge construction contracts let during this period (483 in OK and 177 in TX), we also employ some additional selection rules removing auctions where (i) there was no winner or the DoT rejected all bids (drops 50 in OK, 1 in TX); (ii) the engineer’s estimate was less than \$100,000 or more than \$5 million (drops 121 in OK, 15 in TX); (iii) the data is incomplete in the sense that it records that the number of recorded planholders differs from the number of suppliers that are listed (drops 1 in OK); and (iv) the winning bid was extremely high or extremely low. Specifically we drop auctions where the winning bid was less than 70% of the engineer’s estimate, which drops 3 auctions in Texas and 47 auctions in Oklahoma. While these auctions have similar observable characteristics to the rest of the sample, 16 of these Oklahoma auctions were won by two companies that only won 4 of the remaining auctions in the sample, suggesting that they may have been somewhat unusual. We also drop 2 auctions in Oklahoma and 4 in Texas with winning bids greater than 150% of the engineer’s estimate, as we assume in estimation that the DoTs set a reserve price at this level even though no reserve price was officially announced for these auctions. After these selection rules, we have a sample of 262 auctions in Oklahoma and 154 in Texas.

B.3 Monte Carlo Studies

In this section we describe some Monte Carlo experiments designed to test the performance of a simulated method of moments estimator where we use importance sampling to approximate the moments. We are particularly interested in how many importance sampling draws per auction are required for accurate estimates (in our actual estimation we use 50). In all of our experiments we create data using the specification in (2.6). We draw the number of potential suppliers N uniformly at random from the set $\{4, 6, 8\}$. We set the vector of observed characteristics for auction a to be $X_a \equiv (1, x_a)$ where x_a is a random draw from the set $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The parameters in the specification are assumed to be $\beta_{\mu_C} = (0, 0.65)'$, $\beta_{\sigma_C} = (0.05, 0)'$, $\beta_{\alpha} = (0.5, 0)'$, $\beta_K = (0.04, 0.05)'$, $\omega_{\mu_C} = 0.05$, $\omega_{\sigma_C} = 0.02$, $\omega_{\alpha} = 0.05$, and $\omega_K = 0.015$. We denote this “true” parameter vector as Γ_0 . The truncation points are taken to be $\underline{c}_{\mu_C} = -0.4$, $\bar{c}_{\mu_C} = 1$, $\underline{c}_{\sigma_C} = 0.005$, $\bar{c}_{\sigma_C} = 0.995$, $\underline{c}_{\alpha} = 0.1$, $\bar{c}_{\alpha} = 0.9$, $\underline{c}_K = 0.0001$, and $\bar{c}_K = 0.16$.² We truncate the cost distribution to $[0, 4.75]$, where the probability that a cost draw will be above the upper support is almost always negligible for our draws of the parameters, and we set the reserve price equal to 4.75. For any set of parameters drawn from this distribution, we can solve for the equilibrium bid functions and entry decisions, and simulate outcomes.

As in our empirical application, we estimate the parameters by matching a set of moments describing the outcomes from auctions with their expectations, calculated using importance sampling, with separate moments for each $\{N, x_a\}$ combination. Specifically, we match the distribution of

²These truncation points are chosen to ensure that we can solve the FPAFE model accurately.

the number of entrants (i.e., indicator variables for whether one, two, three, etc. firms enter), the distribution of winning bids (using indicator variables for whether the winning bid is in a particular bin, where we divide the interval $[0.4, 2.65]$ into bins of width 0.075), the distribution of all bids (we divide the interval $[0.4, 4.2]$ into bins of width 0.1) and the number of suppliers that enter but do not submit bids because they have valuations greater than the reserve. We estimate the parameters by minimizing the squared sum of the differences between the observed and predicted moments, where each moment is weighted equally.

We report the results from three experiments. In each experiment we use 500 auctions as data, but we vary the importance sampling density and the number of simulated draws we use for each auction. Table B.1 shows the mean and the standard deviation of each parameter when we run each experiment 100 times.

In the first experiment (A), we use the true distribution of the parameters as the importance sampling density, which is optimal but obviously infeasible in an empirical application where the goal is to estimate the true parameters, and use five simulations for each auction.³ In the second experiment (B), we make the importance sampling density more diffuse, by setting $\omega_{\mu_C} = 0.1$, $\omega_{\sigma_C} = 0.04$, $\omega_{\alpha} = 0.1$, and $\omega_K = 0.03$, and other parameters remain the same. We continue to use 5 simulations per auction.⁴ In the final experiment (C), we use the same importance sampling density as in (B) but we use 30 draws per auction. In all of our experiments we use the true parameters to begin our search but we have verified that using a wide-range of starting values does not change our estimates because the objective function is generally very well behaved. The results of these experiments are shown in columns (A)–(C) of Table B.1, respectively.

In all three experiments we recover the coefficients accurately, and, when we use a more diffuse importance sampling density and 30 draws (C), we recover the parameters almost as well as when we use the true distribution of the parameters (A). There is some evidence that we tend to overestimate the variance ω parameters in (B) when we use the diffuse density and only 5 draws, but even in this case, the magnitude of any bias appears to be quite small.

Based on these results, in our empirical application we choose to use 50 draws for each auction. We also choose the importance sampling density based on a large number of initial runs, using quite diffuse importance sampling densities. This indicated that the parameters of the importance sampling density that we actually use (see Table 2.2) were likely to be close to the parameters that we would end up estimating.

³Note rather than solving a fresh set of 2,500 auctions each time we perform a separate run of the experiment, we solve 25,000 auctions drawn from the distribution up front, and then draw, with replacement, from this pool when creating our data and importance sampling draws for each run.

⁴In this case, we created solved 75,000 auctions and draw from this pool for each run.

Parameter	Variable	True Value	Experiment		
			A	B	C
Location Parameter (μ_C)	Constant	0	0.0000 (0.0100)	0.0015 (0.0186)	-0.0043 (0.0170)
	x_a	0.65	0.6507 (0.0174)	0.6533 (0.0282)	0.6602 (0.0247)
	ω	0.05	0.0488 (0.0044)	0.0575 (0.0073)	0.0504 (0.0059)
Scale Parameter (σ_C)	Constant	0.05	0.0494 (0.0074)	0.0463 (0.0152)	0.0416 (0.0126)
	x_a	0	-0.0002 (0.0113)	0.0085 (0.0204)	0.0195 (0.0194)
	ω	0.02	0.0192 (0.0035)	0.0295 (0.0094)	0.0225 (0.0077)
Degree of Selection (α)	Constant	0.5	0.4994 (0.0205)	0.4985 (0.0354)	0.4984 (0.0231)
	x_a	0	0.0000 (0.0308)	-0.0102 (0.0607)	-0.0130 (0.0571)
	ω	0.05	0.0479 (0.0078)	0.0803 (0.0190)	0.0561 (0.0164)
Entry Cost (K)	Constant	0.04	0.0388 (0.0047)	0.0361 (0.0118)	0.0387 (0.0085)
	x_a	0.05	0.0516 (0.0074)	0.0559 (0.0173)	0.0529 (0.0121)
	ω	0.015	0.0143 (0.0025)	0.0209 (0.0063)	0.0145 (0.0038)

Table B.1: Results of Monte Carlo experiments. Each column reports the mean and standard deviation of the parameter estimates from 100 runs, where there are 500 auctions used as data in each run. In column A the importance sampling density is the true distribution and we use 5 simulations per auction in the data. In columns B and C we use a more diffuse importance sampling density (see text) and either 5 (B) or 30 (C) simulations per auction.

Appendix C

Appendix to Chapter 3

C.1 Formally Adding Ambiguity Aversion

In this appendix, we begin with a standard equilibrium concept to solve the game presented in Section 3.2.1. We then add ambiguity aversion to this setup and show that it leads to the concept of distortion equilibrium as we defined it in Section 3.2.2.

C.1.1 Sequential Equilibrium

In this subsection, we will consider a sequential equilibrium in this game (Kreps and Wilson (1982)), which consists of a set of mixed strategies at each information set along with beliefs over nodes at each set such that actions are sequentially rational, beliefs are formed by Bayes' rule whenever it applies, and beliefs are consistent (in the sense defined by Kreps and Wilson) for information sets that have probability zero in equilibrium. Denote by $p_i(h_i) \in [0, 1]$ the probability that player i accepts μ over \hat{y} at information set h_i , and let $\nu_i(a_{-i}, O|h_i)$ be i 's belief that his opponent played a_{-i} and the order is O , conditional on being at information set h_i . To simplify notation, consider only player 1 and let $\nu(1|h)$ be his belief that he is picking first, at information set h . The expected utility to player 1 from accepting the alternate distribution μ is

$$U_1(\mu|h) = \nu(1|h)\chi(\mu, h) + (1 - \nu(1|h))\mathbb{E}_{\mu[y]}w_1(y), \quad (\text{C.1})$$

where $\chi(1, h)$ is the expected profit of player 1 from choosing μ at information set h , conditional on being first to alter the signal. We can write

$$\chi(\mu, h) = \sum_{a_2 \in A_2} \frac{\nu(a_2, 1|h)}{\nu(1|h)} \mathbb{E}_{\mu[y']} \mathbb{E}_{F(a_2, y')|\mu'} \{ w_1(\mu') p_2(a_2, y', \mu') + w_1(y') (1 - p_2(a_2, y', \mu')) \}, \quad (\text{C.2})$$

where $w_1(\mu')$ is shorthand for $\mathbb{E}_{\mu'[y'']}w_1(y'')$. Of course, the expected utility to player 1 from accepting the temporary signal \hat{y} is simply $U_1(\delta_{\hat{y}}|h)$ from (C.1), where $\delta_{\hat{y}}$ is a distribution that places probability 1 on \hat{y} .

Sequential rationality for player 1 at these information sets, therefore, requires that

$$p_1(a_1, y', \mu) = \begin{cases} 1 & \text{if } U_1(\mu|h) > U_1(\delta_{\hat{y}}|h) \\ [0, 1] & \text{if } U_1(\mu|h) = U_1(\delta_{\hat{y}}|h) \\ 0 & \text{if } U_1(\mu|h) < U_1(\delta_{\hat{y}}|h) \end{cases} \quad (\text{C.3})$$

Bayesian updating requires that

$$\begin{aligned} \nu(a_2, 1|a_1, \hat{y}, \mu) &= \pi(a_1, a_2)(\hat{y}) \cdot dF(\hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a_2)\gamma \left/ \left[\sum_{a'_2 \in A_2} \pi(a_1, a'_2)(\hat{y}) \cdot dF(\hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a'_2)\gamma \right. \right. \\ &\quad \left. \left. + \int_{\mu'} \sum_{a'_2 \in A_2} \pi(a_1, a'_2)(\hat{y}) \cdot dF(\hat{y})(\mu') \cdot (1 - p_2(a_2, \hat{y}, \mu')) \cdot dF(a_2, \hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a'_2)(1 - \gamma) \right. \right. \\ &\quad \left. \left. + \int_{\mu'} \sum_{y' \in Y} \sum_{a'_2 \in A_2} \pi(a_1, a'_2)(y') \cdot dF(y')(\mu') \cdot p_2(a_2, y', \mu') \cdot \mu'(\hat{y}) \cdot dF(a_2, \hat{y})(\mu) \cdot \sigma_1(a_1)\sigma_2(a'_2)(1 - \gamma) \right] \right. \quad (\text{C.4}) \end{aligned}$$

This expression is extensive, but it captures that the game can arrive at the information set (a_1, \hat{y}, μ) in three ways: (i) Player 1 could be first and have been offered the temporary signal \hat{y} and the alternate μ , (ii) Player 1 could be second, and \hat{y} may be the signal because Player 2 refused to change it, and (iii) Player 1 could be second, and \hat{y} may be the signal because it was the draw from Player 1's alternate signal. The probability of reaching this node is zero if $\sigma_1(a_1) = 0$ or if π and the F 's are such that either \hat{y} or μ has probability zero following the action a_1 . If the probability of this node is zero due to $\sigma_1(a_1) = 0$, we follow Kreps and Wilson and define $\nu(a_2, 1|a_1, \hat{y}, \mu)$ using consistency. If certain nodes are not reached due to π or the F 's, then we can simply ignore them in constructing the equilibrium.

Finally, the mixed strategies $\sigma_i \in \Delta A_i$ chosen in the first stage of the game (i.e., when playing G) must be optimal given the future path of play (i.e., p_1 and p_2 defined in (C.3)). and given the opponents' strategy $\sigma_{-i} \in \Delta A_{-i}$. Together with the conditions (C.3), (C.4), and the consistency requirement from Kreps and Wilson (1982), we have a sequential equilibrium of the game.

C.1.2 Adding Ambiguity Aversion

We now add ambiguity-aversion to the payoffs presented in the previous subsection. Throughout this derivation, restrict attention to signal structures π such that $\pi(a)$ has full support on Y ; we will briefly discuss relaxing this restriction at the end of the subsection. As mentioned in Section 3.2.2, we assume that agents only know that $\gamma \in [0, 1]$, that $\xi \in [0, 1]$, and that agents believe each F is an element of $\Delta(\Delta(Y))^0$, or the set of distributions over ΔY that have full support.

Such an agent, therefore, would evaluate payoffs according to

$$\tilde{U}_1(\mu|h) = \inf_{\substack{\gamma \in [0, 1] \\ F \in \Delta(\Delta(Y))^0}} \nu(1|h)\chi(\mu, h) + (1 - \nu(1|h))\mathbb{E}_{\mu[y]}w_1(y), \quad (\text{C.5})$$

where $\nu(1|h)$ and $\chi(1|h)$ are defined as in (C.2) and (C.4), and $p_2(\cdot)$ is taken as fixed. By F in the infimum we refer to all three different distributions, $F, F(a_1, \cdot), F(a_2, \cdot)$ over alternate signal distributions. Note that we have to take the infimum instead of the minimum since $\Delta(\Delta(Y))^0$ is not a closed set. However, because F is an element in $\Delta(\Delta(Y))^0$ and we are currently restricting attention to π such that $\pi(a)$ has full support on Y , we can ensure that all combinations of (\hat{y}, μ) have positive probability following every action profile and every possible ordering.

We are searching for an equilibrium in which (i) p_i is determined by sequential rationality (as in (C.3), with U replaced by \tilde{U}), (ii) p_i is either 0 or 1, that is, for simplicity we restrict our attention to pure strategies in the distortion part of the game, p_i is independent of a_i ,¹ and finally (iii) the actions in the initial game G are optimal given the value of \tilde{U}_i . In the remainder of this subsection, we show how these assumptions motivate the concept of distortion equilibrium (Section 3.2.2).

Note than, in equilibrium, agent 2 will weakly prefer to distort the signal in favor of an alternate signal distribution at a history $h_2 = (a_2, y', \mu')$ if $\tilde{U}_2(\mu'|h_2) \geq U_2(\delta_{y'}|h_2)$. Let us denote the set

$$D_2(y') = \{\mu' \in \Delta(Y) : \tilde{U}_2(\mu'|h_2) \geq U_2(\delta_{y'}|h_2)\}, \quad (\text{C.6})$$

as the set of all μ' that player 2 will weakly prefer to $\delta_{y'}$. Note that this set is independent of a_2 . Finally, we restrict ourselves to a particular tie-breaking rule for all alternate signal distributions $\mu' \neq \delta_{y'}$: we assume that both agents, when indifferent between distorting and not distorting, choose to distort the signal.

To obtain an explicit characterization of the equilibrium, we begin by computing the value of $\tilde{U}_1(\mu|h)$ in (C.5). This is a convex combination of the set of terms

$$\left\{ \left\{ \mathbb{E}_{\mu[y']} \mathbb{E}_{F(a_2, y')|\mu'} [w_1(\mu') p_2(a_2, y', \mu') + w_1(y') (1 - p_2(a_2, y', \mu'))] \right\}_{a_2}, \mathbb{E}_{\mu[y]} w_1(y) \right\}. \quad (\text{C.7})$$

We have that

$$\min_{\mu' \in D_2(y')} w_1(\mu') \leq w_1(y') \quad (\text{C.8})$$

since by construction $\delta_{y'} \in D_2(y')$. Thus, we can minimize the elements of the set in (C.7) term-by-term by setting

$$F(a_2, y') \equiv \delta \left\{ \arg \min_{\mu' \in D_2(y')} w_1(\mu') \right\}, \quad (\text{C.9})$$

where $\delta\{\mu\}$ is the distribution that places full mass on μ .

We now regard the global infimum in (C.5). From (C.9) and the tie-breaking rule that $\{\mu' : p_2(a_2, y', \mu') = 1\}$ is identically $D_2(y')$, we have that the global infimum (over all allowed F and γ) in (C.5) is no less than

$$\min \left\{ \mathbb{E}_{\mu[y']} \min_{\mu' \in D_2(y')} w_1(\mu'), \mathbb{E}_{\mu[y]} w_1(y) \right\} = \mathbb{E}_{\mu[y']} \min_{\mu' \in D_2(y')} w_1(\mu'), \quad (\text{C.10})$$

¹We believe this can be relaxed, so that the *only* equilibrium is one in which p_i does not depend on a_i , but it doesn't seem straightforward to show. We realize that assumptions are piling up, so it might be useful to drop this condition as one we are imposing on the equilibrium and instead it is a property of the equilibrium.

where the equality follows from (C.8) again. It then suffices to argue that we can find γ such that, together with the choice of $F(a_2, y')$ in (C.9), achieves this lower bound. It is easy to see from (C.4) that setting $\gamma = 1$ achieves this bound as long as $\pi(a)$ has full support for all a , and (C.10) gives an expression for $\tilde{U}_1(\mu|h)$ from (C.5),

$$\tilde{U}_1(\mu|h) = \mathbb{E}_{\mu[y']}\min_{\mu' \in D_2(y')} w_1(\mu'). \quad (\text{C.11})$$

We can apply similar logic to compute the payoff from a (possibly mixed) action α_1 in the initial game G . This payoff is given by

$$V_1(\alpha_1, \alpha_2) = \inf_{\substack{\xi \in [0,1] \\ F(\hat{y}) \in \Delta(\Delta Y)^0}} (1 - \delta)g_1(\alpha) + \delta \cdot \mathbb{E}_{\pi(\alpha)[\hat{y}]} \left[(1 - \xi) \cdot w_1(\hat{y}) \right. \\ \left. + \xi \cdot \mathbb{E}_{F(\hat{y})[\mu]} \max\{\tilde{U}_1(\mu|(\alpha_i, \hat{y}, \mu)), \tilde{U}_1(\delta_{\hat{y}}|(\alpha_i, \hat{y}, \mu))\} \right],$$

where we already used the fact that agent 1 believes to go first ($\gamma = 1$).² Similar computations show that this expression is minimized when $\xi = 1$ and $F(\hat{y})$ is such that the second element in the max exceeds the first. In light of (C.11) this lets us write

$$V_1(\alpha_1, \alpha_2) = (1 - \delta)g_1(\alpha) + \delta \cdot \mathbb{E}_{\pi(\alpha)[\hat{y}]} \min_{\mu' \in D_2(\hat{y})} w_1(\mu'). \quad (\text{C.12})$$

All the expressions are symmetric for player 2.

We have thus far restricted attention to signal structures π with full support. The role of this restriction (along with the assumption that $F(a_2, y) \in \Delta(\Delta Y)^0$) is to allow player 1 to hold the belief that he is going first with probability 1 to be rationalizable. Without a full support assumption, it may be possible for player 1 to be certain that he is going second if he is offered a signal \hat{y} that has probability zero given his action a_1 in the game G —and Bayesian updating beliefs about the order of distortion will be ill-defined if $\gamma = 1$. We can instead derive a distortion equilibrium for games where the signal structure π does not have full support as a limit of distortion equilibria for games with the same payoffs but signal structures π'_n that do have full support and converge to π . Such an equilibrium would be identical to the one presented in Section 3.2.2. We omit the details, but the rationale is that the only role of π in the expressions above is in (C.12), and that expression is continuous in π . In this chapter we apply the definition in that section to games with signal structures that do not satisfy this full support assumption.

C.2 Extension to N Players

In this section, we define a distortion equilibrium for $N > 2$ players. We then prove the analogue of Theorem 3.1 for a case with $N = 3$ players, which involves a direct extension of arguments presented in Section 3.3.1.

²The notation $\tilde{U}_1(\mu|(\alpha_i, \hat{y}, \mu))$ is shorthand for the expectation $\mathbb{E}_{\alpha_i[\alpha_i]} \tilde{U}_1(\mu|(a_i, \hat{y}, \mu))$.

A distortion equilibrium as defined in Section 3.2.2 involves strategies for the first stage game along with distortion strategies for how players are willing to alter the signal, with these distortion strategies being disciplined by consistency in Definition 3.1. Implicit in this formulation is some notion of what sort of distortions players fear from their opponents. In the case of $N = 2$ players, there is not much choice in how to model this set: we simply assume that each player i fears that their opponents will distort after them and thus choose the signal y may yield any distortion in $D_{-i}(y)$.³ However, for $N \geq 3$ players, this modeling choice is not obvious. Each player may worry that *exactly* one player will be able to distort after him, but he may be unsure which player it is. Each player may worry that all other players will be able to distort after him but may be uncertain about the sequence. Players may worry that all remaining $N - 1$ players will be able to distort the signal jointly.

We will add this modeling choice to the definition of the game. Note that in the $N = 2$ -player setup of Section 3.2, the game could be summarized by the triple (G, w, π) of first-stage payoffs $G : A \rightarrow \mathbb{R}^N$, continuation payoffs $w : Y \rightarrow \mathbb{R}^N$, and a signal distribution $\pi : A \rightarrow \Delta Y$. In the general N -player setup, the primitives of the game will be a quadruple (G, w, π, Γ) . The new component $\Gamma : \mathcal{D}^N \rightarrow \mathcal{D}^N$ takes in the players' distortion strategies and outputs the set of future distortions each player fears. For instance, in the baseline model in Section 3.2 with $N = 2$, we would have $\Gamma(D_1, D_2) = (D_2, D_1)$. The following are natural examples of Γ for $N > 2$.

- I. Suppose each player fears that exactly one player will be able to distort after him but is unsure who. Then, $\Gamma_i(D)(y) = \bigcup_{j \neq i} D_j(y)$. This is because any μ that *some* player will accept in lieu of y is a possible future distortion from y that player i fears.
- II. Suppose $N = 3$ and each player fears that exactly two players will distort and is sure about the order. For concreteness, say that player 1 believes with certainty that player 2 will get the chance to distort after him, and player 3 will distort after player 2. Then, the set of distortions that player 1 will fear is that of all possible *paths* of distortions. That is, $\Gamma_1(D)(y)$ is the set of all $\mu \in \Delta Y$ such that there exists $\mu_2 \in D_2(y)$ and $\mu_3(y') \in D_3(y')$ for all y' such that μ is the product of the row vector that represents μ_2 and the stacked row vectors that represent $\mu_3(y')$ for all y' .⁴ Define $(D_2 * D_3)(y)$ as the value of $\Gamma_1(D)(y)$ in this case. (To understand this set better, note that $D_2(y)$ and $D_3(y)$ are both obviously subsets of $(D_2 * D_3)(y)$.)
- III. If $N = 3$ and each player fears that exactly two players will distort but is unsure about the order, then $\Gamma_1(D)(y) = (D_2 * D_3)(y) \cup (D_3 * D_2)(y)$.
- IV. If each player fears that the remaining $N - 1$ players will all jointly distort a signal after he does, then $\Gamma_i(D)(y) = \bigcap_{j \neq i} D_j(y)$. This is because any μ that player i fears must be preferred

³As we note in Section 3.3.3, there is *some* degree of flexibility in how to model the distortion phase, since one may imagine a model in which players can distort multiple times. However, as long as each player fears a potential future distortion from his opponent at all times, nothing of import changes.

⁴We can think of this as a two-step non-homogenous Markov process on signals y , where player 2's distortion μ_2 gives the first transition probability vector and the a selection of distortions $\mu_3(y')$ for player 3 gives the second transition matrix.

to y by *all* of his opponents.

Of course, it is simple to consider models where players fear multiple such possibilities by taking unions (just as for moving from Case II to Case III). However, some cases are clearly less restrictive than others. If player i fears both a single opponent distorting the signal and the possibility of all opponents jointly deciding to distort the signal, we would set $\Gamma_i(D)(y) = \bigcup_{j \neq i} D_j(y) \cup \bigcap_{j \neq i} D_j(y) = \bigcup_{j \neq i} D_j(y)$. This is expected, since collusion allows for fewer distortions than under individual distortions. Finally, note that the relevant Γ for each of the above cases can be derived by starting from an extensive form game like in Figure 3-1 and adding ambiguity aversion.

With this formulation in mind, we modify the definitions from Section 3.2.2.

Definition 3.1' (Consistency). *A triple (w, \tilde{w}, D) is said to be consistent with respect to Γ if*

$$\begin{aligned} \tilde{w}_i(y) &= \min_{\mu \in \Gamma_i(D)(y)} \mathbb{E}_{\mu[y']} w_i(y) \text{ and} \\ D_i(y) &= \{\mu \in \Delta(Y) : \mathbb{E}_{\mu[y']} [\tilde{w}_i(y')] \geq \tilde{w}_i(y)\}. \end{aligned}$$

We can define distortion equilibrium analogously to before as well.

Definition 3.2' (Distortion Equilibrium). *A strategy profile (α, D) is a distortion equilibrium in the game (G, w, π, Γ) if*

- (i) *the triple (w, \tilde{w}, D) is consistent with respect to Γ as in Definition 3.1'; and*
- (ii) *for each player, α_i is optimal given α_{-i} and \tilde{w}_i , meaning for all $a_i \in \alpha_i$,*

$$a_i \in \arg \max_{a'_i \in A_i} (1 - \delta) g_i(a'_i, \alpha_{-i}) + \delta \cdot \mathbb{E}_{\pi_y(a'_i, \alpha_{-i})[y]} \tilde{w}_i(y).$$

We now restrict to the case where $N = 3$ and explore the analogue of Theorem 3.1 for the possibilities for Γ presented above. We include the (somewhat repetitive) proofs to highlight that arguments very similar to ones in Section 3.3.1 apply. We leave an analysis of $N > 3$ for future work.

Lemma 3.1'. *Suppose $\tilde{w}(y)$ and D satisfy consistency and Γ satisfies any of Cases I–IV above. Then, the $\tilde{w}(y)$ are strongly Pareto-ranked in that $\tilde{w}_i(y') \geq \tilde{w}_i(y'')$ if and only if $\tilde{w}_j(y') \geq \tilde{w}_j(y'')$ for any players i and j .*

Proof. Note that for any \tilde{w} and pair of signals, at least two players must agree on the relative rankings of the signals. Thus, suppose without loss of generality that $\tilde{w}_1(y') \leq \tilde{w}_1(y'')$ and $\tilde{w}_2(y') \leq \tilde{w}_2(y'')$; Then, we have that $D_1(y'') \subseteq D_1(y')$ and $D_2(y'') \subseteq D_2(y')$. Then, it is easy to check that in each of Cases I–IV, we have $\Gamma_3(D)(y'') \subseteq \Gamma_3(D)(y')$. This implies that

$$\tilde{w}_3(y') = \min_{\mu \in \Gamma_3(D)(y')} \mathbb{E}_{\mu[y']} w_3(y) \leq \min_{\mu \in \Gamma_3(D)(y'')} \mathbb{E}_{\mu[y']} w_3(y) = \tilde{w}_3(y''),$$

as needed. □

Lemma 3.2'. *Suppose $\tilde{w}(y)$ and D satisfy consistency and, among the $\tilde{w}(y)$, we have a unique Pareto-best point \tilde{w}_B such that $\tilde{w}_B \geq \tilde{w}(y)$ for all y and a unique Pareto-worst point \tilde{w}_W such that $\tilde{w}_W \leq \tilde{w}(y)$ for all y . Suppose further that Γ satisfies any of Cases I–III above. Then, the $\tilde{w}(y)$ can be expressed as $\tilde{w}(y) = c + t(y) \cdot d$, where $c \in \mathbb{R}^N$, $d \in \mathbb{R}^N$ with $d_i > 0$ for all i , and $t(y) \in \mathbb{R}$ for all y . That is, \tilde{w} lie on a line with strictly positive slope.*

Proof. The proof proceeds almost exactly like the one to Lemma 3.2. If $\tilde{w}_{W,i} = \tilde{w}_{B,i}$ for any i , then Lemma 3.1' implies that it holds for all i .

Now suppose that $\tilde{w}_W < \tilde{w}_B$, and as in the proof of Lemma 3.2, let $Y_B \equiv \{y \in Y : \tilde{w}(y) = \tilde{w}_B\}$ and $Y_W \equiv \{y \in Y : \tilde{w}(y) = \tilde{w}_W\}$. Consider the space of $(\tilde{w}_1(y), \tilde{w}_2(y))$ and consider the line ℓ connecting \tilde{w}_W to \tilde{w}_B in this plane. Suppose that there exists \hat{y} such that the project of $\tilde{w}(\hat{y})$ lies above ℓ in this plane. Note that for all i and $y_B \in Y_B$, we have that $D_i(y_B) = \Delta Y_B$, and for all i and $y_W \in Y_W$, we have $D_i(y_W) = \Delta Y$. In Cases I–III, this means that $\Gamma_i(D)(y_B) = \Delta Y_B$ for all $y_B \in Y_B$, and $\Gamma_i(D)(y_W) = \Delta Y$ for all $y_W \in Y_W$ too. Then, as in Footnote 10, we have that there exists y_B^* and y_W^* such that $w_2(y_B^*) = \tilde{w}_{B,2}$ and $w_2(y_W^*) = \tilde{w}_{W,2}$. Find the $\alpha \in [0, 1]$ such that $\tilde{w}_1(\hat{y}) = \alpha \tilde{w}_1(y_B^*) + (1 - \alpha) \tilde{w}_1(y_W^*)$ and note that this means $\mu \equiv \alpha \delta_{y_B^*} + (1 - \alpha) \delta_{y_W^*} \in D_1(\hat{y})$. For Cases I–III, this means that $\mu \in \Gamma_2(D)(\hat{y})$ as well. But, $\tilde{w}_2(\hat{y}) > \alpha \tilde{w}_2(y_B^*) + (1 - \alpha) \tilde{w}_2(y_W^*) = \alpha w_2(y_B^*) + (1 - \alpha) w_2(y_W^*)$, which is a contradiction to the definition of $\tilde{w}_2(\hat{y})$ as $\min_{\mu \in \Gamma_2(D)(\hat{y})} \mathbb{E}_{\mu[y]} w_2(y)$. Considering points to the right of ℓ leads to a similar contradiction. This shows that the projection of the points $\tilde{w}(y)$ onto the plane containing payoffs for players 1 and 2 lies on a strictly positive line. Since players 1 and 2 were arbitrary, \square

Aggregating, we have the following analogue of Theorem 3.1.

Theorem 3.1'. *If $N = 3$ and Γ is as in Cases I–IV, then consistency of the triple (w, \tilde{w}, D) with respect to Γ requires that $\tilde{w}(y)$ are strongly Pareto-ranked. If Γ is as in Cases I–III, then they lie on a line of strictly positive slope.*

Note that we have not explored the most general conditions for Γ under which Lemmas 3.1' and 3.2' hold. For instance, Lemma 3.2' will hold for generalizations of Cases I–III for $N > 3$ as well. Instead, we have chosen to show the result for various natural assumptions on the distortion stage. For the case of $N = 3$, Cases I–IV (and combinations of thereof) likely exhaust all reasonable assumptions on this stage. The one case in which our linearity result may not hold (and incentives are instead simply Pareto-ranked) is Case IV. This case, however, is in some sense the one that embodies the lowest degree of fear of signal distortion since collusion places constraints on how opponents can decide to distort the signal. As long as each player fears *some* possibility of individual distortions by his opponents, we can show a linearity result.

C.3 Omitted Proofs

C.3.1 Proof of Theorem 3.2

Proof. This proof has two pieces. In the first step, we show that an RDE is a PPED. Let the RDE be denoted $(\alpha(y^{t-1}), D(y^{t-1}), w(y^{t-1}))$. Define $\mathcal{C}'(y^t)$, $\mathcal{S}'_i(y^t)$, and $\sigma'_i(y^t)$ as the one-period analogues of \mathcal{C} , \mathcal{S}_i , and σ_i from Section 3.5.⁵ We claim that w satisfies the difference equation

$$v_i(y^{t-1}) \equiv w_i(y^{t-2})(y_{t-1}) = \min_{\nu \in \mathcal{S}'_i(y^{t-1})} (1 - \delta)g_i(\alpha(y^{t-1})) + \delta \mathbb{E}_{\nu[y]}[w_i(y^{t-1})(y)]. \quad (\text{C.13})$$

It suffices to show that

$$\min_{\nu \in \mathcal{S}'_i(y^{t-1})} \mathbb{E}_{\nu[y]}[w_i(y^{t-1})(y)] = \mathbb{E}_{\pi(\alpha(y^{t-1}))}[y]}[\tilde{w}_i(y^{t-1})(y)].$$

Note that by consistency and the definition of \tilde{w} , we have

$$\begin{aligned} & \mathbb{E}_{\pi(\alpha(y^{t-1}))}[y_t]}[\tilde{w}_i(y^{t-1})(y_t)] \\ &= \mathbb{E}_{\pi(\alpha(y^{t-1}))}[y_t]} \left[\min_{\mu_i \in D_i(y^{t-1})(y_t)} \mathbb{E}_{\mu_i[\hat{y}]} [\tilde{w}_i(y^{t-1})(\hat{y})] \right] \\ &= \mathbb{E}_{\pi(\alpha(y^{t-1}))}[y_t]} \left[\min_{\mu_i \in D_i(y^{t-1})(y_t)} \mathbb{E}_{\mu_i[\hat{y}]} \left[\min_{\mu_{-i} \in D_{-i}(y^{t-1})(\hat{y})} \mathbb{E}_{\mu_{-i}[\hat{y}]} [w_i(y^{t-1})(\hat{y})] \right] \right] \\ &= \min_{\{\mu_i \in D_1(y^{t-1})(y_t)\}_{y_t \in Y}} \min_{\{\mu_{-i} \in D_{-i}(y^{t-1})(\hat{y})\}_{\hat{y} \in Y}} \mathbb{E}_{\pi(\alpha(y^{t-1}))}[y_t]} \left[\mathbb{E}_{\mu_i[\hat{y}]} \left[\mathbb{E}_{\mu_{-i}[\hat{y}]} [w_i(y^{t-1})(\hat{y})] \right] \right] \\ &= \min_{\nu \in \mathcal{S}'_i(y^{t-1})} \mathbb{E}_{\nu[y]}[w_i(y^{t-1})(y)]. \end{aligned}$$

Hence, w from the RDE satisfies equation (C.13). Expanding out the difference equation gives

$$\begin{aligned} v_i(y^t) &= \min_{\nu \in \mathcal{S}'_i(y^t)} (1 - \delta)g_i(\alpha(y^t)) + \delta \mathbb{E}_{\nu[y]}[w_i(y^t)(y)] \\ &\quad \vdots \\ &= \min_{\nu \in \mathcal{S}'_i(y^t)} (1 - \delta)g_i(\alpha(y^t)) + \mathbb{E}_{\nu} \left[(1 - \delta) \sum_{s=1}^{T-1} \delta^s g_i(\alpha(y^{t+s})) + \delta^T w_i(y^{t+(T-1)})(y^{t+T}) \right]. \end{aligned}$$

Since $w_i(y^{t+s-1})(y)$ is bounded, letting $T \rightarrow \infty$ clearly shows that

$$v_i(y^t) = w_i(y^{t-1})(y_t) = \min_{\nu \in \mathcal{S}'_i(y^t)} \left\{ (1 - \delta)g_i(\alpha(y^t)) + (1 - \delta) \mathbb{E}_{\nu} \left[\sum_{s=1}^{\infty} \delta^s g_i(\alpha(y^{t+s})) \right] \right\}.$$

Hence the continuation values from the RDE match up with the computed continuation values from the associated PPED. Incentive compatibility and consistency of w as a PPED follow immediately

⁵That is, only consider choices μ_i from $D_i(\cdot)(\cdot)$ for a single period, and the induced distribution is only over the next period's signal.

from conditions (i)–(iii) in the definition of RDE. Hence we have shown that an RDE is also a PPED.

Now we show the other direction. Let $v(y^t)$ be the PPED payoffs from (3.4) and let $w(y^{t-1})(y_t)$ be the associated continuation payoffs. All that is required is to show that w satisfies the difference equation in (C.13). Noting that the minimization in (3.4) is recursive, it is clear that

$$v_i(y^t) = w_i(y^{t-1})(y_t) = \min_{\nu \in \mathcal{S}_i(y^t)} \left\{ (1 - \delta) \cdot g_i(\alpha(y^t)) + \delta \cdot \mathbb{E}_{\nu[y]} [w_i(y^t)(y)] \right\}^6$$

Now by reversing the argument that led us to (C.13), we see that

$$\begin{aligned} v_i(y^t) &= w_i(y^{t-1})(y_t) = \min_{\nu \in \mathcal{S}_i(y^t)} \left\{ (1 - \delta) \cdot g_i(\alpha(y^t)) + \delta \cdot \mathbb{E}_{\nu} [w_i(y^t)(y)] \right\} \\ &= (1 - \delta) \cdot g_i(\alpha(y^t)) + \delta \cdot \mathbb{E}_{\pi(\alpha(y^t))} [\tilde{w}_i(y^t)(y)]. \end{aligned}$$

Consistency and incentive compatibility in the definition of RDE follow immediately from the corresponding conditions in the definition of PPED. Thus we have shown that a PPED is also an RDE. It follows that the RDE and PPED formulations are equivalent. \square

C.3.2 Proofs in Section 3.6.1

Proof of Lemma 3.3. We show that any $v \in W$ can be supported as an RDE. Since $v \in W$, there exists (α, w, D) that P-enforces v . In particular,

$$v = (1 - \delta)g(\alpha) + \delta \mathbb{E}_{\pi(\alpha)} [\tilde{w}(y)]$$

with $w(y) \in W$. Since W is self-generating, for each $v(y_0) = w(y_0)$ there exists $(\alpha(y^0), w(y^0), D(y^0))$ with $w(y_0)(y_1) \in Y$ which enforces $v(y_0)$. Iterating this process forward, we will have a sequence $\{(\alpha, w(y^t), D(y^t))\}_{t \geq 1}$. Since $v(y^{t-1})$ is P-enforced by $(\alpha, w(y^{t-1}), D(y^{t-1}))$ and W is a bounded set (so that the $w(y^t)$ are bounded), this exactly shows that the constructed sequence $\{(\alpha, w(y^t), D(y^t))\}_{t \geq 1}$ is an RDE and thus a PPED by Theorem 3.2. \square

Proof of Lemma 3.4. The fact that $E_P(\delta) \subseteq B_P(\delta, E_P(\delta))$ follows trivially from the recursive structure of the PPED. In particular, after any y_0 the continuation strategies are again a PPED. Thus $v(y_0) \in E_P(\delta)$ for all $y_0 \in Y$. Thus it is clear that $E_P(\delta) \subseteq B_P(\delta, E_P(\delta))$.

Now note that if $v \in B_P(\delta, E_P(\delta))$, then it is enforceable with continuation payoffs that are from PPEDs and thus is clearly a PPED. Thus, $E_P(\delta) \supseteq B_P(\delta, E_P(\delta))$ so that we have the equivalence. \square

Proof of Theorem 3.3. The proof of this theorem follows that of Theorem 3.1 in FL very closely. First, we prove that $E_P(\delta) \subseteq Q_P$. Suppose it does not hold. Then, we can find $\lambda \in \mathbb{R}^2$, $\delta \in (0, 1)$, and $v \in E_P(\delta)$ so that $\lambda \cdot v \equiv k > k_A^*(\lambda)$. Let $E_P^*(\delta)$ be the convex hull of $E_P(\delta)$. Let \bar{k} be the

⁶Note that an element of \mathcal{S}_i gives a distribution over all future public signals, but for the purposes of this equation, we only worry about the distribution it induces over the next period's signal.

maximal score in $E_P^*(\delta)$ in direction λ . Note that $\bar{k} > k_P^*(\delta)$. Since the score function is linear, there exists $v \in E_P(\delta)$ with $\lambda \cdot v = \bar{k}$. Then by the fact that $E_P(\delta) \subseteq B_P(\delta, E_P(\delta))$, we know that v is enforceable with continuation values $w(y)$ in $E_P(\delta)$. But, the score \bar{k} is maximal, so we know that $\lambda \cdot v \geq \lambda \cdot w(y)$. Hence the value of the programming problem is at least $\bar{k} > k_P^*(\delta)$, which is a contradiction.

We now move to the second claim. Suppose Q_P has full dimension and pick a smooth convex set W in the interior of Q_P . It is sufficient to prove that for any such set W and any $v \in W$ there is an open set $U \ni v$ and a δ such that $U \subseteq B_P(\delta, W)$. This is easy if point v lies in the interior of W (by using a static Nash equilibrium), so focus on v on the boundary of W . Let λ be the unique vector orthogonal to the tangent of W at v , let $H(\lambda, k)$ be the half-space $\{v' : \lambda \cdot v' \leq k\}$, and let α be an action profile which enforces the score $k^*(\lambda)$ in direction λ . Then, for some δ' and $\epsilon > 0$, (α, v) can be enforced with respect to $H(\lambda, k - \epsilon)$; that is, in particular,

$$v_i = (1 - \delta')g_i(\alpha) + \delta' \mathbb{E}_{\pi(\alpha)[y]} \tilde{w}_i(y),$$

with $w(y) \in H(\lambda, k - \epsilon)$. This holds since the map from w to \tilde{w} is translation-preserving.⁷ Thus, for all $\delta'' \geq \delta'$, (α, v) can be enforced with respect to the half-plane $H(\lambda, k - \delta'(1 - \delta'')/\delta''(1 - \delta')\epsilon)$ so that all continuation values are in an $\kappa(1 - \delta'')$ -ball around v (for some constant $\kappa > 0$), because the map from w to \tilde{w} is scaling-invariant. As W is smooth, this means that there is some large $\delta < 1$ so that (α, v) is enforceable with respect to continuation values in the interior of W . Moving around the continuation values in a small neighborhood yields the desired open neighborhood U containing v . \square

C.3.3 Proof of Theorem 3.4

To prove this theorem, it helps to define the following notion for expositional purposes.

Definition C.1 (Linearly P-enforceable). *The pair (α, v) is linearly P-enforceable with respect to δ and a set W if it is P-enforceable and the w from the definition of P-enforceability can be chosen to lie on a positive sloped line.*

Define the operator $B_{PL}(\delta, \cdot)$ analogously to $B_P(\delta, \cdot)$ in Definition 3.7, replacing P-enforceability with linear P-enforceability. We say a set W is *linearly P-self generating* (linearly P-SG) if $W \subseteq B_{PL}(\delta, W)$. Note that if a set is linearly P-SG, then it is also P-SG and linearly self-generating.

Lemma C.1. *$E_L(\delta)$ is a linearly P-SG set.*

Proof. Choose $v \in E_L(\delta)$ and let σ be a totally linear PPE with payoff v . Let α be the action in the first period. Let $w(y)$ be the continuation payoffs after signal y realizes. Then note that $w(y) \in E_L(\delta)$ since any subgame of a totally linear PPE is also a totally linear PPE. Now choose

⁷By “translation-preserving” we mean that if (w, \tilde{w}, D) is a consistent triple, then $(w + e, \tilde{w} + e, D)$ is also a consistent triple, where $w + e$ denotes adding the ordered pair (e_1, e_2) to each $w(y)$. The only thing one needs to do here is shift the set of optimal w 's from the score problem $k^*(\lambda)$ a little bit towards W until v is enforceable. The \tilde{w} 's merely “follow” the w 's.

$\tilde{w}(y) = w(y)$ and back out the correspondence D from the consistency requirement so that (w, \tilde{w}, D) are consistent. Finally, since no player has an incentive to deviate in the first period of σ we know that

$$\alpha_i \in \arg \max_{a_i} (1 - \delta)g_i(a_i, \alpha_{-i}) + \delta E_{\pi(a_i, \alpha_{-i})[y]}[w_i(y)].$$

Now, since $w \equiv \tilde{w}$, this exactly says that (α, v) is P-enforceable by (w, \tilde{w}, D) . Since $w(y) \in E_L(\delta)$ and w is linear, this exactly shows that $E_L(\delta)$ is linearly P-SG. \square

That $E_L(\delta) \subseteq E_P(\delta)$ follows from Lemma C.1 and Lemma 3.3 since $E_L(\delta)$ is clearly bounded. To show $Q_L \subseteq Q_P$, we simply compare Programs (3.6) and (3.8). Fix α . Let $(v, \{w(y)\})$ be feasible in (3.8). Then, setting $\tilde{w}(y) = w(y)$ and determining D from (3.3) gives us a feasible set $(v, \{w(y)\}, D)$ in (3.6) since $w(y)$ already lie on a line of positive slope. Thus, $k_P(\alpha, \lambda, \delta) \geq k_L(\alpha, \lambda)$ and thus $k_P^*(\lambda) \geq k_L^*(\lambda)$ for all λ , meaning $Q_L \subseteq Q_P$. Together, this proves Theorem 3.4.

C.3.4 Proof of Theorem 3.5

The following lemma is needed for the proof of Theorem 3.5.

Lemma C.2. *If (w, \tilde{w}, D) is consistent, $w(y) \in W$ for all $y \in Y$, and W is convex and meet-closed (i.e., the meet of any set of points in W is also in W), then $\tilde{w}(y) \in W$ for all $y \in Y$.*

Proof. We already know that $\{\tilde{w}(y)\}$ is linear with a positive slope. Hence the points are Pareto-ranked. If $\{\tilde{w}(y)\}$ is a singleton then $D_i(y) = \Delta(Y)$ for all $y \in Y$, and thus $\tilde{w}(y) = \bigwedge \{w(y)\}_{y \in Y} \in W$. Now suppose $\{\tilde{w}(y)\}$ has at least two distinct points. Let $\tilde{w}(y_0)$ be the Pareto-worst point. Then we know that $D_i(y_0) = \Delta(Y)$, so that we still have

$$\tilde{w}(y_0) = \bigwedge \{w(y)\}_{y \in Y} \in W,$$

since W was assumed to be meet-closed. Now let $\tilde{w}(y_1)$ be the Pareto-best point, and let $Y' \subset Y$ be the set of signals such $\tilde{w}(y) = \tilde{w}(y_1)$ for only $y \in Y'$. Then we know that $D_i(y) = \Delta(Y')$ for all $y \in Y'$. But then using the definition of \tilde{w} we exactly have

$$\tilde{w}(y') = \bigwedge \{w(y)\}_{y \in Y'} \in W$$

when $y' \in Y$. Thus, $\tilde{w}(y_1) \in W$ as well. All other points of \tilde{w} lie in a line between $\tilde{w}(y_0)$ and $\tilde{w}(y_1)$, both of which lie in W . Since W is convex, we have the result. \square

Proof of Theorem 3.5. Suppose either $\lambda_1 \geq 0$ or $\lambda_2 \geq 0$. Fix α and consider a triple $(v, \{w(y)\}, D)$ that is feasible in Program (3.6). Since the half plane $H(\lambda, \lambda \cdot v) \equiv \{w' \in \mathbb{R}^2 : \lambda \cdot w' \leq \lambda \cdot v\}$ is meet-closed, Lemma C.2 says that $\tilde{w}(y) \in H(\lambda, \lambda \cdot v)$ for all y as well. It follows that $(v, \{\tilde{w}(y)\})$ is feasible in Program (3.8), meaning $k_L(\alpha, \lambda) \geq k_P(\alpha, \lambda)$, or $k_L^*(\lambda) \geq k_P^*(\lambda)$. For $\lambda_1 < 0$ and $\lambda_2 < 0$, we know that $k_L^*(\lambda) \geq v_{NE}$, since v_{NE} is supported by a Nash equilibrium.

To prove the theorem, note that

$$\begin{aligned} \mathcal{L}(v_{NE}) &\subseteq \bigcap_{\{\lambda: \lambda_1 < 0 \text{ and } \lambda_2 < 0\}} H(\lambda, \lambda \cdot v_{NE}) \cap \bigcap_{\{\lambda: \lambda_1 \geq 0 \text{ or } \lambda_2 \geq 0\}} H(\lambda, k_P^*(\lambda)) \\ &\subseteq \bigcap_{\lambda} \{w' \in \mathbb{R}^2 : \lambda \cdot w' \leq k_P^*(\lambda)\} = Q_L. \end{aligned}$$

□

C.3.5 Proof of Theorem 3.6

By Corollary 3.6.3, $Q_P = Q_L$. Now suppose there existed an action profile a with $g(a) \in Q_L$ and $g(a)$ on the Pareto frontier of V^* . Then, a has full score in some direction $\lambda \gg 0$, that is, $k_L(a, \lambda) = \lambda \cdot g(a)$. By the definition of k_L in (3.8), this requires that all continuation values $w(y)$ lie on the hyperplane defined by λ , so $\lambda \cdot w(y) = \lambda \cdot g(a)$ for all y . This, however, is impossible for values $w(y)$ that lie on a positively sloped line, unless the values are equal, i.e. $w(y) = w(y')$ for all y and y' .

If continuation values were equal, then they could not provide any incentives to enforce a , so a would have to be a stage game Nash equilibrium. This was ruled out in the statement of Theorem 3.6 and completes the proof.

Bibliography

- ABREU, D., D. PEARCE, AND E. STACCHETTI (1986): "Optimal Cartel Equilibrium with Imperfect Monitoring," *Journal of Economic Theory*, 39(1), 251–269.
- (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), 1041–1063.
- ACEMOGLU, D., AND J. LINN (2004): "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *Quarterly Journal of Economics*, 119(3), 1049–1090.
- ACKERBERG, D. (2009): "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*, 7, 343–376.
- AGHION, P., N. BLOOM, R. BLUNDELL, R. GRIFFITH, AND P. HOWITT (2005): "Competition and Innovation: an Inverted-U Relationship," *Quarterly Journal of Economics*, 120(2), 701–728.
- AGHION, P., C. HARRIS, AND J. VICKERS (1997): "Competition and Growth with Step-by-Step Innovation," *European Economic Review*, 41(3), 771–782.
- ANTON, J. J., AND D. A. YAO (1989): "Split Awards, Procurement, and Innovation," *RAND Journal of Economics*, 20(4), 538–552.
- (1992): "Coordination in Split Award Auctions," *Quarterly Journal of Economics*, 107(2), 681–707.
- ARROW, K. (1962): "Economics Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity*, ed. by R. Nelson, pp. 609–625. Princeton University Press, Princeton.
- ATHEY, S., D. COEY, AND J. LEVIN (2013): "Set-Asides and Subsidies in Auctions," *American Economic Journal: Microeconomics*, 5(1), 1–27.
- ATHEY, S., AND P. HAILE (2002): "Identification of Standard Auction Models," *Econometrica*, 70, 2107–2140.
- ATHEY, S., J. LEVIN, AND E. SEIRA (2011): "Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions," *Quarterly Journal of Economics*, 126(1), 207–257.
- BAGWELL, K., AND R. W. STAIGER (1990): "A Theory of Managed Trade," *American Economic Review*, 80(4), 779–795.
- BAJARI, P., H. HONG, AND S. RYAN (2010): "Identification and Estimation of a Discrete Game of Complete Information," *Econometrica*, 78(5), 1529–1568.

- BAJARI, P., D. NEKIPELOV, S. P. RYAN, AND M. YANG (2015a): “Demand Estimation with Machine Learning and Model Combination,” Working paper, National Bureau of Economic Research.
- (2015b): “Machine Learning Methods for Demand Estimation,” *American Economic Review*, 105(5), 481–485.
- BERRY, S., AND E. TAMER (2006): “Identification in Models of Oligopoly Entry,” in *Advances in Economics and Econometrics*, ed. by T. P. Richard Blundell, Whitney Newey, vol. 2, pp. 46–85. Cambridge University Press, Oxford.
- BHATTACHARYA, S., J. GLAZER, AND D. E. M. SAPPINGTON (1990): “Sharing Productive Knowledge in Internally Financed R&D Contests,” *Journal of Industrial Organization*, 39(2), 187–208.
- BIAIS, B., T. MARIOTTI, J.-C. ROCHET, AND S. VILLENEUVE (2010): “Large Risks, Limited Liability, and Dynamic Moral Hazard,” *Econometrica*, 78(1), 73–118.
- BLEI, D. M., A. Y. NG, AND M. I. JORDAN (2003): “Latent Dirichlet Allocation,” *Journal of Machine Learning Research*, 3(4–5), 993–1022.
- BLUNDELL, R., R. GRIFFITH, AND J. VAN REENEN (1999): “Market Share, Market Value, and Innovation in a Panel of British Manufacturing Firms,” *Review of Economic Studies*, 66(3), 529–554.
- BODOH-CREED, A. L. (2012): “Ambiguous Beliefs and Mechanism Design,” *Games and Economic Behavior*, 75(2), 518–537.
- BOLTON, P., AND M. D. WHINSTON (1993): “Incomplete Contracts, Vertical Integration, and Supply Assurance,” *Review of Economic Studies*, 60(1), 121–148.
- BOSE, S., E. OZDENOREN, AND A. PAPE (2006): “Optimal Auctions with Ambiguity,” *Theoretical Economics*, 1(4), 411–438.
- BOSE, S., AND L. RENO (2014): “Mechanism Design with Ambiguous Communication Devices,” *Econometrica*, 82(5), 1853–1872.
- BOUDREAU, K. J., N. LACETERA, AND K. R. LAKHANI (2011): “Incentives and Problem Uncertainty in Innovation Contests: An Empirical Analysis,” *Management Science*, 57(5), 843–863.
- BOUDREAU, K. J., K. R. LAKHANI, AND M. E. MENIETTI (2016): “Performance Responses to Competition Across Skill-Levels in Rank Order Tournaments: Field Evidence and Implications for Tournament Design,” *RAND Journal of Economics*, 47(1), 140–165.
- BULOW, J., AND P. KLEMPERER (1996): “Auctions Versus Negotiations,” *The American Economic Review*, 86(1), 180–194.
- CABRAL, L., G. COZZI, V. DENICOLÓ, G. SPAGNOLO, AND M. ZANZA (2006): “Procuring Innovations,” in *Handbook of Procurement*, ed. by N. Dimitri, G. Piga, and G. Spagnolo, chap. 19, pp. 483–528. Cambridge University Press, Cambridge.
- CARROLL, G. (2015): “Robustness and Linear Contracts,” *American Economic Review*, 105(2), 536–563.
- CHASSANG, S. (2013): “Calibrated Incentive Contracts,” *Econometrica*, 81(5), 1935–1971.

- CHE, Y.-K., AND I. GALE (2003): "Optimal Design of Research Contests," *American Economic Review*, 93(3), 646–671.
- CHE, Y.-K., E. IOSSA, AND P. REY (2016): "Prizes versus Contracts as Incentives for Innovation," Working paper.
- CHOI, J. P. (1991): "Dynamic R&D Competition Under "Hazard Rate" Uncertainty," *RAND Journal of Economics*, 22(8), 596–610.
- CONGRESSIONAL BUDGET OFFICE (2011): "Economic and Budget Issues Brief: Spending and Funding for Highways," Discussion paper, Washington DC.
- CRAWFORD, G. S., AND A. YURUKOGLU (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102(2), 643–685.
- CREMER, J., Y. SPIEGEL, AND C. ZHENG (2009): "Symposium on Auction Theory: Auctions with Costly Information Acquisition," *Economic Theory*, 38(1), 41–72.
- D'ASPROMONT, C., S. BHATTACHARYA, AND L.-A. GÉRARD-VARET (1998): "Knowledge as a Public Good: Efficient Sharing and Incentives for Development Effort," *Journal of Mathematical Economics*, 30(4), 389–404.
- (2000): "Bargaining and Sharing Innovative Knowledge," *Review of Economic Studies*, 67(2), 255–271.
- DE SILVA, D., T. DUNNE, A. KANKANAMGE, AND G. KOSMOPOULOU (2008): "The Impact of Public Information on Bidding in Highway Procurement Auctions," *European Economic Review*, 52, 150–181.
- DOW, J., AND S. R. DA COSTA WERLANG (1994): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64(2), 302–324.
- EDMANS, A., AND X. GABAIX (2011): "Tractability in Incentive Contracting," *The Review of Financial Studies*, 24(9), 2865–2894.
- EVDOKIMOV, K., AND H. WHITE (2012): "Some Extensions of a Lemma of Kotlarski," *Econometric Theory*, 28(04), 925–932.
- FOLEY, J. A. (2003): "Multi-Round Auctions for Institutional Real Estate Assets: Theory and Practice," Master's thesis, MIT.
- FUDENBERG, D., D. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5), 997–1039.
- FUDENBERG, D., AND D. K. LEVINE (1994): "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62(1), 103–135.
- FUDENBERG, D., D. K. LEVINE, AND S. TAKAHASHI (2007): "Public Perfect Equilibrium When Players Are Patient," *Games and Economic Behavior*, 61(1), 27–49.
- FUDENBERG, D., AND Y. YAMAMOTO (2010): "Repeated Games Where the Payoffs and Monitoring Structure Are Unknown," *Econometrica*, 78(5), 1673–1710.

- (2011): “Learning From Private Information in Noisy Repeated Games,” *Journal of Economic Theory*, 146(5), 1733–1769.
- FULLERTON, R. L., B. G. LINSTER, M. MCKEE, AND S. SLATE (2002): “Using Auctions to Reward Tournament Winners: Theory and Experimental Investigations,” *RAND Journal of Economics*, 33(1), 62–84.
- FULLERTON, R. L., AND R. P. MCAFEE (1999): “Auctioning Entry into Tournaments,” *Journal of Political Economy*, 107(3), 573–605.
- GABLONSKY, J. M., AND C. T. KELLEY (2001): “A Locally-Biased Form of the DIRECT Algorithm,” *Journal of Global Optimization*, 21(1), 27–37.
- GENTRY, M., AND T. LI (2012): “Auctions with Selective Entry,” *Working Paper*.
- (2014): “Identification in Auctions with Selective Entry,” *Econometrica*, 82(1), 315–344.
- GENTZKOW, M., AND J. SHAPIRO (2013): “Measuring the Sensitivity of Parameter Estimates to Sample Statistics,” *Working Paper*.
- GENTZKOW, M., AND J. M. SHAPIRO (2010): “What Drives Media Slant? Evidence from U.S. Daily Newspapers,” *Econometrica*, 78(1), 35–71.
- GILBERT, R. J., AND D. M. G. NEWBERY (1982): “Preemptive Patenting and the Persistence of Monopoly,” *American Economic Review*, 72(3), 514–526.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18(2), 141–153.
- GREEN, B., AND C. R. TAYLOR (2016): “Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects,” *American Economic Review*, Forthcoming.
- GREEN, E. J., AND R. H. PORTER (1984): “Noncooperative Collusion Under Imperfect Price Information,” *Econometrica*, 52(1), 87–100.
- GRENNAN, M. (2013): “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review*, 103(1), 145–177.
- GROSS, D. P. (2016a): “Creativity Under Fire: The Effects of Competition on Creative Production,” Working paper, Harvard Business School.
- (2016b): “Performance Feedback in Competitive Product Development,” Working paper, Harvard Business School.
- HALPIN, D. (2005): *Construction Management, 3rd Edition*. Wiley.
- HANSEN, S., M. MCMAHON, AND A. PRAT (2014): “Transparency and Deliberation within the FOMC: A Computational Linguistics Approach,” Working paper 9994, Centre for Economic Policy Research.
- HARRIS, C., AND J. VICKERS (1987): “Racing with Uncertainty,” *Review of Economic Studies*, 54(1), 1–21.
- HATFIELD, J. W., F. KOJIMA, AND S. D. KOMINERS (2016): “Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence,” Working paper.

- HOLMES, T. J., D. K. LEVINE, AND J. A. SCHMITZ (2012): "Monopoly and the Incentive to Innovate when Adoption Involves Switchover Disruptions," *American Economic Journal: Microeconomics*, 4(3), 1–33.
- HOLMSTRÖM, B., AND P. MILGROM (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55(2), 303–328.
- HOWELL, S. (2016): "Financing Innovation: Evidence from R&D Grants," Working paper, New York University.
- HUBBARD, T., AND H. PAARSCH (2009): "Investigating Bid Preferences at Low-Price, Sealed-Bid Auctions with Endogenous Participation," *International Journal of Industrial Organization*, 27, 1–14.
- JOFRE-BONET, M., AND M. PESENDORFER (2003): "Estimation of a Dynamic Auction Game," *Econometrica*, 71(5), 1443–1489.
- JOHNSON, S. G. (2010): "The NLOpt Nonlinear-Optimization Package," Software package, <http://ab-initio.mit.edu/nlopt>.
- KAGEL, J., S. PEVNITSKAYA, AND L. YE (2008): "Indicative Bidding: an Experimental Analysis," *Games and Economic Behavior*, 62(2), 697–721.
- KIREYEV, P. (2016): "Markets for Ideas: Prize Structure, Entry Limits, and the Design of Ideation Contests," Working paper 16-129, Harvard Business School.
- KLEMPERER, P. (2002): "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 16(1), 169–189.
- KLIBANOFF, P. (1996): "Uncertainty, Decision, and Normal Form Games," Working paper, Northwestern University.
- KOTLARSKI, I. (1967): "On Characterizing the Gamma and the Normal Distribution," *Pacific Journal of Mathematics*, 20(1), 69–76.
- KRASNOKUTSKAYA, E. (2011): "Identification and Estimation of Auction Models with Unobserved Heterogeneity," *Review of Economic Studies*, 78(1), 293–327.
- KRASNOKUTSKAYA, E., AND K. SEIM (2011): "Bid Preference Programs and Participation in Highway Procurement Auctions," *American Economic Review*, 101(6), 2653–2686.
- KREPS, D. M., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50(4), 863–894.
- LENER, J. (2000): "The Government as Venture Capitalist: The Long-Run Impact of the SBIR Program," *Journal of Private Equity*, 3(2), 55–78.
- LEVIN, D., AND J. SMITH (1994): "Equilibrium in Auctions with Entry," *American Economic Review*, 84, 585–599.
- LEVIN, J. (2003): "Relational Incentive Contracts," *American Economic Review*, 93(3), 835–857.
- LI, T. (2010): "Indirect Inference in Structural Econometric Models," *Journal of Econometrics*, 157(1), 120–128.

- LI, T., AND Q. VUONG (1998): “Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators,” *Journal of Multivariate Analysis*, 65(2), 139–165.
- LI, T., AND B. ZHANG (2010): “Affiliation and Entry in First-Price Auctions with Heterogeneous Bidders,” *Working Paper*.
- LI, T., AND X. ZHENG (2009): “Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions,” *Review of Economic Studies*, 76, 1397–1429.
- (2011): “Information Acquisition and/or Bid Preparation: A Structural Analysis of Entry and Bidding in Timber Sale Auctions,” *Journal of Econometrics*, 168(1), 29–46.
- LICHTENBERG, F. R. (1995): “Economics of Defense R&D,” in *Handbook of Defense Economics*, ed. by T. Sandler, and K. Harley, vol. 1, chap. 15, pp. 431–457. Elsevier.
- LO, K. C. (1996): “Equilibrium in Beliefs under Uncertainty,” *Journal of Economic Theory*, 71(2), 443–484.
- (1999): “Extensive Form Games with Uncertainty Averse Players,” *Games and Economic Behavior*, 28(2), 256–270.
- LOPOMO, G., L. RIGOTTI, AND C. SHANNON (2010): “Uncertainty in Mechanism Design,” Working paper, Duke University.
- (2011): “Knightian Uncertainty and Moral Hazard,” *Journal of Economic Theory*, 146(3), 1148–1172.
- LU, J., AND L. YE (2013): “Efficient and Optimal Mechanisms with Private Information Acquisition Costs,” *Journal of Economic Theory*, 148(1), 393–408.
- LYON, T. P. (2006): “Does Dual Sourcing Lower Procurement Costs?,” *Journal of Industrial Economics*, 54(2), 223–252.
- MANKIW, N., AND M. WHINSTON (1986a): “Free Entry and Social Inefficiency,” *The RAND Journal of Economics*, pp. 48–58.
- MANKIW, N. G., AND M. D. WHINSTON (1986b): “Free Entry and Social Inefficiency,” *RAND Journal of Economics*, 17(1), 48–58.
- MARMER, V., A. SHNEYEROV, AND P. XU (2011): “What Model for Entry in First-Price Auctions? A Nonparametric Approach,” *Working Paper*.
- (2013): “What Model for Entry in First-Price Auctions? A Nonparametric Approach,” *Journal of Econometrics*, 176(1), 46–58.
- MARSHALL, R. C., AND L. M. MARX (2012): *The Economics of Collusion: Cartels and Bidding Rings*. MIT Press, Cambridge, MA.
- MAURER, S. M., AND S. SCOTCHMER (2004): “Procuring Knowledge,” in *Intellectual Property and Entrepreneurship: Advances in the Study of Entrepreneurship, Innovation, and Growth*, ed. by G. Libecap, vol. 15, pp. 1–31.
- MCCALLUM, A. K. (2002): “MALLETT: A Machine Learning for Language Toolkit,” Software package, <http://mallet.cs.umass.edu>.

- MILGROM, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press.
- OECD (2011): “Size of Public Procurement Market,” in *Government at a Glance 2011*.
- OWEN, A. B. (2001): *Empirical Likelihood*. Chapman and Hall.
- PARK, W., AND W. B. CHAPIN (1992): *Construction Bidding: Strategic Pricing for Profit (Second Edition)*. Wiley.
- POWELL, M. J. D. (2009): “The BOBYQA Algorithm for Bound Constrained Optimization with Derivatives,” Technical report, Department of Applied Mathematics and Theoretical Physics, Cambridge.
- QUINT, D., AND K. HENDRICKS (2013): “Indicative Bidding in Auctions with Costly Entry,” *Working Paper, University of Wisconsin*.
- RADNER, R. (1986): “Repeated Partnership Games with Imperfect Monitoring and No Discounting,” *Review of Economic Studies*, 53(1), 43–57.
- ROBERTS, J., AND A. SWEETING (2013a): “Bailouts and the Preservation of Competition,” *NBER Working Paper 16650*.
- (2013b): “When Should Sellers Use Auctions?,” *American Economic Review*, 103(5), 1830–1861.
- ROBINSON, P. M. (1988): “Root-N-consistent Semiparametric Regression,” *Econometrica*, 56(4), 931–954.
- ROGERSON, W. P. (1989): “Profit Regulation of Defense Contractors and Prizes for Innovation,” *Journal of Political Economy*, 96(6), 1284–1305.
- (1994): “Economic Incentives and the Defense Procurement Process,” *Journal of Economic Perspectives*, 8(4), 65–90.
- (1995): “Incentive Models of the Defense Procurement Process,” in *Handbook of Defense Economics*, ed. by T. Sandler, and K. Harley, vol. 1, chap. 12, pp. 309–346. Elsevier.
- RUDIN, W. (1962): *Fourier Analysis on Groups*. Interscience Publishers.
- SAMUELSON, W. (1985): “Competitive Bidding with Entry Costs,” *Economic Letters*, 17, 53–57.
- SANNIKOV, Y., AND A. SKRZYPACZ (2010): “The Role of Information in Repeated Games With Frequent Actions,” *Econometrica*, 78(3), 847–882.
- SCHUMPETER, J. (1939): *Business Cycles*. Allen Unwin, London.
- SCHWARTZ, M., W. GINSBERG, AND J. F. SARGENT (2015): “Defense Acquisitions: How and Where DOD Spends its Contracting Dollars,” CRS Report, Congressional Research Service.
- SHAKED, A., AND J. SUTTON (1984): “Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 52(6), 1351–1364.
- SU, C., AND K. JUDD (2012): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*, 80(5), 2213–2230.

- TAYLOR, C. R. (1995): "Digging for Golden Carrots: An Analysis of Research Tournaments," *American Economic Review*, 85(4), 872–890.
- TOXVAERD, F. (2006): "Time of the Essence," *Journal of Economic Theory*, 129(1), 252–272.
- WALLSTEN, S. J. (2000): "The Effects of Government-Industry R&D Programs on Private R&D: The Case of the Small Business Innovation Research Program," *RAND Journal of Economics*, 31(1), 82–100.
- WELCH, D., AND O. FREMOND (1998): *The Case-by-Case Approach to Privatization: Techniques and Examples*. World Bank.
- WILLIAMS, H. (2012): "Innovation Inducement Prizes: Connecting Research to Policy," *Journal of Policy Analysis and Management*, 31(3), 752–776.
- YE, L. (2007): "Indicative Bidding and a Theory of Two-Stage Auctions," *Games and Economic Behavior*, 58, 181–207.