

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

VIVEK BHATTACHARYA

Department of Economics, Northwestern University

August 8, 2018

ABSTRACT. The degree of competition and the level of rewards are central aspects of the design of contests for the procurement of innovative products, but their effects are theoretically ambiguous. This paper develops an empirical framework for analyzing R&D procurement contests based on estimates of values, research and delivery costs, and the share of surplus firms receive, using data on research expenditures and procurement contracts. Identification is based on the generally applicable conditions that firms with higher values spend more on R&D, that only positive-surplus projects are delivered, and that research effort is chosen optimally. Applying the model to the Small Business Innovation Research program in the Department of Defense suggests improvements in the design. While the high-powered incentives that firms currently receive are close to socially efficient, increasing competition and mandating that firms share intellectual property would both improve social surplus. However, socially beneficial design changes often harm the DOD.

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

JEL CODES. C51, C78, O31.

CONTACT. vivek.bhattacharya@northwestern.edu. This is a revised version of Chapter 1 of my PhD dissertation at MIT. I would like to thank Nikhil Agarwal, Glenn Ellison, and Mike Whinston for invaluable advice and support throughout this process. Panle Barwick, Gastón Illanes, Karam Kang, Rob Porter, Jimmy Roberts, and numerous seminar participants also provided valuable feedback. Chris Snyder and Elena Krasnokutskaya provided helpful discussions. Conversations with Jonathan Libgober were instrumental in the early stages of the project. I had useful discussions with John Williams and G. Nagesh Rao at the Small Business Administration and Robert Smith at the Office of Naval Research about the institutions and data sources. This material is based on work supported by the NSF Graduate Research Fellowship under Grant No. 1122374 as well by the John Krob Castle 1963 Fellowship and the George P. and Obie B. Shultz Fund. All errors are mine.

1. Introduction

Firms and governments often procure innovative products that are not commercially available and thus must be researched, designed, and developed. The Department of Defense procures unique weapon systems that require substantial research, usually inviting competitive prototypes. Automobile companies require parts suppliers to design products customized for specific models. Outside traditional R&D, local governments solicit costly proposals from multiple architecture firms when procuring services for large construction projects. An important feature of these settings is that competition happens not just in delivery but also in design and development. In particular, firms are often incentivized through mechanisms that resemble *R&D 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.¹

A central decision in the design of these procurement contest is the degree of competition to allow. Introducing an additional competitor increases the chance of a successful innovation and can directly reduce the price the procurer pays for the innovation in delivery. On the other hand, duplicating research efforts can be costly from a social perspective. Moreover, anticipating a lower expected reward, each competitor may reduce their research effort—or more generally their investment in cost-reduction or quality (Taylor, 1995; Fullerton and McAfee, 1999; Che and Gale, 2003). How competition should be modulated—and more generally the level of incentives that should be provided—is an empirical question that depends on the underlying primitives of the setting. It is perhaps not surprising that there is debate in practice over how to structure procurement contests.² This paper provides a first step towards an empirical framework for quantifying these tradeoffs: 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 DOD.

R&D contests are often multistage processes involving separate research and procurement phases. The design can control incentives and competition in various ways. I investigate three natural methods of controlling competition. First, I study the “extensive margin” of competition by investigating the optimal number of competitors that the procurer—in this case, the DOD—should admit to each phase of the contest. In doing so, I decompose the effect of competition into the direct effect of adding competitors and the indirect incentive

¹See Cabral, Cozzi, Denicoló, Spagnolo, and Zanza (2006) and Williams (2012) for reviews.

²Drezner and Huang (2009) discuss the tradeoffs in the context of the Weapon Systems Acquisition Reform Act of 2009, which effectively mandated competitive prototyping in early stages of major defense acquisition programs. McMillan (1990) notes that American car manufacturers would encourage competition between many suppliers in the 1980s, before moving towards the Japanese model of inviting fewer bidders to compete and rewarding them with longer contracts in an effort to incentivize quality.

effect of allowing these competitors to adjust their research effort. Second, I consider the “intensive margin” of competition, modulated by the portion of the 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. Finally, I consider changes in the prize structure, first to partially decouple the incentives for research from the final procurement contract and then to study the benefits of sharing intermediate research breakthroughs.

I evaluate these mechanisms in the context of the DOD Small Business Innovation Research (SBIR) program. In this program, a set of firms first 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 an innovative idea and learn its value to the DOD. In the second phase, the most promising firms 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 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 and sets the contract amount via a natural extension of Nash bargaining, where the bargaining parameter provides a reduced-form method to capture the share of the surplus the firm receives in contracting.

The aforementioned structure is representative of other settings in which R&D is followed by procurement. The empirical strategy developed here is based on general assumptions that are likely applicable more broadly. I show that the underlying parameters of the model—the distributions 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 and the delivery contract amounts. I provide a constructive identification proof to make the argument transparent, and the key conditions are relatively weak: the conditions that (i) firms with higher-value projects have more of an incentive to exert effort and (ii) the procurer would never purchase a project whose delivery cost exceeds its value are sufficient to nonparametrically identify all parameters conditional on the bargaining parameter. The condition that (iii) the research effort is set optimally then identifies the bargaining parameter. This argument can be extended to generalizations of the model.

I then use this identification argument to develop a multistep, computationally tractable estimator that allows a researcher with knowledge of the contracting procedure to remain agnostic about the actual process that determines the effort schedule (the map from values to research efforts). In this case, knowledge of the bargaining parameter lets me estimate values, delivery costs, and research costs. In the final steps of the procedure, I utilize the condition that the firm sets research efforts optimally to estimate the bargaining parameter.

When adapting this procedure to contests outside the one studied here, this condition can readily be replaced by one more appropriate to the setting without influencing the earlier steps. Moreover, the procedure controls for unobserved heterogeneity that affects both values and costs, borrowing techniques from auctions (Li and Vuong, 1998; Krasnokutskaya, 2011).

SBIR provides an important laboratory for studying R&D contests for two reasons. First, the DOD SBIR program is an interesting example of government procurement of R&D, which is an especially large market: in 2014, the US federal government accounted for almost 10% of all R&D expenditures worldwide. The DOD in particular spent \$28 billion on R&D, more than all other agencies combined (Schwartz, Ginsberg, and Sargent, 2015). DOD SBIR—and SBIR more generally—is important in its own right. 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 is a core part of the acquisition process, and the DOD solicits research on technologies related to all major defense acquisition programs. While the DOD is relatively unique in its focus on acquisition through SBIR, ten other federal agencies and many other countries run programs that follow the SBIR model.³ Second, the program is well-structured and data is readily available. At each stage, the set of competing firms is known to the econometrician, data on R&D and procurement contracts is readily available, and research can be tied directly to the product developed from it—which is a difficulty in studies of innovation.

Using this data and the empirical procedure described above, I estimate that the DOD values successful projects at an average of \$14–\$18 million and tends to invite more competitors to contests that it finds more valuable. The within-contest variation in values is fairly small; most of the final variation in contract amounts comes from variation in delivery costs drawn in the development phase. Finally, firms capture about two-thirds of the surplus generated by the program. The identification argument allows me to clearly comment on the patterns in the data that lead to these estimated parameters.

I then quantify the inefficiencies inherent in this design. Research in the later phase is underprovided due to a holdup effect; removing this holdup cost improves social efficiency by 10–14%. Research in the early stages 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, reducing social surplus by as much as 6%. These inefficiencies are informative by themselves, but they also feed into the analysis of alternate contest designs. I find that the social planner would like to admit a large number of competitors into both

³All federal agencies with an extramural R&D budget of at least \$100 million run a SBIR program. The DOD SBIR notably is the largest such program, as the total budget of all other SBIR programs amounts to about \$1.5 billion a year. The United Kingdom’s Small Business Research Initiative and the European Commission’s “SME Instrument” via Horizon 2020 have an almost identical multistage structure for most contests. China’s Innofund also provides financing to SMEs working on particular technology areas.

phases of the contest whereas the DOD prefers to restrict competition severely; the incentive effect is usually beneficial for social surplus and DOD surplus, but the DOD does not capture much of the direct effect of adding late-stage competitors. I show evidence that the current design is both on the efficient side of the Laffer curve and in fact fairly close to the social optimum, although DOD profits would increase considerably by reducing the share the firms capture. Decoupling research and delivery increases social surplus considerably but may reduce DOD profits. These counterfactuals suggest that the social planner and the DOD often have conflicting incentives for design.

1.1. Related Literature

The nontrivial interaction between competition and innovation has been of interest to economists in many fields.⁴ 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—an added chance of success and an increase in total research costs—with the indirect *incentive effect* on efforts. Taylor (1995), Fullerton and McAfee (1999), and Che and Gale (2003) present models in which the salient conclusion is the importance of restricting entry into contests to counteract this incentive effect. Of course, my setting has some differences. First, the incentives in my setting come from a procurement contract instead of a fixed prize.⁵ Second, my setting is explicitly a multistage process in which breakthroughs (or draws of values and costs) happen sequentially and innovation requires successes in both stages: progress on values influences the effort exerted on minimizing costs.⁶ This paper is the first to use the foundations of R&D contests to build a rich structural model of R&D procurement with costly effort and multistage progress.

The empirical setting is related to two strands of the R&D literature. First, a growing empirical literature studies online “ideation” contests, in which many competitors offer solutions to narrow, short-term problems.⁷ While the motivation of looking at controlled environments with R&D is similar, these papers differ from my setting in a number of ways: I focus on multistage contests; the projects in my dataset differ in multiple dimensions (values and costs); and the incentive structure I consider is different, as competitors are rewarded by procurement contracts rather than fixed prizes. Moreover, the approaches in

⁴See Schumpeter (1939), Arrow (1962), and Gilbert and Newbery (1982). These theoretical analyses have inspired a number of cross-firm studies, such as Blundell, Griffith, and van Reenen (1999) and Aghion, Bloom, Blundell, Griffith, and Howitt (2005).

⁵In this sense, Che and Gale (2003) and, more recently, Che, Iossa, and Rey (2016) and Liu and Lu (2018) consider the closest incentive scheme. Koh (2017) adds stochasticity to Che and Gale (2003), an important component here.

⁶“Leaders” and “laggards” in R&D races have differential incentives for research. See Harris and Vickers (1987) and Choi (1991), and Green and Taylor (2016) on multiple breakthroughs.

⁷See, for instance, Boudreau, Lakhani, and Menietti (2016), Gross (2016, 2017), Kireyev (2016), and Lemus and Marshall (2018).

these papers (structural or otherwise) are starkly different from the one I develop here. The second strand is that of the SBIR program itself. Lerner (2000) and Howell (2017) document long-term effects on businesses, and Wallsten (2000) suggests that SBIR financing crowds out private investment. Unlike these papers, I study competition within the SBIR program itself, and I also focus on an agency that uses it as part of procurement—as the Navy aspires to—rather than as a substitute for private R&D or venture capital funding.

This paper relates to defense procurement. An advantage of studying SBIR relative to other DOD procurement is that projects are smaller in scope, 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. Yet, SBIR retains salient features of defense procurement (Rogerson, 1994, 1995; Lichtenberg, 1995). Defense procurement involves contracting for both R&D and delivery, and the DOD considers multiple prototypes before narrowing the competition for delivery.⁸ Contracts are structured so that firms earn economic profits, providing them incentives for investment in early stages (Rogerson, 1989). Innovation and delivery can be decoupled, and the DOD may contract with separate firms for them. The counterfactuals here speak to all three methods for controlling incentives.

2. Empirical Setting and Data

2.1. Overview of the Navy SBIR Program

The SBIR program is a federal program that provides small firms funding to commercialize early-stage research projects—either on the private market or, as will primarily be the case for this paper, to the government. Federal agencies with extramural R&D budgets of more than \$100 million must allocate approximately 3% (or more) of it competitively through this program to small businesses.⁹ I focus on the DOD—and in particular the Navy—because, unlike other 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 needs of their acquisition programs.¹⁰ Furthermore, the market for technologies produced through the DOD SBIR program is more limited than with other agencies, since they are defense-specific. 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.

The DOD posts 150–250 solicitations for specific research projects each year for each of its main services. These solicitations include a description of the required technology, including fairly detailed technical requirements; goals for Phases I, II, and III; and potential

⁸Early- vs. late-stage competition relates to dual sourcing (Anton and Yao, 1989, 1992; Lyon, 2006).

⁹“Small business” in the DOD setting is often a misnomer, since the limit is 500 employees.

¹⁰See <http://www.navysbir.com/natconf14f/presentations/3-09-Navy-Comm-Williams.pdf>.

to be delivered to the DOD. The Navy in particular connects almost all solicitations to not just systems commands (Naval Air Systems) but also acquisition programs (Virginia Class Submarine). The solicited products are fairly specific to military applications and often are smaller components of major weapons systems.¹¹

Firms interested in competing for a Phase I contract must submit a technical proposal discussing an approach to meeting the goals of the solicitation. 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. 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 [Navy]’s need or requirement.”¹² It involves preliminary prototyping, benchtop testing, computer simulations, and other low-cost preliminary research. The Navy currently awards approximately \$80,000 for the base Phase I contract, and there is little variation across competitors and projects in this amount. Approximately six months after the award date, the firms submit a 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 excludes any consideration of the proposed cost of Phase II research,¹³ in the case of the Navy, the PEO itself is in charge of making Phase I and Phase II selections.¹⁴ 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.¹⁵

Phase II awardees conduct intensive research to test prototypes to assess commercial viability. Contracts are larger (~ \$1 million) and vary considerably both across projects and across competitors within a project: the Navy guidelines note that Phase II is structured in a way “that allows for increased funding levels based on the project’s transition potential.” The firm submits progress reports and a final report after about two years. Finally, 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

¹¹Recent solicitations 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 for Advanced Amphibious Assault and the latter for the Ship-to-Shore Connector.

¹²See the Navy SBIR Program Overview at <http://www.navysbir.com/overview.htm>.

¹³Section 8 of the DOD SBIR solicitation guidelines (<http://www.acq.osd.mil/osbp/sbir/solicitations/sbir20162/preface162.pdf>) notes the primary dimension of evaluation is “the soundness, technical merit, and innovation of the proposed approach and its incremental progress toward topic or subtopic solution,” which is “significantly more important than cost or price.”

¹⁴See <http://www.navysbir.com/natconf14f/presentations/3-09-Navy-Comm-Williams.pdf>.

¹⁵The Phase II desk reference (<http://www.acq.osd.mil/osbp/sbir/sb/resources/deskreference/12.phas2.shtml>) notes, “[DOD] anticipates that at least 40% of its Phase I awards will result in Phase II projects. This is merely an advisory estimate and [DOD] reserves the right and discretion not to award to any or to award less than this percentage of Phase II projects.”

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 set aside specifically for SBIR but is instead funded by the specific acquisition program in charge of the contest. Few contests result in a Phase III contract. 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:¹⁶ technologies developed by Phase II competitors are 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 and location; the topic number (which maps contracts to solicitations and contests); the systems command (“SYSCOM”) in charge of the contract; the phase of the contracts; and dates of execution. It also includes the title and keywords of the proposal from the firm and an abstract of the project as well as a description of the benefit to the Navy. I then match this data using the contract number 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. I restrict the analysis to contests between 2000 and 2012.¹⁷

I collect the full text of all DOD SBIR solicitations and match them to the information above. Each solicitation is a one- to two-page document containing the solicitation title, a very broad technology area, keywords, and the acquisition program in charge. The solicitation also includes a large amount of text describing the project, including an objective, a description of the problem and specific technical requirements, and guidelines for the goals for each phase. This free-flowing text allows me to construct detailed project-level covariates to control for the topic of the contest via an unsupervised machine learning algorithm.¹⁸ I use a Latent Dirichlet Allocation algorithm for topic modeling implemented in MALLET. This algorithm infers topics as collections of words that appear together and then classifies documents as mixtures of topics (Blei, Ng, and Jordan, 2003). Details regarding the topic algorithm, sample topics, and other aspects of data cleaning are in Appendix G.

¹⁶A number of the exceptions in the dataset can be explained by idiosyncratic reasons.

¹⁷Earlier, the Navy was not especially careful about classifying follow-on delivery projects as Phase III, and restricting to projects before 2012 ensures projects have enough time to enter Phase III.

¹⁸This paper contributes to recent work exploring unstructured data. See Bajari, Nekipelov, Ryan, and Yang (2015), Gentzkow and Shapiro (2010), and Hansen, McMahon, and Prat (2014).

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

2.3. Descriptive Statistics

Table 1 shows the distribution of competitors in each phase.¹⁹ About 75% of the contests in the dataset have 2 or 3 Phase I competitors, and fewer than 4% have more than 4. The transition to Phase II is usually not the constraining factor in whether the contest succeeds: over 80% of contests proceed to Phase II, but about 75% of contests that enter Phase II have only one competitor.²⁰ However, very few contests—about 11% of the ones that enter Phase II—lead to a Phase III contract.

Figure 1 shows substantial variation in Phase II and III contract amounts. Phase II contracts can be as small as \$200,000 and as large as \$2 million.²¹ Phase III contract amounts have a long right tail and can exceed \$25 million. The variation is partially due to cross-contest heterogeneity in the types and costs of technologies, but it is also present *within*-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, within-contest. Because this comparison controls perfectly for contest-level heterogeneity, I interpret these differences as suggestive of variation in the value to the Navy of the each competitor’s project.²² These differences can be large: the best-funded competitor often receives more than 50% more funding than the worst-funded competitor.

How does competition affect the probability that the contest transitions into the subsequent stage? Adding a competitor increases the number of draws and should increase the probability that at least one firm succeeds. However, there may be an equilibrium response

¹⁹Table C.1 in Appendix C.1 provides complete summary statistics.

²⁰Note 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.

²¹There 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. While these contracts tend to be earlier in the sample, I have not found any other systematic explanation for these contracts. However, most contracts are for other amounts, and the distribution of residuals when controlling for project-level effects shows no similar masses.

²²This 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, even within contest. 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)
R^2	0.083	0.128			0.133	0.422
N	2773	2292	2773	2292	2292	151

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

in research effort: firms may anticipate a lower probability of capturing the return and thus reduce effort in response, or they may increase effort on the margin in response to the competitive pressure. The net effect is ambiguous.

Table 2 reports OLS regressions of contest-level “success” rates from Phase I to II and from Phase II to III. I run linear probability models of the contest transitioning to a particular phase on measures of funding and competition, controlling for contest-level heterogeneity using year and SYSCOM fixed effects and the topics information. 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 pp—compared to a mean of 83%. Column (2) shows that adding a Phase II competitor is associated with an increase in the probability of transitioning to Phase III by 7.6 pp, a large number

compared to the mean success rate of 10.5%.²³ Somewhat counterintuitively, contests with one more competitor in Phase I have a (slight) lower rate of transitioning from Phase II to III. 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) suggest that this effect is reasonable. If anything, this correlation highlights the endogeneity concern that contests with different numbers of Phase I competitors could be systematically different.²⁴ I allow for this possibility in the structural model.

Columns (3) and (4) investigate the probability that an *individual* competitor generates successful research. Because successes are not observed,²⁵ I use censoring models 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 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 $N_{2j} = \min\{N_{Sj}, \bar{N}_{2j}\}$. I estimate this model via MLE, controlling for the same contest-level covariates, and I report $p(\cdot)$. 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 $\bar{N}_2 = 1$ 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 set $\bar{N}_2 = N_1$. 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 pp. Column (4) indicates that contestants in contests with one additional Phase II competitor have a *higher* probability of success, by 2.8 pp.²⁶ 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) regresses the average Phase II contract per firm within-contest on the number of competitors. Contests with one more Phase I competitor have on average 1.6% larger contracts, which is small (and imprecise).

²³I control for Phase II funding but not Phase I because there is no variation in Phase I funding.

²⁴In principle, this correlation could be explained by stronger competition leading to lower incentives for research, which leads to a lower success rate. This explanation is, however, at odds with the final two columns of Table 2: contests with more Phase I competitors have slightly more funding in Phase II and lead to larger Phase III amounts. Appendix C.1 models the dependence on N_1 more flexibly, and the source of the negative coefficient on N_1 is primarily contests with $N_1 = 4$.

²⁵That 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.

²⁶I model the transition from Phase II to III as follows: a contestant i generates a successful innovation in contest j with probability $p(X_j; t_{ij})$, where t_{ij} is the Phase II funding; if multiple contestants succeed, one contestant is awarded the Phase III contract uniformly at random. In Section 3, I develop a model for how the DOD award the Phase III contract to in case multiple firms succeed. Uniform-at-random simply provides a useful baseline for descriptive analysis.

Adding more Phase II competitors has little impact on average funding, but Appendix C.1 notes a large and significant drop when moving from contests with 3 to contests with 4 Phase II competitors. I also explore implications of Phase II funding: to the extent that Phase II contracts are an indication of latent value (as Section 2.1 suggests that firms with more promising research projects are given more funding) and funding directly increases success probabilities, we would expect that funding correlates positively with success in Phase III. Indeed, Columns (2) and (4) show that increasing funding (contest- and individual-level, respectively) by 10% is associated with an increase in the contest-level success rate of 1.6 pp and in individual success by 2.5 pp. Moreover, Appendix C.1 shows that even *within contest*, firms with larger Phase II contracts are more likely to enter Phase III.

Column (6) regresses the Phase III amount against Phase II 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 amount: if there are multiple Phase II competitors, the DOD can capture a larger portion of the surplus by threatening to go to another competitor with a successful innovation. The predictions related to Phase III contract amounts are therefore threefold: (i) a larger number of Phase I competitors would indicate that firms with more valuable projects survive into later rounds and thus would lead to larger Phase III contracts, (ii) having more Phase II competitors would give the DOD more chances for a lower draw of the delivery cost—and also let it leverage competition—and thus lead to lower Phase III contracts, and (iii) 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 contracts. The coefficients agree with these predictions, although estimates are somewhat imprecise. Adding one Phase I competitor is associated with an increase in the Phase III contract by about 26%. Adding a Phase II competitor is associated with a reduction in the Phase III contract by about 35%. Finally, a 10% increase in average Phase II funding is associated with a 3.3% increase in the Phase III contract.

I use these correlations primarily as motivation for developing a model with features consistent with them. In the model, firms 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 phase. 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.

3. Model

I develop a model of a multistage R&D contest that captures the features of the DOD SBIR program. Section 3.1 presents the primitives and the timing, detailing how research efforts translate to values, costs, and awards. Section 3.2 then discusses two assumptions for how efforts are determined. I highlight the identifying power of each assumption in Section 4.1 and develop an estimation procedure in Section 4.3.²⁷

3.1. Model Timing and Primitives

Each SBIR contest consists of three phases. The primitives 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). In this section I simply discuss how these primitives, together with endogenous efforts, map to outcomes and relegate a discussion of the information structure until Section 3.2, which presents models for how efforts are endogenized.

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 to the DOD. The DOD invites N_1 firms to 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 Phase I, 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

²⁷Section 4.1 shows the weak assumption is sufficient to identify many of the model primitives.

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 $c > v$).²⁸

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 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 effort does not depend on opponents' values, and I thus discuss an effort function $\hat{t}(v)$.²⁹ I begin with the more general assumption.

²⁸It 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, and in general the low success rate suggests it is 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 (Shaked and Sutton, 1984; 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 (by a small amount).

²⁹One institutional justification is that firms know their own values at the start of Phase II but do not know

Assumption M. *The research effort $\hat{t}(v)$ is an increasing function of the value v . This map may depend on all of the primitives of the contest, as well as on the realization of N_2 .*

Assumption M (monotonicity) places no restrictions on Phase I efforts \hat{p} . The restriction that is placed on Phase II efforts 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 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.³⁰ Outside the context of 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.

In this paper, I impose an additional assumption: the contract amounts for Phases I and II coincide with the efforts the firm would choose itself, i.e., the amounts are the *firm-optimal* ones. Below, I specify the firm’s problem to define these amounts and then discuss why this assumption seems appropriate in this setting.

Phase I. Firms are aware of the number of Phase I competitors N_1 , the limit \bar{N}_2 on the number of Phase II competitors, and the primitives of the contest (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 $t^{DOD}(v)$.³¹ It then decides on effort to reduce delivery costs. In doing so, it knows its own value v_i and the number N_2 firms that entered Phase II. However, it knows neither the number of successes N_S nor its opponents’ values.³² It forms beliefs (with cdf $F(\cdot; v_i, N_2, p)$) about its

their opponents’. Firms do not see opponents’ contract amounts initially either.

³⁰This 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 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.

³¹This 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.

³²To justify this informational structure, note that firms see the DOD’s evaluation of its Phase I technical proposal—a proxy for v_i —and are in contact with the DOD liason about it. The number of competitors in Phase II is announced. However, the DOD does not share information about firms with other firms.

opponents' values, where p is its belief of the Phase I effort of each of their competitors.³³ Based on these beliefs as well as its value, it exerts 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

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

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

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

where $G(s; v, t_{N_2}^*(\cdot), p^*)$ is the cdf of a type v competitor's beliefs about the highest surplus

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.

of its competitors. The cdf of the surplus that a type v' firm generates is

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

and the cdf of the maximum surplus of a type- v firm's opponents can be computed by combining (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(ptimality) states that the amounts spent on research—i.e., the amounts that determine the probability of success in Phase I and the distribution of cost draws in Phase II—are chosen by the firm. When taking the model to the data under Assumption O, I will assume that the Phase II research award coincides with this firm-optimal amount as well, so the DOD reimburses the cost of effort. In the case of Phase II, for instance, this amounts to saying that $t^{DOD}(v) = t^*(v)$. While there is admittedly a tension in assuming that the DOD transfer is firm-optimal, this assumption can be justified 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 given that the DOD can only imperfectly monitor how the firms spend the money, the firms can redirect some excess resources. 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 this project.³⁵ Secondly, the DOD actively tries to encourage firms to participate in the defense industrial base through this program, and as such, it would like to limit ex-post losses. Were the DOD to award less than the firm-optimal amount, the firm would try to use money from other sources and suffer losses if the project does not enter Phase III. Even in a setting in which firms may have positive expected profits, Phase III is sufficiently rare that firms may have to enter many contests before realizing a payoff.^{36,37}

The following proposition shows that in many cases Assumption O implies M.

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.*

Intuitively, higher-value firms have a higher probability of winning and 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.³⁸

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.

³⁵It should also be noted that Phase II contracts do not have an explicit profit margin worked into the terms, unlike DOD procurement contracts.

³⁶That firms would substitute internal funds for SBIR funding is consistent with the results of Wallsten (2000). Furthermore, other evidence in this paper suggests that the DOD is generous to firms: for instance, Section 5 estimates that firms capture two-thirds of the surplus. Finally, note that Phase I contracts cover Phase I research costs, the firms do not sustain any ex-post losses throughout the process—not just in Phase II.

³⁷One may also consider a model in which the Phase II amount is set, say, as a convex combination of the firm- and DOD-optimal amounts. Footnote 48 conjectures that such a model would still have empirical content. Throughout the body of the paper, however, I maintain Assumption O as it seems reasonable in this empirical setting. I report results with Assumption M in the Appendix, which would be weaker than Assumption O as well as this slightly more general model of Phase II.

³⁸If 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. I have not been able to find a counterexample where the computed equilibrium is nonmonotone.

4. Identification and Estimation

4.1. Identification

Suppose that in the model in Section 3, we observe the numbers of players N_1 and \bar{N}_2 , 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 $N_2 \leq \bar{N}_2$, and the Phase III contract amount (if the project enters Phase III).³⁹ 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 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 different models. The second part applies to the specific model with Assumption O.

The rationale behind separating the identification argument into two steps is that one can analyze results from the model without imposing the stronger Assumption O and still be confident that the model has empirical content. Indeed, I provide results without Assumption O in Appendix C.3. Moreover, the two steps help motivate the modular estimation procedure discussed in Section 4.2 and highlight that Assumption O can easily be replaced by alternate assumptions if other empirical settings demand it.

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

³⁹Research efforts are measured as the Phase I and II dollar amounts. I discuss identification with and without knowledge of $\psi(p^*)$, because this is set institutionally and exhibits little variation, and thus it may be unrepresentative of the true expenditures on research in this empirical setting.

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.2.

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

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 $(1 - \hat{p})^{N_1}$.⁴⁰ Since the Phase I research effort is also observed, we identify the single point $\psi(\hat{p})$. Variation that affects \hat{p} but not $\psi(\cdot)$ can identify the entire cost function.⁴¹ The following proposition summarizes this 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 $[0, v(t)]$; and (iii) a single point $\psi(\hat{p})$ on $\psi(\cdot)$ is identified (where \hat{p} is the success probability in Phase I), and variation in \hat{p} identifies $\psi(\cdot)$ entirely.*

Interestingly, Assumption M also gives information about a lower bound on 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) instead of exploiting this partial identification argument, I relegate the discussion of the identification of this lower bound from Assumption M to Appendix D.1.

⁴⁰I 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 $\bar{N}_2 = 1$, but it is testable with enough data on contests with $\bar{N}_2 > 1$. Departures from the binomial distribution on N_2 will point towards correlation. 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.

⁴¹We can 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 $\psi(\hat{p})$.

4.1.2. Identification Under Assumption O

Suppose further that the research efforts are set optimally for the firm, as per Assumption O. Then, η 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 Assumption M identifies the marginal benefit of a dollar of research (conditional on η). Any wedge between this and the marginal cost (i.e., a dollar) must be due to the firm not capturing the full surplus through η .⁴²

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.*

In fact, identification of $\psi'(\cdot)$ within a single-parameter family of functions does not

⁴²Mathematically, (5) is an equation in a single variable, although the full argument that it has a unique (relevant) solution is based on rearranging the terms in terms of observables and quantities that have already been identified. See Appendix E.2.

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: the DOD’s willingness to pay for a product must exceed what it pays. This selection happens on a two-dimensional set of Phase II efforts and Phase III transfers instead of being simply based on Phase III transfers, and the point at which selection occurs informs values. The residual variation in the cross-section of Phase III contracts conditional on values is informative of outcomes in Phase II, i.e., delivery costs. Variation in outcomes at the end of Phase I yields differential incentives for research, which allows for identification of the response of delivery costs to research efforts.⁴³

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.⁴⁴ This identification argument is slightly different from ones used in other empirical papers involving Nash bargaining. Grennan (2013) identifies the bargaining parameter roughly by comparing distributions of transfers that are generated by different value distributions (“added value” in his paper) but similar cost distributions: if the transfer distributions change dramatically, then the effect of the value on the transfer—governed by η —would be high.⁴⁵ Crawford and Yurukoglu (2012) identify bargaining parameters by matching the model-implied outcomes to estimated outcomes with auxiliary knowledge about one of the components of the transfer.⁴⁶ I do not have similar knowledge, because delivery costs are nonzero and unobserved in my setting, but unlike Crawford and Yurukoglu (2012), I can leverage the optimality of the investment that I do observe.

Note further that I will use estimates from this model to decompose the effect of increasing

⁴³This intuition highlights that the fact that $H(\cdot; \cdot)$ does not depend directly on v is important for identification. This assumption is reasonable in this empirical setting, in which projects are similar and there likely are not drastically different approaches to research. Otherwise, we would need access to a source of variation that changes either t while holding v constant. An instrument for research costs would be an example, but as discussed in Footnote 45, variation in N_2 is in principle another example.

⁴⁴This identification strategy leverages 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.

⁴⁵In 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 $\bar{N}_2 = 1$.

⁴⁶These “outcomes” correspond to channel input costs, with true marginal cost known to be zero.

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) .⁴⁷ 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 .

Appendix D 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. 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).⁴⁸

4.2. Empirical Model

For each contest, I observe the realized number N_1 and N_2 of contestants in Phase I and II and 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. The research efforts $\psi(p)$ and t are monetary and map to contracts observed in the data. The Phase III contract amount is also observed and maps 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 $t_{\bar{N}_2}^*(v)$ in the model, respectively, as described in Section 3.2.⁴⁹

I add two components to the model in Section 3 to take it to the data: (i) observed

⁴⁷The exception is the distribution of unobserved heterogeneity, which could not be estimated separately when there is a single informative data point in a contest, e.g., when $\bar{N}_2 = 1$.

⁴⁸I conjecture that this set of overidentifying restrictions will allow for identification of a more general model of effort-setting in Phase II as well. Suppose the observed R&D contract is, say, a convex combination of the firm-optimal one and the DOD-optimal one—perhaps to capture a more complex bargaining process for the Phase II amount. In such a model, we would need to identify a bargaining parameter in Phase III, the parameter of the convex combination in Phase II, and—for each research amount in the data—what the corresponding firm-optimal and DOD-optimal ones would be. For a particular level of R&D expenditures, we would have optimality conditions for the firm and the DOD as well as the condition that the research contract is a combination of the firm and the DOD amounts. This gives us three equations for four unknowns, so the parameters are fundamentally underidentified. However, a different level of research expenditure would introduce three new equations (the optimality conditions and the convex combination) but only two new parameters (the firm- and DOD-optimal levels of research at this new value). Thus, exploiting the cross-section of research amounts would offer a battery of extra restrictions. Moreover, the estimation procedure would apply with a modification to Step 4 below. Developing such arguments formally would be an interesting area for future work to apply this model to other settings.

⁴⁹I will not use the observed values of $\psi(p^*)$ in estimation, because they exhibit no variation.

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} , $\tilde{\psi}(\cdot)$, $\tilde{H}(\cdot; \cdot)$, 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 (or more uncertain) value. I set $\tilde{\psi}'(p) \equiv \alpha p$ and estimate α . I restrict η to be constant across contests.⁵⁰

Specification (6) induces a correlation between values, implementation costs, and costs of research: certain projects are more valuable to the DOD but also more costly to implement and research. Controlling for (θ_j, X_j) , however, the residual values \tilde{v} are still mutually independent, 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 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 orthogonal to the research and implementation costs. Adding unobserved heterogeneity also “softens” the hard constraint induced by the fact that Phase III does not happen if $v \leq c$.⁵¹

The multiplicative specification in (6) yields the following property, which follows directly from substitution into the equilibrium conditions (2) and (4).

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.*

Proposition 4 is reminiscent of scaling properties of auction models.⁵² By Proposition 4

⁵⁰The 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. Furthermore, the other parameters (such as α and η) could depend on coarse quantities like N_1 once again without affecting the estimation procedure, although I would need a larger sample size to implement such an estimator.

⁵¹One especially large transfer need not signify that values are high; rather, they may signify that the particular contest in question had a large value of θ_j .

⁵²See Krasnokutskaya (2011) for an example in previous work. Unlike auctions models—in which there is a single dimension of heterogeneity—I let both values and costs scale.

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 $\tilde{t}_{N_2}(\cdot)$. The advantage of this specification is that it allows me to control for heterogeneity tractably by regressing Phase II efforts on covariates.

Distributional Assumptions. I place parametric restrictions to assist in estimation. I assume that (i) V is lognormal with location parameter μ_{N_1} and scale parameter σ_{N_1} ; (ii) $H(\cdot; t)$ is lognormal with mean parameter $\mu(t)$ with $\mu'(t) < 0$ (the parameterization is in Appendix F.2) and scale parameter σ_C ; and (iii) $\psi(p) = \alpha p^2/2$. I place no parametric restrictions on the distribution of θ .

The 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 $\psi(0) = 0$.⁵³ I make this decision because, unlike the Phase II contract amount, the Phase I contract amount is set 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.⁵⁴ 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

⁵³I 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.

⁵⁴Since 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.

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.

A multistep procedure avoids the computational burden of a full solution approach. Furthermore, the procedure is modular: if a researcher is unwilling to impose Assumption O and instead is willing to assume a bargaining parameter, she can terminate the algorithm at Step 3 and still recover meaningful estimates of values and costs without imposing any model on how research efforts are determined.⁵⁵

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, i.e., $\nu_{ij} \equiv \log \tilde{t}_{N_2}(\tilde{v}_i) + \log \theta_j \equiv \log \tilde{t}_i + \log \theta_j$, along with an estimate $\hat{\beta}$ of the impact of the covariates. I then residualize the Phase III transfer by dividing by $\exp(X_j \hat{\beta})$.⁵⁶

Step 2 (Estimating Θ). I use a deconvolution argument from auctions (Li and Vuong (1998), Krasnokutskaya (2011)) to estimate the distributions of Θ and 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 mutually independent and the distribution of θ_j normalized to mean zero,⁵⁷ 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) for estimation, and details are in Appendix F.1.

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

⁵⁵Similarly, Steps 4 and 5 can be replaced by models that satisfy Assumption M but not Assumption O if it is more appropriate in other settings.

⁵⁶I use year fixed effects, SYSCOM fixed effects, and the topics generated from the unsupervised machine learning algorithm as elements of X_j .

⁵⁷Note that t_{i_1} and t_{i_2} would be dependent in cases with selection—albeit in a way that can be modeled—in which case we would have to modify this procedure. However, this step of the estimation also excludes contests with selection.

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 . However, rather than solving the model explicitly—which is computationally cumbersome and also economically strong—I leverage monotonicity to recover an implied effort function. 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 compute the inverse effort function $\tilde{v}(\cdot)$ without solving the model directly, for $\theta_j = 1$. This inverse effort function, together with the multiplicativity assumption on unobserved heterogeneity, then lets me compute the likelihood efficiently. Details are in Appendix F.2.

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 FOC. In this step, I impose the firm’s FOC by solving the Phase II model explicitly. I do so at each value of η on a fine grid, at the estimated parameters from Step 3. I then use a simulated method-of-moments procedure to match the failure rate and the Phase III transfers of the observed data with simulated values from each of the solved models for the various values of η , detailed in Appendix F.3.

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 $N_{2j} = \min\{N_{Sj}, \bar{N}_{2j}\}$. Upon estimating $\hat{p}_{(N_1, \bar{N}_2)}$, I compute the profits from Phase II by solving the model using the estimated parameters from Step 4 for all values of N_1 at the estimated $p_{(N_1, \bar{N}_2)}^*$. I then use the FOC associated with (4) as the estimating equation for α , as discussed in Appendix F.3.

5. Structural Estimates

Table 3 reports the parameter estimates for the equilibrium model, following Steps 1–5.^{58,59} Panel (a) shows quantiles of the distribution of unobserved heterogeneity θ . The distribution

⁵⁸I construct standard errors by a nonparametric bootstrap. I sample with replacement from the dataset, fixing the distribution of (N_1, \bar{N}_2) , and repeat the whole procedure 100 times.

⁵⁹Appendix C.3 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.

Percentile	2.5%	10%	25%	50%	75%	90%	97.5%
θ	0.387 (0.089)	0.728 (0.057)	0.876 (0.036)	1.012 (0.020)	1.165 (0.038)	1.346 (0.124)	1.938 (0.382)

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

Values (\$M)	$N_1 = 1$	$N_1 = 2$	$N_1 = 3$	$N_1 = 4$
Mean	14.17 (8.18)	16.76 (5.68)	18.03 (6.27)	18.39 (5.25)
Standard Deviation	0.44 (0.26)	0.54 (0.24)	0.56 (0.20)	0.57 (0.17)
95% Range	1.74 (1.04)	2.11 (0.95)	2.19 (0.78)	2.25 (0.65)

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

$\Pr(c < v)$		$\mathbb{E}[c c < v]$		Quantiles (\$M)			
Value	Semi-Elasticity	Value	Elasticity	1%	5%	10%	Elasticity
0.067 (0.011)	0.012 (0.005)	9.38 (2.99)	-0.015 (0.006)	4.07 (1.25)	13.24 (4.09)	24.83 (7.69)	-0.159 (0.062)

(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)	0.63	(0.19)
Phase I Cost Parameter (α , \$M)	0.242	(0.130)
Average Phase I Cost (\$M)	0.031	(0.017)

(d) Phase I and bargaining parameters

Table 3: Structural estimates

is fairly concentrated around 1: 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) describes value distributions as a function of N_1 . I scale the estimates of \tilde{v} by the estimated mean of $\theta_j \cdot \exp(X_j\beta)$ from Steps 1 and 2. First, average projects have mean values of around \$14–\$18 million. Projects in which the DOD selects a larger number of Phase I competitors tend to have larger values, although the difference is somewhat imprecise. Second, these value distributions are fairly narrow; standard deviations are about \$450,000–\$550,000. 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, of about \$2 million, or 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 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 (as a function of Phase II effort) 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 below.

Panel (c) 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.⁶⁰ Just as values are positively selected conditional on success, the cost draws are negatively selected; because so few contests succeed in Phase II, the mean unconditional 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 \$9.4 million. The 1st percentile of the *unconditional* cost distribution is about \$4.1 million and the 5th percentile about \$13.2 million. The elasticity of the quantiles with respect to effort is relatively low: if research efforts increase by 1%, the quantiles of the delivery cost distribution decrease by 0.2%.⁶¹ This value translates to an elasticity of about 0.016 for the conditional expectation of costs and a semi-elasticity of 0.012 for the probability that the cost draw is less than the value draw.⁶²

Conditional on η , 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—and 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

⁶⁰I fix a value of unobserved heterogeneity θ and compute moments of the cost distribution at the implied value of $\hat{t}_{2ij} = t_{2ij}/\theta$ for each contestant i in each contest j . I then average across all these data points and integrate out over θ and scale the estimates to millions of dollars.

⁶¹It 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.

⁶²Note that $\Pr(c < v)$ is not exactly a failure rate (although it is quite close), 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 with respect to the research efforts, since v would increase as well. I view these moments as descriptive features of the cost distributions.

the observed transfers to drop.

Panel (d) first reports an estimated bargaining parameter for the firms of 0.63: the DOD gives the winning firm about two-thirds of the (incremental) surplus generated from the project, although this estimate is rather imprecise.⁶³ This estimate directly uses information about the firm choosing research efforts optimally. 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). Using the functional form assumption that $\psi(p) = \alpha p^2/2$, the Phase I expenditure amounts to approximately \$31,000. While this value is slightly lower than the DOD-specified amount of \$80,000, it is 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 suggest 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. I will maintain this functional form, with the caveat that I may be underestimating the cost of Phase I research slightly.

The estimates suggest these contests have moderate value but fairly small variation across competitors within contests, consistent with the notion that these projects are well-specified ex-ante: while there is heterogeneity across contestants, there is not much room for innovation on the dimension of the quality of the proposal at the end of Phase I. Delivery costs, however, are substantially different across firms, although the map from effort to costs 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

I first explore whether the equilibrium of the R&D contests features underprovision or overprovision of R&D from a social standpoint. I discuss the sources of inefficiency to help interpret the design counterfactuals studied in Sections 7–8 and Appendix A. I also compute the optimal social surplus to quantify the surplus left on the table due to the current design of the contest. 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. I initially focus on social surplus as the outcome of interest; Section 9 and Appendix B consider the

⁶³This bargaining parameter feeds into the value and cost estimates, but not the unobserved heterogeneity distributions. Thus, its imprecision feeds into the larger standard errors of other estimates. Standard errors tend to be smaller if conditioning on η , as shown in Appendix C.3.

N_1	\bar{N}_2	Baseline		Phase I			Phase II			Optimum	
		p^*	SS	\hat{p}	SS	%	γ	SS	%	SS	%
1	1	0.33	0.015	0.35	0.016	1.9%	1.80	0.017	13.9%	0.018	17.9%
2	1	0.46	0.065	0.39	0.066	2.8%	1.71	0.073	12.4%	0.074	14.1%
3	2	0.64	0.203	0.58	0.205	0.8%	1.67	0.224	10.5%	0.225	10.6%
4	2	0.64	0.269	0.52	0.284	5.8%	1.65	0.294	9.5%	0.307	14.3%

Table 4: 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.

profits of the DOD instead.

To assess the efficiency of 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 socially 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 equilibrium (holding first-stage behavior fixed). In the next experiment, I compute the socially optimal first-stage entry probability \hat{p} , keeping the second-stage effort function fixed at the equilibrium. I compare this value to p^* . Table 4 shows the surplus in the equilibrium of the contest and the optimal values of γ and \hat{p} and the surplus at these values.⁶⁴

What are the sources of inefficiency in Phase II? If $N_2 = 1$, 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 surplus and effectively amounts to selling the project to the firm, maximizing social surplus. 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, Phase II efforts should be underprovided in the R&D contest when $N_2 = 1$. Accordingly, Table 4 notes that in the cases where $\bar{N}_2 = 1$ (so that N_2 must equal 1 when Phase II occurs), the efficient level of R&D is about 71–80% larger than the equilibrium level of R&D. Across contests with $N_1 = 1$ or 2, this amounts to a gain in social surplus of around 12–14% relative to the equilibrium, a ballpark estimate of the “cost of holdup” in this setting.

A less obvious implication is that a similar conclusion holds for $N_2 > 1$: the social planner’s optimum is supportable by the firms in equilibrium if $\eta = 1$. The key observation is that in Phase II, the winning firm’s 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 social surplus, meaning the firm is rewarded in a manner that coincides with the social planner’s objective function.

⁶⁴I show various values of (N_1, \bar{N}_2) use the parameters for the associated value of N_1 .

Thus, for Phase II, the social planner would always prefer $\eta = 1$. I codify this below.⁶⁵

Proposition 5. *Consider a contest that begins in Phase II. The social planner’s solution (when the 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 4 also shows γ for cases where $\bar{N}_2 = 2$; by Proposition 5, we expect $\gamma > 1$ here as well. I find magnitudes similar to the instances when $\bar{N}_2 = 1$: socially efficient research efforts would be about 66% larger than the ones in the R&D contest, and the social surplus would increase by about 10% off the baseline.

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 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. This *reimbursement effect*, in which later-stage research expenditures are not internalized when early-stage expenditures are decided, would lead to overprovision of Phase I 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 is in principle ambiguous.

Comparing the equilibrium p^* to the optimal \hat{p} in Table 4 suggests that there is usually overprovision of R&D in equilibrium. The sum of the reimbursement effect and the business-stealing effect usually outweighs the holdup effect, and social surplus can increase by 1–6% when reducing p^* to the optimal \hat{p} . The one exception is $N_1 = 1$: here there is no business-stealing effect, and holdup outweighs the reimbursement effect so that R&D is (slightly) underprovided.⁶⁶ The final two columns of Table 4 show the optimal social surplus, subject to the information constraints that the agents face. In particular, I set $\eta = 1$ to maximize surplus in Phase II and then simultaneously choose the effort p in Phase II to maximize the surplus generated in the entire contest.⁶⁷ These columns provide a measure of the surplus left on the table due to the current design of the contest and provide a benchmark

⁶⁵See Hatfield, Kojima, and Kominers (2017) for a discussion of this issue in general models.

⁶⁶This example is somewhat like a proof-of-principle, though, as R&D is overprovided at $N_1 = \bar{N}_2 = 1$ when using estimated parameters from other values of N_1 .

⁶⁷This 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” in Table 4 because it is closer to the current design of the contest.

	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	-0.000	0.115			$N_1 = 2$	-0.001	0.146		
$N_1 = 3$	0.007	0.131	0.220		$N_1 = 3$	0.007	0.165	0.277	
$N_1 = 4$	0.012	0.143	0.249	0.319	$N_1 = 4$	0.013	0.178	0.312	0.393

(a) Change in social surplus (Baseline is 0.126 \$M) (b) Change in total research costs (Baseline is 0.168 \$M)

Table 5: 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.

against which the design counterfactuals in the subsequent sections can be compared. Social efficiency improves by 10%–18% by setting this optimal design.

7. The Effect of Early- and Late-Stage Competition

The addition of a competitor to a contest has two effects. First, there is a *direct* effect of 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. In this section, I quantify the effect on surplus of adding competitors in both the early (Phase I) and late (Phase II) stages of the program.⁶⁸ I then decompose this effect into direct and incentive effects.

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.⁶⁹ At the baseline, the expected social surplus per contest is \$126,000. About \$168,000 of this is due to R&D cost reimbursements, so each contest generates \$294,000 of surplus, ignoring research costs.

Table 5(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 . Fixing $\bar{N}_1 = 1$, social surplus is barely changed if increasing to $N_1 = 2$ and increases by a total of \$12,000 if changing to $N_1 = 4$. 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 (slightly) socially beneficial, as Section 6 shows that there is overprovision

⁶⁸In a multistage contest, the design variables are the number N_1 of competitors in the first stage and the *limit* \bar{N}_2 in the second stage. The counterfactual experiment here is that there is a pool of ex-ante identical competitors that the DOD could choose to let into Phase I. This is reasonable for small values of N_1 : after all, considerably more firms apply to SBIR contests than are awarded Phase I contracts. If marginal firms are weaker, that would be a force against increasing N_1 .

⁶⁹I use the parameter estimates with $N_1 = 4$ in this section.

of R&D in Phase I—but does lead to larger *total* Phase I effort expenditures.⁷⁰ However, much of the failure rate is due to failure in Phase II, and limiting entry to exactly one competitor in Phase II only leverages the benefit of having one additional 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 into Phase II 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 high in this setting, firms are effectively *not* substitutes; the two firms would only be substitutable in the unlikely event that they both succeed and have similar values. 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 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 \$115,000, slightly less than the base of \$126,000. Adding one more competitor to each stage increases it by \$105,000. This decrease is due to firms becoming slightly more substitutable as competition increases.⁷¹ In addition, there are effects on equilibrium incentives that I discuss in Section 7.2, but the fact that efforts increase almost linearly (see Table 5(b)) suggest that they are quite small.

7.2. Decomposing the Effect of Competition

Consider a contest with (N_1, \bar{N}_2) and any outcome $S(N_1, \bar{N}_2, p, \{t_{N_2}(\cdot)\}_{N_2 \leq \bar{N}_2})$, 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 functions $t_{N_2}(\cdot)$. 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 total effect of moving from a contest with one contestant to one with (N_1, \bar{N}_2) is

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

⁷⁰Moving from $N_1 = 1$ to $N_1 > 1$ does introduce the business-stealing effect in Phase I, which leads to further overprovision of Phase I R&D.

⁷¹Note 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.

	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	-0.056	-0.056			$N_1 = 2$	-	0.176		
$N_1 = 3$	-0.137	-0.137	-0.137		$N_1 = 3$	-	0.224	0.368	
$N_1 = 4$	-0.224	-0.224	-0.224	-0.224	$N_1 = 4$	-	0.239	0.441	0.562
(a) Direct (Phase I)					(b) Direct (Phase II)				
	$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$		$\bar{N}_2 = 1$	$\bar{N}_2 = 2$	$\bar{N}_2 = 3$	$\bar{N}_2 = 4$
$N_1 = 2$	0.056	-0.002			$N_1 = 2$	0.000	-0.002		
$N_1 = 3$	0.144	0.047	-0.005		$N_1 = 3$	0.000	-0.002	-0.006	
$N_1 = 4$	0.236	0.131	0.039	-0.008	$N_1 = 4$	0.000	-0.003	-0.007	-0.010
(c) Incentive (Phase I)					(d) Incentive (Phase II)				

Table 6: 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 (8). All values are in millions of dollars, and the baseline value of social surplus (at $N_1 = \bar{N}_2 = 1$) is \$126,000.

$$\begin{aligned}
& + \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_1^*(\cdot)\}\right) - S\left(N_1, \bar{N}_2, p_{(1,1)}^*, \{t_1^*(\cdot)\}\right)}_{\text{incentive effect from Phase I competition}} \\
& + \underbrace{S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_{\bar{N}_2}^*(\cdot; p_{(N_1, \bar{N}_2)}^*)\}_{N_2 \leq \bar{N}_2}\right) - S\left(N_1, \bar{N}_2, p_{(N_1, \bar{N}_2)}^*, \{t_1^*(\cdot)\}\right)}_{\text{incentive effect from Phase II competition}}. \quad (8)
\end{aligned}$$

In words, the direct effect of Phase I adds Phase I competitors without changing efforts. The direct effect of Phase II subsequently increases the limit on Phase II contestants, again without any change in equilibrium efforts.⁷² The incentive effect from Phase I allows firms to adjust their research efforts *in Phase I* to the equilibrium effort given by the new competitive structure. The incentive effect from Phase II allows firms to adjust their Phase II efforts and arrives at the new equilibrium.

Table 6 quantifies these four effects. Panel (a) shows the direct effect of adding Phase I competitors, which is definitionally independent of \bar{N}_2 . 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, this effect is negative and substantial, from \$56,000 to $N_1 = 2$ to \$224,000 for $N_1 = 4$. Panel (b) shows the direct effect of increasing entry into Phase II. This effect is even larger and *positive* (but definitionally 0 for $\bar{N}_2 = 1$). 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 positive as long as $\bar{N}_2 > 1$.

Panel (c) of shows the incentive effect for Phase I. Phase I effort decreases with N_1 and

⁷²In 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.

N_1	\bar{N}_2	η^*	Baseline	η^*	% Increase to η^*	% Increase to Opt
1	1	0.72	0.015	0.016	4.8%	17.9%
2	1	0.65	0.065	0.065	0.9%	14.1%
3	2	0.65	0.203	0.203	0.1%	10.6%
4	2	0.59	0.269	0.270	0.6%	14.3%

Table 7: 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 $\eta = 0.63$ 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 4 and reports the percent increase in social surplus from changing to the optimum.

increases with \bar{N}_2 . 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. The Phase I incentive effect on social surplus, which is (usually) large and positive, is increasing in N_1 but decreasing in \bar{N}_2 . 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.

In short, the planner prefers invite more firms to enter *both* phases, and the main benefits come from the direct effect in Phase II and the incentive effect in Phase I.

8. The Effect of the Bargaining Parameter

The share η of surplus firms receive provides a second way to control competition without resorting to finding more competitors—which may be costly or impossible, especially if there are few firms capable of conducting specialized research. In this section, I fix estimates of values and costs and vary η to identify to what extent we improve efficiency purely by changing the rewards the firms earn from procurement.

Is it possible for the social planner to face a nonmonotonicity in η ? Increasing η ameliorates the holdup problem by giving the firm a greater claim to the surplus. Since this is the only inefficiency in Phase II, Proposition 5 notes that increasing η is unambiguously beneficial for social surplus in Phase II. However, larger η increases both the business-stealing and reimbursement effects in Phase I. Research is already overprovided in Phase I. Thus, increasing η could exacerbate these two effects to the point where they overshadow the gain from addressing the holdup problem.

Fundamentally, η is simply one lever that simultaneously controls two inefficiencies: increasing η increases Phase II research (socially beneficial) but also Phase I research efforts (socially harmful). Table 7 shows the value η^* that maximizes social surplus for a number of

different values of (N_1, \bar{N}_2) .⁷³ Interestingly, η^* is actually rather close to the point estimate of 0.63: the DOD is setting η close to a point where these two effects approximately cancel each other.⁷⁴ In the cases where $N_1 > 1$, social surplus can be improved by at most 1% by changing η . I find that the DOD could improve surplus for $N_1 = 1$ by setting a *higher* η : this reflects the result in Table 4 that $N_1 = 1$ had the largest cost of holdup. In all cases, it seems the social planner cannot improve surplus much further purely through η : the final two columns of Table 7 shows that the gains are modest relative to moving to the social optimum in which both inefficiencies can be rectified separately.

I briefly make one point related to firm and DOD profits. Varying η yields a natural “Laffer” curve for the DOD profits. If the firm is not promised any part of the surplus ($\eta = 0$) and thus has no incentives to exert effort, DOD and firm profits are both 0. Setting $\eta = 1$ will give the firm high-powered incentives, but the DOD will not capture any of the surplus. An interior value of η will thus optimize DOD profits; values less than this amount are actually *Pareto* inefficient for the firms and the DOD. In Appendix B.2, I estimate this Laffer curve and provide evidence that the current design lies on the efficient side (even in the case $N_1 = 1$ when modest gains in social surplus are possible from increasing η).

9. DOD Profits Under Alternate Contest Designs

Finally, I extend the analysis beyond social surplus by considering two objective functions for the DOD. The first one is a natural measure of DOD profits: the value of project the DOD acquires in Phase III, less the Phase III contract amounts, less total expenses it pays along the way (Phase I and II R&D contracts and prizes, if relevant).⁷⁵ I also consider “Phase III DOD profits,” which is simply the value of the product less Phase III contract, and it ignores research costs and prizes.⁷⁶

Table 8 lists outcomes for alternate contest designs, using contests with $\eta = 0.63$, $N_1 = 4$, and $\bar{N}_2 = 2$ as the baseline. I collect results for setting \bar{N}_2 optimally (Section 7), setting η optimally (Section 8), mandating firms share intellectual property after Phase I (Appendix A),⁷⁷ and implementing the optimal design (Section 6).⁷⁸ The first column summarizes results on social surplus: η was already near optimal so the gains are small, IP

⁷³See Appendix B.2 for plots of social surplus against η .

⁷⁴There is nothing in the estimation procedure that pushes towards η being close to optimal.

⁷⁵If $\bar{N}_2 = 1$, 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.

⁷⁶This provides an interesting comparison for institutional reasons: the DOD must spend approximately a fixed proportion of its R&D budget on Phase I and Phase II research, and thus the surplus generated in delivery (the “bang for the buck”) may be independently of interest.

⁷⁷I model IP sharing by saying that all firms get access to project plans with the highest value v at the end of Phase I. Firms may be compensated by prizes from the DOD to incentivize effort to counteract the induced free-rider effect, and I choose the socially optimal level of this prize.

⁷⁸The DOD sets $\eta = 1$, sets \bar{N}_2 optimally, and charges firms to enter Phase II.

Design	Social Surplus (\$M)	DOD (\$M)	Phase III DOD (\$M)
Baseline	0.269	-0.113	0.233
$\eta = \eta^*$	0.270	-0.067	0.241
IP Sharing	0.314	-0.109	0.258
$N_2 = \bar{N}_2^*$	0.445	-0.168	0.393
Social Optimum	0.499	-0.389	0.066

Table 8: 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.

sharing has a larger effect, but the largest gains are from allowing more Phase II competitors or setting the social optimal design. The second column tabulates DOD profits. At baseline, the DOD runs a loss of \$113,000 per contest; about \$346,000 is due to research reimbursements, meaning Phase III profits are \$233,000. Like the planner, the DOD internalizes the *full* costs of research but, unlike the planner, it internalizes only one-third of the surplus generated in delivery.

The two changes that are most beneficial for social surplus harm DOD profits considerably. At the social optimum, firms have a larger incentive to exert effort because $\eta = 1$, so the DOD pays a larger amount to reimburse these research efforts and recovers little from the procurement contract. Losses increase almost threefold.⁷⁹ Setting $N_2 = \bar{N}_2^*$ increases the DOD’s losses by about 50% when accounting for research efforts. Accordingly, a DOD that maximizes profits would not want to move from the baseline to either of these designs that have significant social benefits.

The two designs yielding more modest improvements in social surplus are (at least slightly) beneficial to DOD profits. IP sharing increases surplus by 17% but leaves DOD profits approximately unchanged. Furthermore, reducing η slightly to the socially optimal one of 0.59 does increase DOD profits. 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. While this raises the question of why the DOD does not implement a design change that is both socially beneficial and improves its own profits, it is important to remember that reducing η would harm *firm* profits. The DOD may well incorporate firm profits into its objective—perhaps as a mechanism to ensure that these small businesses stay as part of the defense industrial base.

⁷⁹If only one firm succeeds in Phase III, the DOD earns nothing from the delivery process because it pays the firm its value. If multiple firms succeed, the DOD recovers the inframarginal surplus generated by the firms, but the winning firm captures the entire incremental surplus it generates.

The takeaway from this summary table is that while some design changes can improve both social surplus and DOD profits, the two objectives seem at odds when considering the designs that are most beneficial to surplus. This stems from the DOD capturing only one-third of the surplus from procurement to ensure that firms have the incentives to conduct research—while also paying out the full research costs through R&D contracts. Indeed, paying out research costs is substantial: the final column of Table 8 shows that Phase III DOD profits are aligned (almost) fully with the social planner’s objective at these parameters,⁸⁰ which need not be the case for general parameters. 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 often select starkly different ones. Details are provided in Appendix B.

10. Conclusion

This paper proposed a model of R&D contests incentivized by procurement contracts. I provided constructive identification for the distributions of values, research costs, delivery costs, and the share of the surplus the firms receive. I leveraged a monotonicity condition that firms with higher-value projects spend more on research, a selection condition that the procurer does not contract with a firm generating negative surplus, and an optimality condition that research efforts are set optimally. I find evidence that in the DOD SBIR program, most of the uncertainty in the research process happens in the late-stage “development” phase and that firms are able to capture about two-thirds of the surplus generated. Moreover, R&D efforts are underprovided in the late phase but overprovided earlier on.

Social surplus would improve by adding more contestants, due to a combination of the direct effect of more draws in the late phase and an indirect incentive effect of firms adjusting their efforts in the early phase. Changing the firms’ share of the surplus in procurement bears limited gains. Mandating that firms share intellectual property can improve social surplus. The DOD, however, would often be at odds with the social planner in terms of whether it would prefer a particular design change.

I envision two avenues for future work. First, this paper adopts a short-term view of the benefits to firms from the DOD SBIR program: I take the procurement contract as the sole source of incentives. One may wonder about other benefits that may accrue many years down the line.⁸¹ While it is beyond the scope of the current dataset and paper to study the effects of the contest on the life cycle of defense firms, an advantage of the approach is that richer

⁸⁰The exception is the optimum, when $\eta = 1$ and the DOD rarely extracts any surplus in procurement.

⁸¹A natural concern is that a firm may win non-SBIR contracts (or subcontracts) many years later for tangentially related research. Investigating this requires a much more extensive dataset on defense contracts and a way to understand the specific technologies used. Other concerns, such as potential for M&A or non-DOD commercialization down the line, are anecdotally less important.

objective functions—incorporating payoffs beyond the procurement contract—can readily replace the optimality condition adopted in this paper. Indeed, this observation relates to a second avenue for research. Other settings (e.g., FDA trials or venture capital funding) can be conceptualized as multistage contests, and natural models may well have analogues to the monotonicity, selection, and optimality conditions leveraged here. Considering both the theoretical interest in them and their empirical relevance, contests have been understudied in the structural literature. The optimal design of such contests is an empirical question that depends on primitives, and this paper provides a roadmap for how to understand these primitives.

References

- 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.
- 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 (2015): “Machine learning methods for demand estimation,” *American Economic Review*, 105(5), 481–485.
- 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., 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.
- 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.
- DREZNER, J. A., AND M. HUANG (2009): *On Prototyping: Lessons from RAND Research*. RAND Corporation.

- FULLERTON, R. L., AND R. P. MCAFEE (1999): "Auctioning entry into tournaments," *Journal of Political Economy*, 107(3), 573–605.
- 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*, 106(12).
- GRENNAN, M. (2013): "Price discrimination and bargaining: empirical evidence from medical devices," *American Economic Review*, 103(1), 145–177.
- GROSS, D. P. (2016): "Creativity under fire: the effects of competition on creative production," Working paper, Harvard Business School.
- (2017): "Performance feedback in competitive product development," *RAND Journal of Economics*, 48(2), 438–466.
- HANSEN, S., M. MCMAHON, AND A. PRAT (2014): "Transparency and deliberation within the FOMC: a computational linguistics approach," *Quarterly Journal of Economics*, Forthcoming.
- 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 (2017): "Strategy-proofness, investment efficiency, and marginal returns: an equivalence," Working paper.
- HOWELL, S. T. (2017): "Financing innovation: evidence from R&D grants," *American Economic Review*, 107(4), 1136–1164.
- KIREYEV, P. (2016): "Markets for ideas: prize structure, entry limits, and the design of ideation contests," Working paper 16-129, Harvard Business School.
- KOH, Y. (2017): "Incentive and sampling effects in procurement auctions with endogenous number of bidders," *International Journal of Industrial Organization*, 52, 393–426.
- 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.
- LEMUS, J., AND G. MARSHALL (2018): "Dynamic tournament design: an application to prediction contests," Discussion paper, UIUC.
- 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.
- LIU, B., AND J. LU (2018): "Pairing provision price and default remedy: Optimal two-stage procurement with private R&D efficiency," Discussion paper, National University of Singapore.
- 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.
- MCMILLAN, J. (1990): "Managing suppliers: incentive systems in Japanese and U.S. industry," *California Management Review*, 32, 38–55.
- 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.
- 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.
- 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.