

An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program

VIVEK BHATTACHARYA

Department of Economics, Massachusetts Institute of Technology

December 7, 2016

ABSTRACT. This paper 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.

KEYWORDS. R&D procurement, contests, Small Business Innovation Research program, holdup problem, intellectual property, nonparametric identification.

JEL CODES. C51, C78, L24, O31, O32, O38.

Job Market Paper. See <http://economics.mit.edu/grad/bhattacv> for the most recent version.

CONTACT. bhattacv@mit.edu. I would like to thank Nikhil Agarwal, Glenn Ellison, and Mike Whinston for invaluable advice and support throughout this process. This paper 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 George P. and Obie B. Shultz Fund. All errors are mine.

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 paper, 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 paper, 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 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

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.

“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 paper.

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 paper—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. Related Literature

The conceptual framework for this paper 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 paper 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 relevant to this paper 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

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 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.

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 paper 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 paper 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 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 paper speak to all three of these methods

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.

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.

for controlling incentives.

2. Empirical Setting and Data

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 paper, 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 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

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

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.

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 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

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.”

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

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.

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 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

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

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 paper, 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

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: 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.

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 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 F.

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 F.

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 3.

Table 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

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).

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 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 3 indicate that this further restriction does not drop much of the data at all and the sample used in structural

	<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 2: Summary statistics for the dataset of all solicitations posted between 2000 and 2012, dropping ones in which multiple Phase III contracts were awarded.

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 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 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

²¹estimation is representative of the entire population of contests in this time period.

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

	0	1	2	3	4	≥ 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 3: Distribution of the number of competitors in each phase. As in Table 2, I restrict to solicitations posted between 2000 and 2012 and only consider ones in which at most one Phase III contract was awarded.

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 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 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

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.

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.

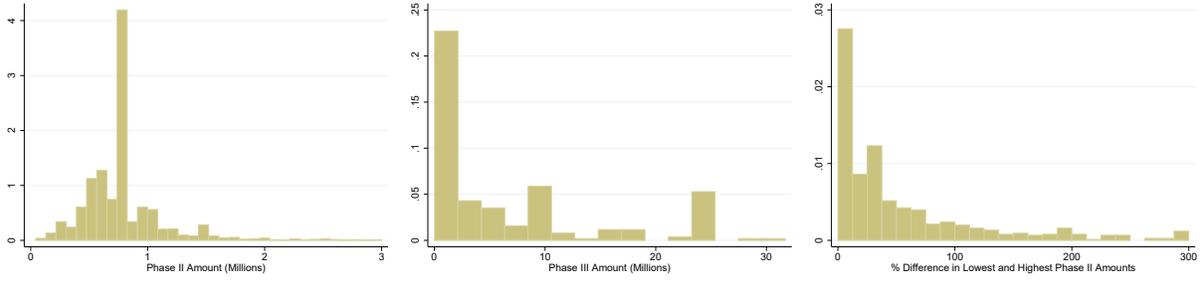


Figure 1: Distribution of (a) Phase II award amounts and (b) Phase III award amounts. The histogram in (a) includes a datapoint for each contract and can thus include multiple contracts for a particular contest. Panel (c) shows the percent difference between the highest and lowest Phase II award amounts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

	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 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 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 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% 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 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

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.

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 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 B.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.

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.

What affects Phase II contract amounts? Column (5) of Table 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 B.1 notes a large drop when moving from contests with 3 to contests with 4 Phase II competitors. The institutional details provided in Section 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 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 B.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 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

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 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.

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 B.1.

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 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.

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 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 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 paper, 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 4.1 and develop a natural estimation procedure in Section 4.3. In particular, the weak assumption is sufficient to identify a large subset of the primitives discussed in Section 3.1.

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 paper, 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 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

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 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 4.1 about which aspects of the structure imposed

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 3.2 (although, once again, by a small amount in this setting).

in Assumption O are used to identify which parameters. Furthermore, because Assumption M 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 paper.

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(otonicity) 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 paper. 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 paper, 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

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

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 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

$$\arg \max_t \left\{ \eta \int_{\underline{c}}^v \int^{v-c} (v - c - \max\{s, 0\}) dG(s; v, t_{N_2}^*(\cdot), p^*) dH(c; t) - t + t^{DOD}(v) \right\}, \quad (2)$$

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

$$S(s; v', t_{N_2}^*(\cdot)) = 1 - H(v' - s; t_{N_2}^*(v')) \quad (3)$$

and the cdf of the maximum surplus of a type- v firm's opponents can be computed by combining

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.

(3) and (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 (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

$$p^* = \arg \max_{p \in [0,1]} \left\{ p \cdot \left[\sum_{N_S=1}^{\bar{N}_2} \binom{N_1-1}{N_S-1} (p^*)^{N_S} (1-p^*)^{N_1-N_S} \int_0^{\bar{v}} \lambda(v, N_S, \bar{N}_2) \pi(v, N_2, p^*) dF(v) \right] - \psi(p) \right\}, \quad (4)$$

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 (2) and (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 (2) and (4).*

Assumption O(optimality) 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 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

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.

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 (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 (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 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 encourage them to exert more effort than firms with higher values, and the proof of Proposition 1 does not apply.³⁶

³⁵That firms would substitute internal funds for SBIR funding is consistent with the results of Wallsten (2000). Furthermore, other auxiliary evidence provided in this paper suggests that the DOD is reasonably generous to firms throughout this process. For instance, the estimated bargaining parameter in Section 5 indicates that firms capture three-fourths of the surplus.

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

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 3 and present an estimation procedure that is based on this identification argument.

4.1. Identification

Suppose that in the model in Section 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.

4.1.1. Identification Under Assumption M

Consider the model timing model described in Section 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 $v(\cdot)$ 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 E.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

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.

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 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 paper utilizes an optimality condition to recover information about the bargaining parameter (see Proposition 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 C.1.

4.1.2. 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

$$\eta \int_{\underline{c}}^{v(t_2)} (v(t_2) - c) \frac{dh}{dt}(c; \eta, t_2) dc = 1. \quad (5)$$

The intuition is that (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

³⁸Throughout this paper, 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.

rearranging (5) in terms of observables and quantities that have already been identified. I relegate it to Appendix E.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 (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 2 and summarizes the arguments in this section.

Proposition 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 4.2.

4.1.3. 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 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 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 C provides a number of extensions of this result. Proposition 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 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 (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

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 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 $\bar{N}_2 = 1$.

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

4.2. Empirical Model

Before discussing the empirical specification I will use to take the model in Section 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 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 3.2.⁴³

I add two components to the model in Section 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

$$\begin{aligned} v_{ij} &\equiv \tilde{v}_i \cdot \theta_j \cdot \exp(X_j \beta), \text{ where } \tilde{v}_i \sim \tilde{V}; \\ \psi_j(p) &\equiv \theta_j \cdot \exp(X_j \beta) \cdot \tilde{\psi}(p); \text{ and} \\ c_{ij} &\equiv \theta_j \cdot \exp(X_j \beta) \cdot \tilde{c}_i, \text{ where } \tilde{c}_i \text{ has cdf } \tilde{H}(\cdot; t/(\theta_j \cdot \exp(X_j \beta))). \end{aligned} \tag{6}$$

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 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 (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

begin in Phase II but always have at least two competitors). In addition, a combination of the overidentifying restrictions in (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 paper 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.

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.

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

Proposition 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 (2) and (4). \square

Proposition 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 4 and the specification in (6), we have that in equilibrium,

$$t_{N_2}^*(v_{ij}; X_j, \theta_j) = \theta_j \cdot \exp(X_j \beta) \cdot \tilde{t}_{N_2}(\tilde{v}_i) \quad (7)$$

for some effort function $\tilde{t}_{N_2}(\cdot)$. As will be discussed in Section 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

- 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 G.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

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

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 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.

identification discussion in Section 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.

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 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 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 regression, (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 (7) gives

$$\log t_{N_2}^*(v_{ij}; X_j, \theta_j) = X_j \beta + \log \theta_j + \log \tilde{t}_{N_2}(\tilde{v}_i).$$

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

$$\nu_{ij} \equiv \log \tilde{t}_{N_2}(\tilde{v}_i) + \log \theta_j \equiv \log \tilde{t}_i + \log \theta_j,$$

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.

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 paper 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

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

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

$$\hat{\Phi}_{\Theta, (N_1, N_2)}(t) = k \cdot \exp \left(\int_0^t \frac{\hat{\Psi}'_{(N_1, N_2)}(0, u)}{\hat{\Psi}_{(N_1, N_2)}(0, u)} du \right), \quad (8)$$

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 (8) over all pairs (N_1, N_2) with $N_2 \geq 2$ to compute the characteristic function

$$\hat{\Phi}_{\Theta}(t) = \frac{\sum_{(N_1, N_2): N_2 \geq 2} n_{(N_1, N_2)} \cdot N_2(N_2 - 1) \cdot \hat{\Phi}_{\Theta, (N_1, N_2)}(t)}{\sum_{(N_1, N_2): N_2 \geq 2} n_{(N_1, N_2)} \cdot N_2(N_2 - 1)}. \quad (9)$$

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

$$\hat{\Psi}_{(N_1, N_2)}(t) = \frac{1}{n_{(N_1, N_2)} \cdot N_2} \sum_{j: (N_{1j}=N_1, N_{2j}=N_2)} \sum_{i'} \exp(it \nu_{i'j}).$$

Then, since $\nu_{ij} = \theta_j + \tilde{t}_i$, and the characteristic function of θ_j is given by (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

$$\begin{aligned} f_{\log \Theta}(\theta) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(-it\theta) \hat{\Phi}_{\Theta}(t) dt \\ f_{\log \tilde{t}_{(N_1, N_2)}}(u) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(-itu) \hat{\Phi}_{\tilde{t}_{(N_1, N_2)}}(t) dt. \end{aligned} \tag{10}$$

In practice, these integrals are approximated on a compact interval $[-T, T]$ chosen in a data-driven fashion. See Appendix G.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: 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 G.

Step 4 (Estimation of the Bargaining Parameter). So far, estimation has only relied on Assumption M and Proposition 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 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 $\widetilde{\text{Pr}}(\text{failure}; \eta, \theta^*(\eta))$, where $\theta^*(\eta)$ are the MLE estimates conditional on η from Step 3, and the (partial) expectation of the observed transfer, $\left(1 - \widetilde{\text{Pr}}(\text{failure}; \eta, \theta^*(\eta))\right) \cdot \tilde{\mathbb{E}}[t_3; \eta, \theta^*(\eta)]$. I match these moments conditional on (N_1, N_2) . Thus, for a contest j , the relevant set of moments is

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

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

$$\eta^* = \arg \min_{\eta} \left(\sum_j g_j(\eta) \right)' \cdot \hat{\Omega}^{-1} \cdot \left(\sum_j g_j(\eta) \right), \quad (11)$$

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 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.

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 (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

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

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 $\hat{p}_{(N_1, \bar{N}_2)}$. I then use the FOC associated with (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) .

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 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 B.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 4, for comparison.

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.

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 5: Structural estimates

Panel (a) of Table 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 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 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 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(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 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

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.

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 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 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 paper, I will maintain this functional form, with the caveat that I may be underestimating the cost of Phase I research slightly.

Figure 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

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.

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

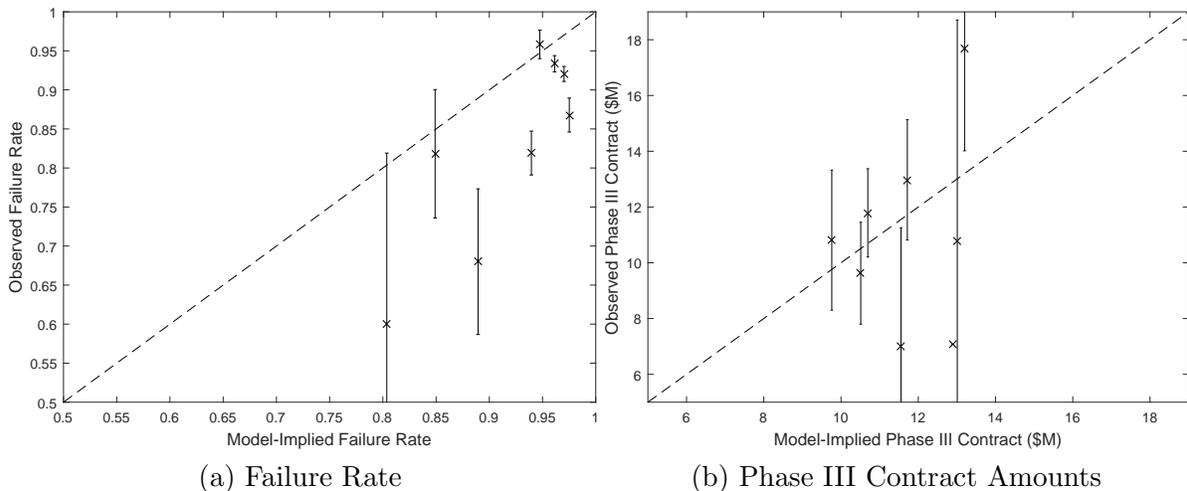


Figure 2: Observed and model-implied (a) failure rates and (b) Phase III contract amounts. The error bars indicate *one* standard error; the point in (b) without standard error bars corresponds to a bin with just one data point.

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.

6. Social Inefficiency in R&D Contests

In this section, I use the estimated values from Section 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 design counterfactuals studied in Sections 7 through 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 7 through 9 focus on social surplus as the outcome of interest; Section 10 considers the profits of the DOD as well, and Appendix A 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

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 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.

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 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, we would expect Phase II research efforts to be underprovided in the R&D contest when $N_2 = 1$. Accordingly, Table 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

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.

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

the following proposition, whose proof is in Appendix E.2.⁶¹

Proposition 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 6 also shows that values of γ for cases where $\bar{N}_2 = 2$. In these situations, we would have 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 (1986): 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 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 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.

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.

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 6 throughout the paper 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 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.

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 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.

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 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 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 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 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

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.

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 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 7.2, but the fact that research efforts increase almost linearly (see Table 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

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 paper, 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 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 (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

$$\begin{aligned}
& \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_{N_2}^*(\cdot; p_{(N_1, \bar{N}_2)}^*)\}_{N_2 \leq \bar{N}_2}\right) - S\left(1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{total effect}} \\
&= \underbrace{S\left(N_1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right) - S\left(1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{direct effect of Phase I competition}} \\
&+ \underbrace{S\left(N_1, \bar{N}_2, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right) - S\left(N_1, 1, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{direct effect of Phase II competition}} \\
&+ \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_1^*(\cdot)\}\right) - S\left(N_1, \bar{N}_2, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{incentive effect from Phase I competition}} \\
&+ \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_{N_2}^*(\cdot; p_{(N_1, \bar{N}_2)}^*)\}_{N_2 \leq \bar{N}_2}\right) - S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_1^*(\cdot)\}\right)}_{\text{incentive effect from Phase II competition}}. \quad (12)
\end{aligned}$$

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 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 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.

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 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 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 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 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.

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.

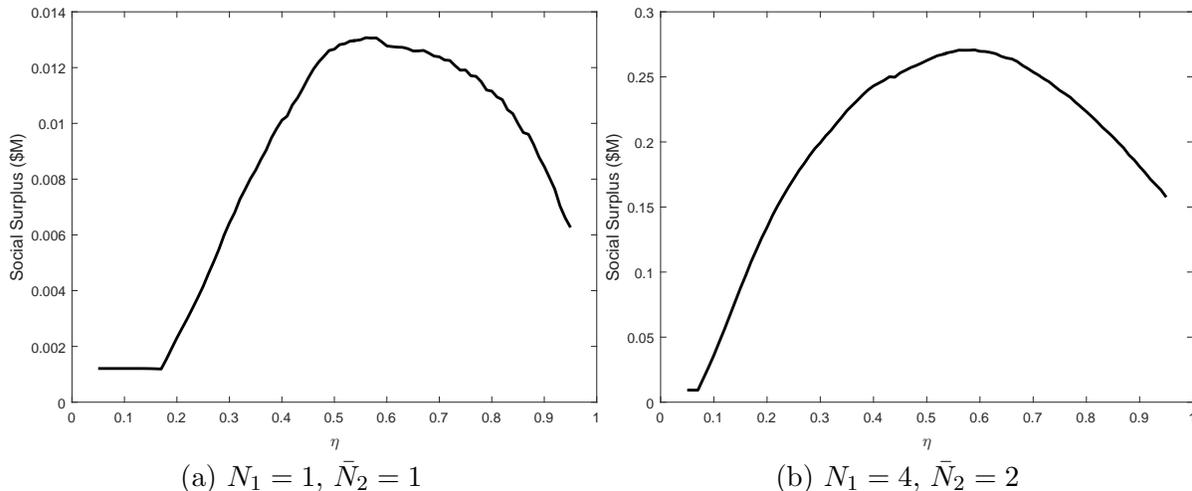


Figure 3: Social surplus as a function of η for two different levels of competition.

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 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 6 and reports the percent increase in social surplus from changing to the optimum.

Figure 3 shows social surplus as a function of η for two levels of competition, and Table 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 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 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 put this gain into perspective, we can compare the social surplus from optimizing η in Table 9 with the optimal values in Table 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

⁶⁹While it is possible to invoke the language of direct and incentive effects to discuss Figure 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, so the derivative of the social surplus function with respect to η is the incentive effect.

$(N_1, \bar{N}_2) = (2, 1)$ 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. 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.

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.

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

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.

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 3.2 can be modified immediately. In Section 9.1, I present a model of IP sharing. I comment briefly on the issue of mandatory vs. optional IP sharing as well. Section 9.2 presents the value of IP sharing in the empirical setting of the DOD SBIR program.

9.1. A Model of IP Sharing

To develop a model of IP sharing, I consider the same timeline as in Section 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, 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.

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 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.

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

$$t^*(v) = \arg \max_t \left\{ \eta \int_{\underline{c}}^v \int_{\underline{c}}^c (c' - c) dH_{\bar{N}_2-1:\bar{N}_2-1}(c'; t^*(v)) dH(c; t) - t \right\}, \quad (13)$$

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

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

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

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

Then, in equilibrium, p^* satisfies

$$p^* = \arg \max_p \{p \cdot \mathbb{E}[\pi_{\text{success}}(v; p^*, K)] + (1-p) \cdot \pi_{\text{failure}}(p^*) - \psi(p)\}. \quad (14)$$

An equilibrium of the R&D contest with IP sharing is a pair $(p^*, t^*(\cdot))$ such that $t^*(\cdot)$ satisfies (13) and p^* satisfies (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 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 D.

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 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 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 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 contests studied through the rest of the paper) 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 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 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

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

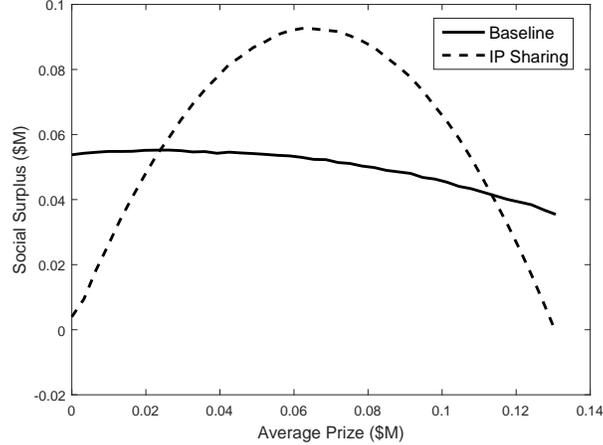


Figure 4: Social surplus as a function of the Phase I prize, without IP sharing and with mandatory IP sharing. I use the estimated parameters for $N_1 = 2$ and use $N_1 = \bar{N}_2 = 2$.

		No IP Sharing			Mandatory IP Sharing				
		$K = 0$	K_{SS}^*		$K = 0$	K_{SS}^*		K_{IC}	
N_1	\bar{N}_2	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 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.

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.

⁷⁷Note that this is not a forgone conclusion. It is possible in this model for IP sharing to be socially suboptimal

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.

10. DOD Profits Under Alternate Contest Designs

How does the DOD evaluate the alternate contest designs proposed in Sections 7 through 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 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 , 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 7 through 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

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 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.

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

⁷⁹As shown in Appendix A.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.

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.

11. Conclusion

In this paper, 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 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 paper 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.

References

- 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.
- 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.
- 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.
- 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.
- BOLTON, P., AND M. D. WHINSTON (1993): “Incomplete contracts, vertical integration, and supply assurance,” *Review of Economic Studies*, 60(1), 121–148.
- 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.
- 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.

- CRAWFORD, G. S., AND A. YURUKOGLU (2012): “The welfare effects of bundling in multichannel television markets,” *American Economic Review*, 102(2), 643–685.
- D’ASPREMONT, 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.
- EVDOKIMOV, K., AND H. WHITE (2012): “Some extensions of a lemma of Kotlarski,” *Econometric Theory*, 28(04), 925–932.
- 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.
- 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.
- GREEN, B., AND C. R. TAYLOR (2016): “Breakthroughs, deadlines, and self-reported progress: contracting for multistage projects,” *American Economic Review*, Forthcoming.
- 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.
- 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.

- HOWELL, S. (2016): “Financing innovation: evidence from R&D grants,” Working paper, New York University.
- JOHNSON, S. G. (2010): “The NLOpt nonlinear-optimization package,” Software package, <http://ab-initio.mit.edu/nlopt>.
- KIREYEV, P. (2016): “Markets for ideas: prize structure, entry limits, and the design of ideation contests,” Working paper 16-129, Harvard Business School.
- 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.
- LERNER, J. (2000): “The government as venture capitalist: the long-run impact of the SBIR program,” *Journal of Private Equity*, 3(2), 55–78.
- LI, T., AND Q. VUONG (1998): “Nonparametric estimation of the measurement error model using multiple indicators,” *Journal of Multivariate Analysis*, 65(2), 139–165.
- 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.
- LYON, T. P. (2006): “Does dual sourcing lower procurement costs?,” *Journal of Industrial Economics*, 54(2), 223–252.
- MANKIW, N. G., AND M. D. WHINSTON (1986): “Free entry and social inefficiency,” *RAND Journal of Economics*, 17(1), 48–58.
- 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>.
- OWEN, A. B. (2001): *Empirical Likelihood*. Chapman and Hall.
- POWELL, M. J. D. (2009): “The BOBYQA algorithm for bound constrained optimization with derivatives,” Technical report, Department of Applied Mathematics and Theoretical Physics, Cambridge.
- 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.
- 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.
- 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.
- WILLIAMS, H. (2012): “Innovation inducement prizes: connecting research to policy,” *Journal of Policy Analysis and Management*, 31(3), 752–776.

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

In this Appendix, I repeat the exercises of Sections 7 through 9 but report the effects on DOD profits (as well as Phase III DOD profits, as defined in Section 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.2.

A.1. Changing the Number of Competitors

I first repeat the analysis in Section 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 12 supplements Table 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.⁸⁰ About \$200,000 of this is due to R&D cost reimbursements (see Table 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 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 13 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 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.

⁸⁰For 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 12: 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.

	Social Surplus				DOD Profits			
Baseline ($N_1 = 1 = \bar{N}_2 = 1$)	0.144				-0.103			
	\bar{N}_2							
	1	2	3	4	1	2	3	4
Direct (Phase I)								
$N_1 = 2$	-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)								
$N_1 = 2$	–	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)								
$N_1 = 2$	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)								
$N_1 = 2$	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 13: 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 (12).

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.2. Changing η : The Procurer’s Laffer Curve

As discussed at the end of Section 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 to exert effort, and there will be very little surplus generated in the R&D contest.⁸¹ 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.⁸² 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 5 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 14 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 8 and included in Table 14 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

⁸¹As long as exerting no effort corresponds to failure with probability 1, there will indeed be zero surplus generated through the R&D contest.

⁸²In principle, there could be other nonmonotonicities in the intermediate region.

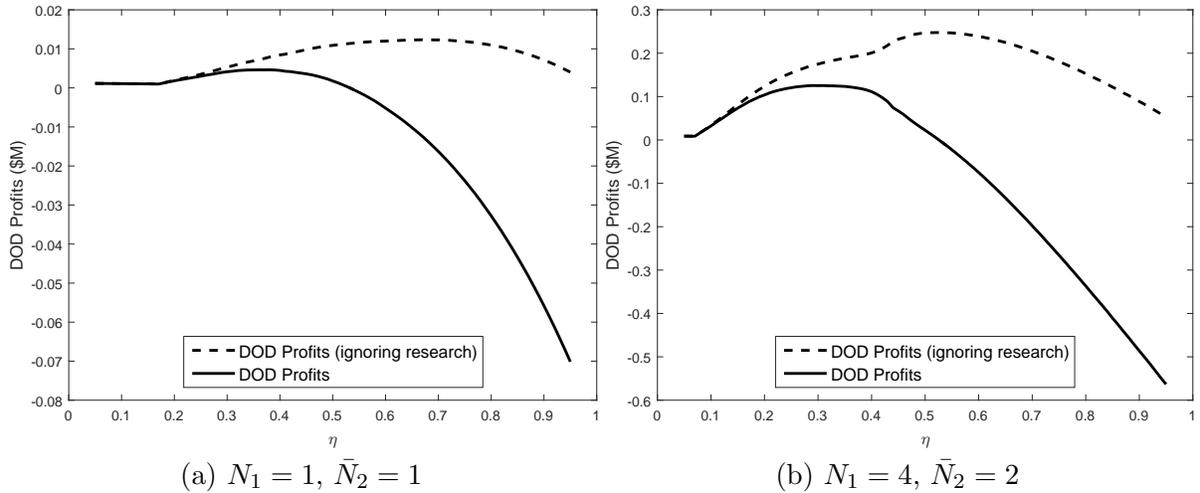


Figure 5: Profit of the DOD (both ignoring and incorporating reimbursing the cost of effort) as a function of η for two different levels of competition.

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 14: 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 $\eta = 0.73$.

firms and the DOD.

Following Section 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.⁸³

⁸³To 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 .

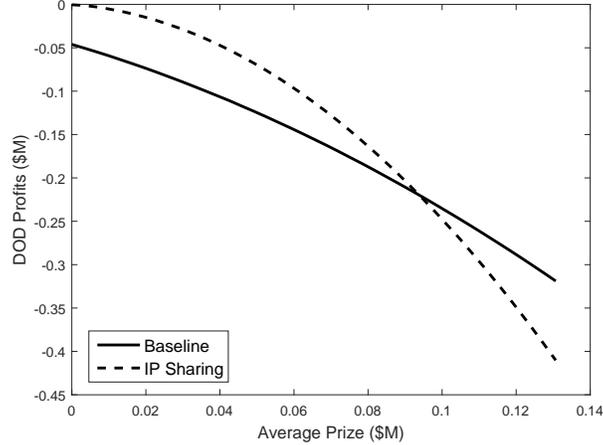


Figure 6: DOD profits as a function of the Phase I prize, with and without mandatory information sharing. I use the estimated parameters for $N_1 = 2$ and use $N_1 = \bar{N}_2 = 2$.

A.3. Mandating IP Sharing

Finally, I study whether the DOD would benefit from mandating IP sharing. Figure 6 shows a similar plot to Figure 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 15, like Table 10, conducts the analysis for other values of N_1 and \bar{N}_2 , reporting 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 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 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.

N_1	N_2	$K = 0$		K_{SS}^*		
		SS	DOD	$\mathbb{E}[K(\cdot)]$	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	$K = 0$		K_{SS}^*			K_{IC}		
		SS	DOD	$\mathbb{E}[K(\cdot)]$	SS	DOD	$\mathbb{E}[K(\cdot)]$	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 15: 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.

B. Additional Empirical Results

In this Appendix, I provide further details about the descriptive statistics presented in Section 2.3. I also present and discuss structural estimates conditional on fixed values of η , to apply the baseline model to this setting.

B.1. Details on the Descriptive Statistics

Table 16 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 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 17 investigates the probability that an *individual* competitor generates successful research,

	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 16: 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.

and it extends columns (3) and (4) of Table 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 18 extends the results shown in column (5) of Table 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 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 7 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 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 17: 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.

Phase II funding, and the estimated function is monotonically increasing both when uncontrolled and when controlling for covariates in a partially linear model.⁸⁴ 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.

Finally, column (6) of Table 18 extends the results of column (6) of Table 4 by adding dummies for the number of competitors, and the results are not substantively different. Columns (3) and (4) of Table 18 run the same regressions but include Phase III contracts with values less than \$1 million.⁸⁵ The qualitative results do not change, and the quantitative results are similar for Phase I

⁸⁴The 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.

⁸⁵Some of these contracts are actually unreasonably small and sometimes amount to \$50,000.

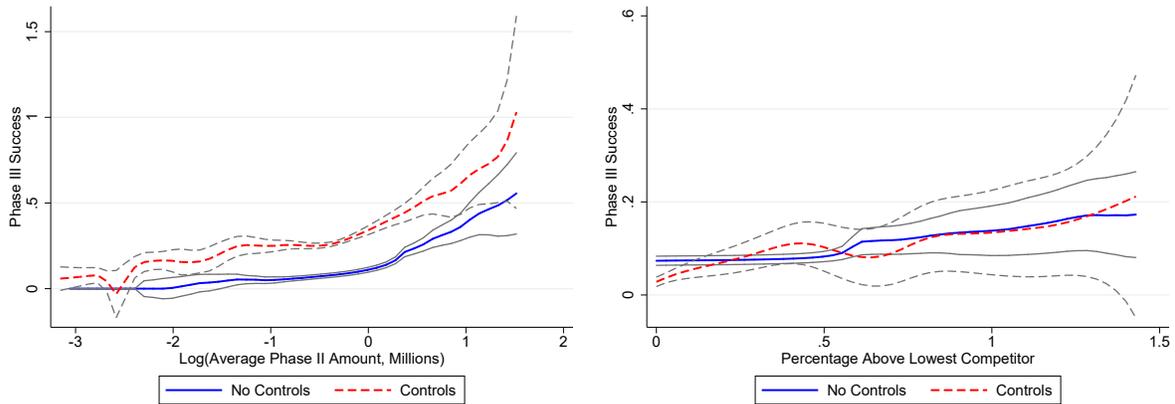


Figure 7: (a) Local linear regression of contest-level success from Phase II to Phase III on the log of average Phase II funding, both controlling and not controlling for covariates. (b) Local linear regression of individual success on the ratio between the particular individual’s Phase II funding amount and the lowest funding amount within that contest.

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 paper. 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.

B.2. Structural Estimates without Assumption O

In this section, I provide estimates using Assumption M and an analogue of the scaling property of Proposition 4. In particular, I make the following assumption, which amounts to assuming the result of Proposition 4 without imposing Assumption O. The following assumption is also natural and, like Proposition 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(calng), Steps 1–3 of the estimation procedure outlined in Section 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 19 presents these new structural results for the value and delivery cost

	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 18: 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.

distribution. Since the distribution of unobserved heterogeneity is independent of η , I do not report those results again and instead refer to Table 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 19(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. 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 5 roughly lie between the estimates for $\eta = 0.50$ and $\eta = 0.80$.

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 (0.010)	0.013 (0.005)	14.83 (1.19)	-0.017 (0.005)	6.42 (0.77)	20.90 (2.53)	39.21 (4.76)	-0.172 (0.047)
	0.062 (0.011)	0.012 (0.005)	13.45 (1.52)	-0.017 (0.005)	6.26 (0.83)	20.37 (2.75)	38.22 (5.20)	-0.182 (0.041)
0.80	0.064 (0.010)	0.012 (0.002)	4.57 (0.51)	-0.017 (0.002)	2.08 (0.29)	6.78 (0.95)	12.72 (1.79)	-0.176 (0.024)

(b) Delivery cost distributions

Table 19: Structural estimates for the baseline model, for three different values of η . The parameters presented are the same ones as in Table 5.

The identification argument presented in Section 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 depressed

when η is higher.

Table 19(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.

C. Extensions to the Identification Result

In this Appendix, I discuss a number of extensions to the identification results in Section 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 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 paper, 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.

C.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 $g(t')^{\text{th}}$ quantile of

the distribution of Phase III research efforts conditional on t'' , denoted $T_3(g(t'), t'')$.⁸⁶ 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

$$\eta \geq \frac{T_3(g(t'), t'') - v(t')}{v(t'') - v(t')}. \quad (15)$$

Since (15) has to apply for all t' and t'' , we have the lower bound

$$\eta \geq \max_{t', t'' > t'} \frac{T_3(g(t'), t'') - v(t')}{v(t'') - v(t')}. \quad (16)$$

Thus, we have the following proposition.

Proposition 6. *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 (16).*

C.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 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 4.1 directly to identify η , V_k , and $H_k(\cdot; \cdot)$;

C.3. Unobserved Heterogeneity

Suppose that, as in the empirical specification in Section 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.⁸⁷

⁸⁶Do 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.

⁸⁷While I have not shown that these technical conditions are necessary, they are needed for the proof I provide.

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.⁸⁸ 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 $g(\tilde{t}) \equiv 1 - \tilde{H}(\tilde{v}(t); \tilde{t})$ (and zero if \tilde{t} is outside 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

$$k(t) \equiv \int_0^\infty f_\Theta(\theta) \cdot g\left(\frac{t}{\theta}\right) d\theta, \quad (17)$$

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 $d\theta/\theta$.⁸⁹

Suppose that $g(\cdot)$ is such that $\int_0^\infty g(\theta) d\theta < \infty$.⁹⁰ 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 (17), 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 $\int \ell(T|t) =$

⁸⁸Evdokimov and White (2012) provide alternate conditions under which this identification result remains true.

⁸⁹Fourier 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.

⁹⁰Since $g(\cdot)$ 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.

$1 - k(t)$. This quantity is observed in the data, and we can express it as

$$\begin{aligned} \ell(T|t) &= \int_0^\infty f_\Theta(\theta) \cdot \Pr\left(\eta\tilde{v}\left(\frac{t}{\theta}\right) + (1-\eta)c = T; \frac{t}{\theta}\right) d\theta \\ &= \int_0^\infty f_\Theta(\theta) \cdot \frac{1}{1-\eta} \tilde{h}\left(\frac{T/\theta - \eta\tilde{v}(t/\theta)}{1-\eta}; \frac{t}{\theta}\right) \frac{d\theta}{\theta} \\ &\equiv \int_0^\infty f_\Theta(\theta) \cdot q\left(\frac{T}{\theta}; \frac{t}{\theta}\right) \frac{d\theta}{\theta}, \end{aligned}$$

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 .⁹¹ 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 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 7. *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

⁹¹It 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.

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.

C.4. External Benefits of Research

In the model presented in Section 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 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 2 applies immediately to this model.

However, the first order condition of the firm changes from (5) to

$$b'(t_2) + \eta \int_{\underline{c}}^{v(t_2)} (v(t_2) - c) \frac{dh}{dt}(c; \eta, t_2) dc = 1, \quad (18)$$

so applying Proposition 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 3 to identify η . Then, for all other t_2 , (18) identifies $b'(t_2)$. We codify this in the following proposition.

Proposition 8. *Consider the model in Section 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 $b(0) = 0$.

D. 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.⁹²

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 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.

⁹²Specifying 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.

E. Omitted Proofs

I collect the proofs of Propositions 2, 3, and 5 in this appendix.

E.1. Proof of Propositions 2 and 3

The one remaining step to prove Proposition 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 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

$$F_{V_S}(v) = \frac{1}{(1-p^*)^{N_1}} \sum_{N_S=1}^{N_1} \binom{N_1}{N_S} (p^*)^{N_S} (1-p^*)^{N_1-N_S} F_V(v)^{N_S}. \quad (19)$$

The right-hand side of (19) is a convex combination of increasing functions of $F_V(v)$. Since $F_{V_S}(v)$ is identified, we can invert (19) to identify the cdf of V .

The missing step in Proposition 3 is to show that (5) has a unique solution for η . Recall that the firm sets its research effort in response to the first-order condition in Section 4.1, given by

$$\eta \int_{\underline{c}}^{v(t_2)} (v(t_2) - c) \frac{dh}{dt}(c; \eta, t_2) dc = 1 \quad (5)$$

Integrating (5) by parts, we have

$$\eta \int_{\underline{c}}^{v(t_2)} \frac{dH}{dt}(c; \eta, t_2) dc = 1. \quad (20)$$

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 (20), 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

$$\frac{\eta}{1-\eta} \int_{\underline{I}}^{v(t_2)} \frac{d\hat{F}}{dt}(u; t_2) du + \frac{\eta^2}{1-\eta} \int_{\underline{I}}^{v(t_2)} v'(t_2) \cdot \hat{f}(u; t_2) du = 1, \quad (21)$$

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 (21) 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 $c \leq v(t_2)$.

E.2. Proof of Proposition 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

$$\begin{aligned} & \max_t \left\{ \mathbb{E} \left[\max\{S_i(v, t), \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i}))\}^+ \right] - t \right\} \\ &= \max_t \left\{ \mathbb{E} \left[\{S_i(v, t) - \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i}))\}^+ + \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i}))^+ \right] - t \right\} \\ &= \max_t \left\{ \mathbb{E} \left[\{S_i(v, t) - \max_{-i} S_{-i}(v_{-i}, t_{-i}^*(v_{-i}))\}^+ \right] - t \right\}, \end{aligned} \quad (22)$$

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, (22) 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 .

F. 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.

F.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.

F.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.

F.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.⁹³ 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.

⁹³Note that these variables are not fixed effects because they are proportions rather than binary variables. However, they do still sum to 1.

G. Computational Methods

G.1. Computing and Optimizing the Likelihood Function

G.1.1. 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 paper, I use $\log \pi_{\text{penalty}} = -100$. None of the data points in the main estimates of the paper are affected by this penalty term, and at most three of the data points in each of the estimates reported in this paper have a likelihood given by this penalty term. Indeed, estimates do not change for penalties in an neighborhood of this value, but if it the penalty is taken to be too low, then a large portion of points are simply rationalized by this penalty.

G.1.2. 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.

G.2. Inverting the Characteristic Function

In practice, computing (10) is numerically challenging, and I follow suggestions outlined by Krasnokutskaya (2011). Applying (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.⁹⁴

⁹⁴I have noticed through visual inspection that this procedure sometimes still leads to distributions that oscillate 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 \tilde{t} by 50% for some additional smoothing.

G.3. Solving for the Equilibrium Effort Function

When solving the equilibrium model in Section 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 (2), with $t_{N_2}^*(\cdot)$ in (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} .⁹⁵
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). 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 0.99.⁹⁶
4. In the final step, I compute p^* from the first-order condition associated with (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.

⁹⁵I 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 (2) I update it to a convex combination of the solution and the previous iteration, the algorithm is more likely to converge.

⁹⁶In practice, I have found it is sufficient to compute the equilibria for $p = \{0.01, 0.1, 0.2, \dots, 0.9, 0.99\}$.