Optimal Research Contracting with Observable Firm Characteristics

Yitong Wang*

Abstract

Government research funding programs face a fundamental challenge: firms' true research ability is private information. This paper shows that mechanisms based on observable firm characteristics (OFC) significantly improve the performance of NASA's Small Business Innovation Research (SBIR) program. Empirically, I find that patent intensity and past SBIR success rates are strong positive predictors of research ability, yet they receive minimal weight in the current contract. Moreover, the optimal OFC mechanism captures 61% of potential government payoff gains—defined as the gap between the pooling benchmark and the first-best—compared to only 5% under the current contract. These improvements arise primarily from reallocating budget resources away from high-ability firms, which would participate regardless, and toward lower-ability firms with higher participation elasticities.

Keywords: Mechanism design, asymmetric information, observable firm characteristics, innovation policy, Small Business Innovation Research (SBIR)

JEL codes: O38, D82, O31

*Department of Economics, Texas A&M University. E-mail: wangyitong@tamu.edu.

1 Introduction

Governments often sign research contracts with firms to support innovation and meet technological needs, yet studies have questioned the effectiveness of these investments (David et al., 2000; Jacob and Lefgren, 2011). A key challenge is that the government cannot directly observe firms' research ability when allocating contracts (Aboody and Lev, 2000; Akcigit et al., 2022). Publicly observable firm characteristics may contain information about underlying ability, but which signals matter and how to optimally incorporate them remains an open question. This paper addresses this question in the context of the Small Business Innovation Research (SBIR) program, one of the largest public R&D programs in the U.S.

Specifically, I propose a mechanism that optimally incorporates observable firm characteristics (OFC)—including patents, past contract performance, firm size, and firm age—into contract design, and compare it with the current SBIR contract. The OFC mechanism is derived from a contract design framework with a principal and heterogeneous firms (agents). Each firm has private research ability, and the principal observes signals that provide a posterior belief about this ability. A successful innovation generates a social value, while each firm incurs a research cost—both are publicly observable. The principal offers success-contingent contracts, paying only upon innovation completion, and maximizes expected payoff subject to a budget constraint.

For comparison with the OFC mechanism, I first recover the current contract's allocation rule, as the program has not explicitly announced it. Specifically, I estimate a mapping from firm characteristics to awards using machine learning methods, achieving a close fit with observed awards (mean difference $< 0.1\%$). The estimated rule reveals that the current contract depends primarily on calendar year, firm age, and firm size, while innovation-related characteristics, such as patent intensity and past success rates, play minimal roles. Yet empirically, I find that these innovation-related characteristics positively predict research ability, suggesting room for improvement.

To quantify the potential gain from optimally incorporating observable firm characteristics, I simulate each mechanism while holding the budget constant at the 2022 SBIR level. I find that the OFC mechanism captures 61.25% of the potential government payoff gain, while the current contract achieves only 4.79%.[1] The poor performance of the current contract comes from its underutilization of innovation-related firm characteristics documented above, resulting in a highly concentrated award distribution that

[1]The maximal potential improvement is the gap between the pooling benchmark and the first-best outcome, which are the lower bound and the upper bound of government payoffs, respectively.

treats firms nearly identically. The OFC mechanism improves upon this by optimally incorporating these characteristics to infer research ability and design contracts accordingly.

I further examine the sources of this improvement in Section 6.2. The gains primarily come from reallocating budget from high-ability firms to low-ability firms whose participation decisions are more sensitive to awards. Specifically, 76% of firms (211 firms) receive increased awards, while 24% (68 firms) receive decreased awards. Firms receiving increased awards have lower participation rates but higher participation elasticities, so small increases in awards substantially raise their participation probabilities. The improvement decomposes into an "entry effect" of $35.78 million from increased participation probability among low-ability firms, an "exit effect" of -$11.27 million from reduced participation probability among some high-ability firms, and net "allocation effects" of $2.19 million from payment efficiency changes. The "entry effect" dominates, showing that improvement comes primarily from expanding the participant pool rather than optimizing payments to existing participants. Moreover, these efficiency gains are broadly consistent with other Congressional objectives: fostering participation by small businesses and underrepresented firms, as well as encouraging new entrants.

These findings carry direct policy implications (Section 6.3). The OFC mechanism's advantage grows with budget size: its relative improvement rises from 61% at the current NASA budget to 75% under a $30 million expansion. Since the Department of Defense operates a substantially larger SBIR program under the same Congressional mandate, potential efficiency gains there could be even greater. More broadly, the approach applies to any public R&D procurement setting in which the government contracts with privately informed firms and observes noisy signals of their research ability.

Related Literature. The research focuses on the SBIR program, which aims to support small businesses. Although large firms such as Apple, Tesla, and Boeing contribute significantly to global technological progress, small businesses also play a crucial role in innovation. Small firms often serve as sources of breakthrough innovations and technological disruptions due to their flexibility, willingness to take risks, and ability to pursue niche markets that larger firms may overlook (Acs and Audretsch, 1988, 1990). They also contribute to innovation through their role as suppliers to larger firms and through knowledge spillovers to the broader economy (Stam and Wennberg, 2009; Fischer and Reuber, 2003). Moreover, small businesses are particularly important in the early stages of technology development, where uncertainty is high and commercial viability is unclear, making them natural partners for government research programs seeking to explore

emerging technologies.

Government-sponsored research funding, including grants and research contracts such as the NASA SBIR, plays an important role in stimulating innovation, particularly among small firms (Link and Scott, 2010). However, the impact of research grants and contracts on innovation outcomes remains debated, with estimated effects often modest. For instance, investigating twenty years of NIH funding, Jacob and Lefgren (2011) find that an average NIH grant of approximately $1.7 million results in only about 1.2 additional publications, and merely 0.2 first-author publications, suggesting a limited return on investment in terms of academic output. Similarly, in a cross-country analysis of universities in eight European nations, Auranen and Nieminen (2010) report no straightforward connection between the volume of public R&D funding and publications, further highlighting the heterogeneity of outcomes across different institutional contexts. Furthermore, David et al. (2000) review a wide body of econometric studies and conclude that public R&D investments sometimes complement, but often crowd out, private sector research activities, leading to ambiguous net effects on overall innovation.

A major challenge underlying these limited results is the difficulty in identifying high-potential applicants. As emphasized by Aboody and Lev (2000) and Akcigit et al. (2022), firms' research ability is private information to the government. Thus, the selection of truly high-potential innovators is inherently constrained by asymmetric information.

To mitigate these information problems, peer evaluation is the standard method used to assess grant or contract applications. However, the effectiveness of peer review is increasingly questioned. Li and Agha (2015) show that peer review scores are associated with better future research outcomes primarily among high-potential applicants, while offering limited discriminatory power among low-potential applicants. In contrast, Fang et al. (2016) argue that even among top-scoring applications, peer review scores have limited ability to predict future scientific productivity, suggesting a broader concern about the predictive validity of peer evaluation. Moreover, Pier et al. (2018) document substantial disagreement among NIH grant reviewers evaluating the same proposals, raising further concerns about the reliability of peer assessment.

Given these challenges, another promising selection approach is to utilize applicants' past observable performance as signals of future innovation success. This idea is supported by the literature on the "Matthew Effect" in innovation, which posits that initial success tends to reinforce future achievements. Howell (2017) finds that SBIR grant recipients are significantly more likely to secure private investment and produce future innovations, suggesting that early success signals underlying research ability. Bol et al. (2018) study Dutch science funding data and find that researchers who received funding

in the initial stage were more likely to receive more funding later than those who did not, even after controlling for other variables. This is direct evidence of the classic Matthew effect in scientific research funding. In addition, Li and Agha (2015) also document that peer-review scores in NIH research funding evaluations are strongly correlated with applicants' prior publication performance, implying that even formal evaluation systems implicitly rely on historical achievements as indicators of future potential.

From a theoretical perspective, it is common to use observable outcomes to gradually reveal private information about an agent's ability. For instance, Holmström (1999), in his dynamic incentive models, assumes that a manager's unobservable talent could be learnt through observed output over time. Similarly, Laffont and Tirole (1986) use observed accounting cost reports as imperfect signals of a firm's true efficiency, and show how a regulator can design dynamic incentive contracts to gradually screen and discipline firms based on these noisy observations. This theoretical framework suggests that, in practice, external observations can help mitigate information asymmetry. In the context of innovation policy, identifying relevant historical performance metrics provides a promising approach to reducing information asymmetry regarding an applicant's underlying research ability.

Empirical studies on business innovation frequently use patents as signals of research ability (Hall et al., 2001, 2005; Aghion et al., 2013; Howell, 2017; Acemoglu et al., 2018). Patents capture the tangible output of R&D processes and are publicly observable prior to contracting, making them well-suited as one of the observable firm characteristics in this paper's framework.

Despite these insights, a significant gap remains. In the broader context of public procurement, Szucs (2024) shows that bureaucratic discretion leads to the selection of less productive contractors, generating inefficiency through misallocation of contracts. The present paper identifies a different source of inefficiency: the current SBIR contract funds high-ability firms but provides awards exceeding their participation thresholds, leaving insufficient budget to attract additional participants. Bhattacharya (2021) structurally estimates the DoD SBIR procurement process, but focuses on contest design rather than on how observable firm characteristics can be optimally incorporated into contract allocation. More broadly, existing studies have not systematically evaluated the predictive power of past performance signals—such as patents and past contract performance—on firms' underlying research ability, nor how optimally incorporating these signals into contract design can improve program outcomes.

This paper focuses on the NASA SBIR program and provides empirical evidence that patent intensity, past contract performance and firm size could positively predict research

ability. Most importantly, this paper finds that adopting the OFC mechanism could substantially increase the government payoff while preserving access for underrepresented firms, new entrants, and small businesses.

The rest of the paper is organized as follows. Section 2 describes the NASA SBIR data and observable firm characteristics. Section 3 develops the theoretical framework for contract design under asymmetric information. Section 4 estimates the current contract's allocation rule. Section 5 estimates distributions and calibrates model parameters. Section 6 presents results, including counterfactual comparisons, decomposition of improvement sources, and policy implications. Section 7 concludes.

2 Context and Data Sources

R&D plays a central role in economic competitiveness and national security and contributes to addressing a range of global challenges. The SBIR program, mandated by Congress, is one of the largest examples of U.S. public-private partnerships (National Academies of Sciences, Engineering, and Medicine, 2016). The program allocates federal research funding to domestic small businesses to foster innovation aligned with national priorities. Federal agencies are required to allocate at least 3.2% of their extramural R&D budgets to this program. NASA is the second-largest contracting agency in this program, allocating 3.67% of its R&D budget to SBIR in 2022. Crucially, NASA maintains detailed payment records that enable tracking of contract completion at the firm level, making it an ideal setting for studying research contracting under asymmetric information.

The NASA SBIR program operates through a two-phase process. Each year, NASA releases solicitations targeting specific research needs. Phase I awards provide up to $150,000 over a six-month period to establish technical feasibility. Phase I participants may apply for Phase II awards, which provide up to $850,000 over a 24-month period to further develop, demonstrate, and deliver technology innovations. During Phase II, as described by NASA, "prototypes are built and tested." In each cycle, NASA selects approximately 150 small firms to receive Phase II contracts, with payments conditional on meeting contract requirements.

This study focuses on Phase II contracts and treats research costs and commercial value as public information. This is because firms disclose detailed cost breakdowns in their proposals, including labor costs, general and administrative expenses, material costs, consulting fees, and overhead; and NASA, as the primary customer, possesses full knowledge of the commercial value of each project. However, firms' true research ability remains unobservable, posing a challenge for contract design within the program.

To address this challenge, this study uses observable firm characteristics to infer research ability. These characteristics include the number of prior SBIR contracts, historical SBIR performance, patents, firm size, firm age, and firm classifications, which together provide a measurable profile of firms' innovative potential. All variables used to predict a firm's research ability are realized before the contract is awarded and are publicly observable.

This paper combines multiple datasets to construct a comprehensive picture of firm characteristics and project outcomes. The primary source is the NASA SBIR Award Database (1984-2024), which provides contract-level details, including contract ID, award date, end date, award amount, awarded company's name, address, the number of employees and firm classification, such as whether the firm is Historically Underutilized Business Zone (HUBZone) owned, socially and economically disadvantaged, or women-owned.

Project success is measured using the USAspending database, which tracks federal contract outlays. When firms meet the requirements outlined in their SBIR contracts, payments are disbursed and reflected as the outlayed amounts. A project is identified as successful if the firm receives at least the contracted amount, indicating fulfillment of contractual commitments.

The analysis is restricted to Phase II contracts awarded from 2018 to 2022 with complete payment data, yielding 804 contracts. This restriction is due to two reasons. First, although NASA has maintained the SBIR Award Database since 1984, payment information is only available beginning in 2014 and becomes reliable starting in 2018. Payment data from 2014 to 2017 contains many incomplete records, limiting the reliability of outcome measurement. Second, the 24-month duration of Phase II contracts means those awarded after 2022 are largely incomplete or have incomplete payment records. For instance, only 31 out of 149 contracts awarded in 2023 have disclosed payment data, illustrating the incomplete nature of more recent observations.

Using this combined dataset, I construct measures of firms' prior SBIR contract experience and performance. "Past Contract Count" is the total number of SBIR contracts awarded to a firm before the current contract. In the analytical sample, 352 contracts are associated with firms that have observable past contract records. For these firms, "Past Success Rate" is the share of earlier contracts that resulted in full payment. For firms without prior contracts, this variable is undefined. This variable reflects how well the firm has performed in the SBIR program in the past and serves as a signal of its research ability.

Patents are also widely used as signals of research ability. Patent data comes from the

PatentsView database, matched to SBIR contracted firms using fuzzy string matching algorithms following Hall et al. (2005). For each contract, I obtain the number of patents granted to the firm prior to the contract award year. In total, 348 contracts are matched to at least one historic patent. Contracts that cannot be matched to any patent record are assigned zero previous patents.

Additional observable firm characteristics are obtained from the OpenCorporates database by matching the firm name and address.[2] From this source, I collect firm's incorporation date, branch status—whether the registered entity is a headquarters or a local branch, and corporate type—limited liability company, corporation or professional entity. Firm age is calculated as the difference between the award year and incorporation date. Patent intensity is computed relative to firm age to normalize for operational duration.

Figure 1 illustrates the distribution of the number of previous SBIR contracts since 1984 for each firm. Approximately 57% of firms are new entrants, while 43% are experienced firms. The final analytical sample consists of 804 Phase II SBIR contracts awarded between 2018 and 2022 with complete payment information and observable firm characteristics. For experienced firms in the final sample, project success is correlated with firms' past success rate at 19%. This pattern suggests that contract history contains valuable information about firms' research ability, motivating the use of past performance as a predictor of future success.

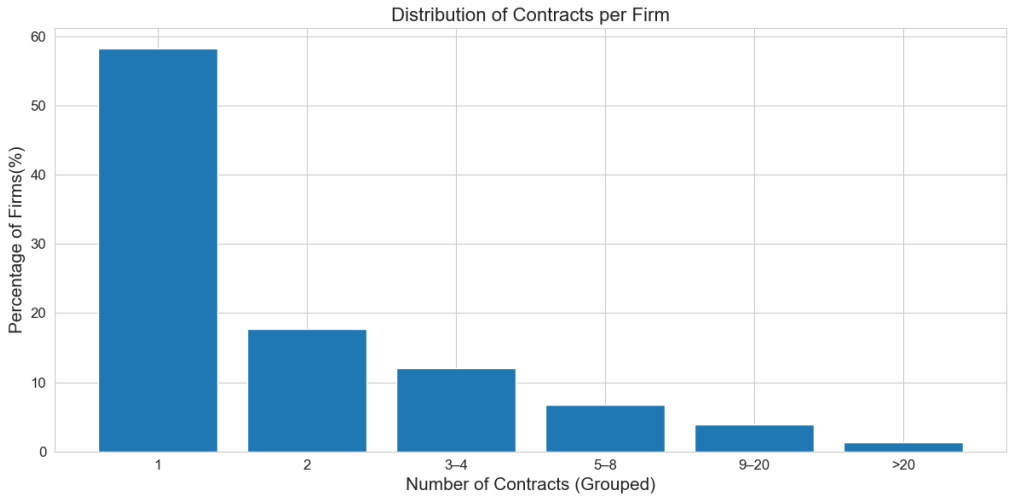


Figure 1: Distribution of Previous SBIR Records

[2]Data sourced from OpenCorporates (`https://opencorporates.com`), the open database of the corporate world, under the Open Data Commons Attribution License: `https://opendatacommons.org/licenses/by/1.0/`.

Table 1: Summary Statistics

Panel A. NASA SBIR Program Data

	Type	Mean	SD	Min	Max	Obs.
Past Contract Count	Count	1.063	1.745	0	11	804
Firm Size	Count	43.406	60.317	1	406	802

Panel B. Non-SBIR Data

	Type	Mean	SD	Min	Max	Obs.
Success	Binary	0.473	0.500	0	1	804
Past Success Rate	Ratio	0.227	0.344	0	1	352
Previous Patents	Count	7.878	18.393	0	231	804
Firm Age (Years)	Count	17.426	14.001	0	97	795
Patent Intensity	Count	0.876	9.071	0	190	795

Table 1 presents summary statistics. The average firm has 1.06 prior SBIR contracts and 7.88 patents. Only 47.3% of projects achieve success, underscoring the need for performance improvement through better selection mechanisms.

3 Theoretical Framework

This section develops a contract design framework to evaluate how public signals about firms' research ability can improve the government's payoff from research funding under a fixed budget.

Three informational settings span the range of possible government payoffs: (1) a first-best benchmark, where each firm's research ability is fully observable, establishing an upper bound on government payoffs; (2) a pooling benchmark, in which the government observes no information about firms, yielding a lower bound; and (3) an Observable-Firm-Characteristics (OFC) mechanism, which optimally designs contracts based on noisy public signals about firm ability. These theoretical benchmarks, analyzed in Sections 3.1–3.3, provide a framework for evaluating the current contract, which is estimated from data in Section 4.

The economy consists of a government (principal) and a continuum of heterogeneous firms (agents) with unit mass. Each firm is characterized by a pair $(\lambda, \mathbf{s})$: a research

success probability $\lambda \in [0,1]$, which captures the firm's underlying research ability and is its private information, and a vector of observable firm characteristics $\mathbf{s} \in \mathbb{R}^d$, including past contract performance, patent intensity, among others, as described in Section 2. The pair $(\lambda, \mathbf{s})$ is drawn from a joint distribution with marginal density $f(\lambda)$ and $g(\mathbf{s})$, and conditional density $h(\lambda|\mathbf{s})$, satisfying $f(\lambda) = \int h(\lambda|\mathbf{s})\, g(\mathbf{s})\, d\mathbf{s}$. The vector $\mathbf{s}$ serves as noisy public signals about λ: upon observing $\mathbf{s}$, the government updates its belief to $h(\lambda|\mathbf{s})$. What varies across the settings analyzed below is the government's information set: it may observe both λ and $\mathbf{s}$ (Section 3.1), observe nothing (Section 3.2), or observe the public signals $\mathbf{s}$ (Section 3.3).

A firm that undertakes research incurs a fixed cost $\theta > 0$ and, with probability λ, delivers a successful innovation that generates social value $\gamma > 0$; both θ and γ are publicly observable (see Section 2). The government offers contracts with success-contingent payment $A(\cdot)$, paying only upon innovation completion; this contract form reflects the institutional design of the SBIR program, where Phase II payments are conditional on successful completion (see Section 2). The award $A(\cdot)$ can depend only on what the government observes, as specified in each setting below.

Each firm's outside option is normalized to zero. A firm with type λ that accepts a contract receives expected payoff $\lambda A(\cdot) - \theta$ and participates if the individual rationality (IR) constraint holds: $\lambda A(\cdot) \geq \theta$. The implementation set is therefore $I = \{(\lambda, \mathbf{s}) : \lambda \geq \theta / A(\cdot)\} \subseteq [0,1] \times \mathbb{R}^d$.

The government's expected payoff from contracting with a firm of type λ is $\lambda(\gamma - A(\cdot))$. The government maximizes its total expected payoff subject to the budget constraint B:

$$\max_{A(\cdot)} \iint_I \lambda(\gamma - A(\cdot)) h(\lambda|\mathbf{s}) g(\mathbf{s}) d\lambda d\mathbf{s}$$
$$\text{s.t.} \iint_I \lambda A(\cdot) h(\lambda|\mathbf{s}) g(\mathbf{s}) d\lambda d\mathbf{s} \leq B.$$

Each informational setting below is a special case of this problem, differing only in what the government observes and hence in what $A(\cdot)$ may depend on.

3.1 First-Best Benchmark

To establish the upper bound on the government's payoff, I assume the government observes each firm's true success probability λ and the signals $\mathbf{s}$. Since λ is fully known, the signals $\mathbf{s}$ provide no additional information, and the award $A(\lambda)$ is tailored to λ. The participation constraint requires $\lambda A(\lambda) \geq \theta$, yielding the implementation set $I =$

$\{\lambda : \lambda \geq \frac{\theta}{A(\lambda)}\}$. The government maximizes:

$$\max_{A(\lambda)} \int_{\lambda \geq \frac{\theta}{A(\lambda)}} \lambda(\gamma - A(\lambda)) f(\lambda) d\lambda$$

$$\text{s.t.} \int_{\lambda \geq \frac{\theta}{A(\lambda)}} \lambda A(\lambda) f(\lambda) d\lambda \leq B.$$

Theorem 1 *The first-best optimal award is* $A^*(\lambda) = \frac{\theta}{\lambda}, \forall \lambda \geq \lambda^*$*, where* λ^* *is the cutoff determined by the budget constraint* $\int_{\lambda \geq \lambda^*} \theta f(\lambda) d\lambda = B$*.*

Proof of Theorem 1: The proof proceeds in two steps.

Step 1: Optimal award. For any participating firm, the government's payoff $\lambda(\gamma - A(\lambda))$ is strictly decreasing in $A(\lambda)$. The government therefore minimizes $A(\lambda)$ subject to the IR constraint $\lambda A(\lambda) \geq \theta$, which yields $A^*(\lambda) = \theta/\lambda$. At this payment, the firm's expected payoff is $\lambda \cdot \theta/\lambda - \theta = 0$, so the government extracts all surplus.

Step 2: Optimal implementation set. Under $A^*(\lambda) = \theta/\lambda$, the government's expected payoff from funding a type λ firm is $\lambda\gamma - \theta$, which is positive if and only if $\lambda > \theta/\gamma$. Since this payoff is strictly increasing in λ, the government ranks firms by λ and funds from the top until the budget is exhausted. Each funded firm costs θ in expectation, so the budget constraint $\int_{\lambda \geq \lambda^*} \theta f(\lambda)\, d\lambda = B$ pins down the cutoff λ^*. ■

Under this first-best optimal award, the government's expected payoff is:

$$\bar{\Pi} = \int_{\lambda \geq \lambda^*} (\lambda\gamma - \theta) f(\lambda) d\lambda,$$

which serves as the upper bound for evaluating other mechanisms.

3.2 Pooling Benchmark

In the pooling benchmark, the government observes neither each firm's success probability λ nor the signals $\mathbf{s}$, and relies only on the prior $f(\lambda)$. Since the government has no firm-specific information, it can only offer a uniform award A to all firms. Each firm then self-selects: a firm with type λ participates if and only if $\lambda A \geq \theta$, yielding the implementation set $I = \{\lambda : \lambda \geq \theta/A\}$.

The government maximizes its expected payoff:

$$\max_{A} \int_{\lambda \geq \frac{\theta}{A}} \lambda(\gamma - A) f(\lambda)\, d\lambda$$

$$\text{s.t. } A \int_{\lambda \geq \frac{\theta}{A}} \lambda f(\lambda)\, d\lambda \leq B.$$

The optimal uniform award A^* is determined by $A^* \int_{\lambda \geq \theta/A^*} \lambda f(\lambda)\, d\lambda = B$.

Under this uniform award, the government's expected payoff is:

$$\underline{\Pi} = \int_{\lambda \geq \frac{\theta}{A^*}} \lambda(\gamma - A^*) f(\lambda)\, d\lambda,$$

which serves as the lower bound for evaluating other mechanisms.

3.3 Observable-Firm-Characteristics (OFC) Mechanism

In the OFC mechanism, the government observes signals $\mathbf{s}$ but not λ. Because $\mathbf{s}$ is publicly observable, the government can condition contracts directly on $\mathbf{s}$ without relying on firm reports, eliminating incentive compatibility concerns regarding $\mathbf{s}$. Upon observing signals $\mathbf{s}$, the government forms posterior belief $h(\lambda|\mathbf{s})$ and offers a contract with payment $A(\mathbf{s})$.[3] A firm with type λ participates if $\lambda A(\mathbf{s}) \geq \theta$, yielding the implementation set $I = \{(\lambda, \mathbf{s}) : \lambda \geq \theta/A(\mathbf{s})\}$.

The government maximizes:

$$\max_{A(\mathbf{s})} \iint_{\lambda \geq \frac{\theta}{A(\mathbf{s})}} \lambda(\gamma - A(\mathbf{s}))h(\lambda|\mathbf{s})g(\mathbf{s})d\lambda d\mathbf{s}$$

$$\text{s.t. } \int A(\mathbf{s})g(\mathbf{s}) \int_{\lambda \geq \frac{\theta}{A(\mathbf{s})}} \lambda h(\lambda|\mathbf{s})d\lambda d\mathbf{s} \leq B.$$

3.3.1 Optimal OFC Award

The Lagrangian is:

$$\mathcal{L} = \iint_{\lambda \geq \frac{\theta}{A(\mathbf{s})}} \lambda(\gamma - A(\mathbf{s}))\, h(\lambda|\mathbf{s})g(\mathbf{s})\, d\lambda\, d\mathbf{s} + \mu_{OFC} \left[B - \int A(\mathbf{s})g(\mathbf{s}) \int_{\lambda \geq \frac{\theta}{A(\mathbf{s})}} \lambda\, h(\lambda|\mathbf{s})\, d\lambda\, d\mathbf{s} \right],$$

[3]Since λ is the firm's private information, one might consider awards of the form $A(\hat{\lambda}, \mathbf{s})$ that also depend on the firm's self-reported type $\hat{\lambda}$. However, the incentive compatibility constraint requires that no firm gains by misreporting: $\lambda A(\lambda, \mathbf{s}) - \theta \geq \lambda A(\hat{\lambda}, \mathbf{s}) - \theta$ for all $\lambda, \hat{\lambda}$, which implies $A(\lambda, \mathbf{s}) \geq A(\hat{\lambda}, \mathbf{s})$ for all $\lambda, \hat{\lambda}$. This holds for all types only if A is independent of the reported $\hat{\lambda}$, so the optimal award conditions on $\mathbf{s}$ alone without loss of generality.

where $\mu_{OFC} \geq 0$ is the Lagrange multiplier. The optimal award $A^*(\mathbf{s})$ satisfies the first-order condition, together with complementary slackness:

$$\begin{aligned} -(1+\mu_{OFC}) \int_{\lambda \geq \frac{\theta}{A^*(\mathbf{s})}} \lambda \cdot h(\lambda|\mathbf{s})\, d\lambda + \frac{\theta^2}{A^*(\mathbf{s})^3} \cdot (\gamma - (1+\mu_{OFC})A^*(\mathbf{s})) \cdot h\left(\frac{\theta}{A^*(\mathbf{s})} \,\middle|\, \mathbf{s}\right) &= 0 \\ \mu_{OFC}\left[B - \int A^*(\mathbf{s}) g(\mathbf{s}) \int_{\lambda \geq \frac{\theta}{A^*(\mathbf{s})}} \lambda h(\lambda|\mathbf{s}) d\lambda d\mathbf{s}\right] &= 0. \end{aligned}$$

Equivalently, defining $\lambda^*(\mathbf{s}) = \theta / A^*(\mathbf{s})$, which represents the marginal participating type for each signal realization, the shadow price satisfies:

$$\frac{1}{\mu_{OFC}+1} = \frac{\theta}{\gamma}\left(\frac{\int_{\lambda^*(\mathbf{s})}^1 \lambda h(\lambda|\mathbf{s}) d\lambda}{\lambda^{*3}(\mathbf{s}) h(\lambda^*(\mathbf{s})|\mathbf{s})} + \frac{1}{\lambda^*(\mathbf{s})}\right), \forall \mathbf{s}.$$

This condition, together with the budget constraint, pins down the optimal OFC award. For a given μ_{OFC}, the shadow price condition determines the cutoff $\lambda^*(\mathbf{s})$ for each signal realization $\mathbf{s}$. The multiplier μ_{OFC} is then chosen so that the resulting expected expenditure exhausts the budget B.

Therefore, under the optimal OFC award, the government's expected payoff is:

$$\Pi_{OFC} = \iint_{\lambda \geq \frac{\theta}{A^*(\mathbf{s})}} \lambda(\gamma - A^*(\mathbf{s}))\, h(\lambda|\mathbf{s})\, g(\mathbf{s})\, d\lambda\, d\mathbf{s}.$$

3.4 Comparison of Shadow Prices

The optimality conditions from Sections 3.1–3.3 can each be expressed as a shadow price condition, revealing the cost of information frictions across mechanisms. To establish this, I first derive the shadow price of the budget constraint for the first-best and pooling mechanisms, then compare all three.

First-best. Under $A^*(\lambda) = \theta/\lambda$, the government's problem reduces to choosing the cutoff λ^*:

$$\max_{\lambda^*} \int_{\lambda^*}^1 (\lambda\gamma - \theta) f(\lambda)\, d\lambda \quad \text{s.t.} \quad \int_{\lambda^*}^1 \theta f(\lambda)\, d\lambda \leq B.$$

Letting $\mu_{FB} \geq 0$ denote the Lagrange multiplier and differentiating with respect to λ^*:

$$\frac{1}{1+\mu_{FB}} = \frac{\theta}{\gamma} \cdot \frac{1}{\lambda^*}.$$

Pooling. Under a uniform award A^* with cutoff $\lambda_P^* = \theta/A^*$, let $\mu_P \geq 0$ denote the Lagrange multiplier. Similarly,

$$\frac{1}{1+\mu_P} = \frac{\theta}{\gamma}\left(\frac{\int_{\lambda_P^*}^1 \lambda\, f(\lambda)\, d\lambda}{\lambda_P^{*3}\, f(\lambda_P^*)} + \frac{1}{\lambda_P^*}\right).$$

OFC. From Section 3.3,

$$\frac{1}{1+\mu_{OFC}} = \frac{\theta}{\gamma}\left(\frac{\int_{\lambda^*(\mathbf{s})}^1 \lambda\, h(\lambda|\mathbf{s})\, d\lambda}{\lambda^{*3}(\mathbf{s})\, h(\lambda^*(\mathbf{s})|\mathbf{s})} + \frac{1}{\lambda^*(\mathbf{s})}\right), \quad \forall \mathbf{s}.$$

All three conditions share the form:

$$\frac{1}{1+\mu} = \frac{\theta}{\gamma}\left(R + \frac{1}{\lambda^*}\right),$$

where the common term $\frac{\theta}{\gamma} \cdot \frac{1}{\lambda^*}$ is the ratio of the marginal firm's research cost θ to its ex-ante social value $\lambda^*\gamma$. $R \geq 0$ is the information-rent term—the additional payment to inframarginal firms that results from raising the award to attract the marginal firm. In the first-best, $R = 0$; in the pooling benchmark, $R = \frac{\int_{\lambda_P^*}^1 \lambda\, f(\lambda)\, d\lambda}{\lambda_P^{*3}\, f(\lambda_P^*)}$; in the OFC mechanism, $R(\mathbf{s}) = \frac{\int_{\lambda^*(\mathbf{s})}^1 \lambda\, h(\lambda|\mathbf{s})\, d\lambda}{\lambda^{*3}(\mathbf{s})\, h(\lambda^*(\mathbf{s})|\mathbf{s})}$.

In the first-best, the government observes each firm's type and sets $A^*(\lambda) = \theta/\lambda$ individually. Adjusting the award for the marginal firm has no effect on any other firm's award, so $R = 0$. In the pooling benchmark, the government offers a uniform award to all firms. Raising A^* to attract the marginal firm simultaneously increases the award to every inframarginal firm with $\lambda > \lambda_P^*$, yielding the largest R. In the OFC mechanism, the signals $\mathbf{s}$ partition firms into groups with similar research ability. Raising $A^*(\mathbf{s})$ affects only the inframarginal firms within that signal group, so $R(\mathbf{s})$ lies between the two extremes: if $\mathbf{s}$ is uninformative about λ, i.e., $h(\lambda|\mathbf{s}) = f(\lambda)$ for all $\mathbf{s}$, $R(\mathbf{s})$ reduces to the pooling expression; if $\mathbf{s}$ perfectly reveals λ, firms within each signal group share the same research ability, $R(\mathbf{s}) \to 0$, and the OFC mechanism achieves the first-best.

This characterization identifies the precise channel through which information frictions reduce efficiency: information rents consume part of the budget, forcing the government to set a higher cutoff λ^* and crowding out firms that would otherwise be funded.

4 Current Contract

Comparing the current contract with the OFC mechanism requires computing the government's expected payoff under both regimes. I focus on expected rather than

realized payoffs because the OFC mechanism is a counterfactual.[4] A key challenge is that computing the government's expected payoff under the current contract requires knowing what award each candidate would have received had it participated. However, awards are observed only for participating firms. I therefore use machine learning methods to estimate unobserved awards for nonparticipating firms.

Estimation Strategy. Using SBIR contract data from 2018–2022, I estimate the current contract as a mapping from observable firm characteristics to awards. The characteristics include patent intensity, firm size, firm age, ownership structure, demographic classifications (women-owned, HUBZone, and underrepresented status), geographic location, and award year, as described in Section 2. For firms with prior SBIR experience, I additionally include past contract count and success rate. Accordingly, I partition firms into two groups—new entrants without previous contracts and experienced firms with at least one prior contract.

For each group, I estimate the award prediction function using several models, including OLS, LASSO, Ridge, Random Forest, and Neural Network, and select the specification with the lowest out-of-sample mean squared error (MSE). Random Forest achieves the best performance for new entrants (MSE = 0.00097), while LASSO performs best for experienced firms (MSE = 0.00019).

Determinants of Current Contract. For new entrants, the Random Forest model identifies three primary determinants: award year (34.5% of importance), firm age (28.1%), and firm size (13.8%). Geographic location in the South also contributes notably (11.4%). Patent intensity, a widely used signal of research ability, explains only 2.6% of the current award variation. This suggests that patents play a limited role in allocation decisions for new entrants.

For experienced firms, the LASSO model selects 7 of 14 observable firm characteristics as relevant predictors, including award year, geographic location, and demographic indicators such as women-owned and disadvantaged status. Two findings merit particular attention. First, repeated awardees receive smaller awards: each additional prior contract reduces the award by \$1,275, suggesting diminishing support for firms that have already benefited from the program. Second, past performance, measured by past success rates, receives zero weight. The current contract penalizes repeated participation

[4]Expected and realized payoffs are close under the current contract: the realized government payoff (\$182.32 million) falls within the 95% confidence interval of the model's predicted expected payoff (\$181.52 million), with a difference of only 0.4%.

without rewarding demonstrated ability to deliver innovations.

Validation. Table 2 compares the predicted award distribution to observed 2022 awards. The model closely matches key moments: the predicted mean differs from the observed by 0.04%, and quantile estimates fall within 4% of actual values. The root mean square error (RMSE) of 0.022 million dollars represents approximately 2.82% of the mean award, confirming the model's predictive accuracy.

Table 2: Model Validation: Comparison of Simulated and Observed Awards

Statistic	Observed	Simulated	Difference (%)
Mean	0.767	0.766	−0.04
Median	0.750	0.765	2.04
25th Percentile	0.750	0.760	1.40
75th Percentile	0.800	0.767	−4.05
Min	0.732	0.709	−3.06
Max	0.800	0.799	−0.13

Notes: This table compares observed 2022 awards with model predictions. All values are in million dollars.

Implications. The current contract depends primarily on calendar year, firm age, and firm size, while innovation-related characteristics such as patent intensity and past success rates play minimal roles. In contrast, the OFC mechanism in Section 3 optimally uses observable firm characteristics to design contracts that maximize expected government payoff.

5 Estimation

This section estimates the key components for the analysis: (1) the conditional distribution of success probability $h(\lambda|\mathbf{s})$ given observable firm characteristics $\mathbf{s}$; (2) the unconditional success probability distribution $f(\lambda)$; (3) the research cost θ; (4) the social value γ; and (5) the budget B.

In the 2022 NASA SBIR program, 279 candidate firms participated in a preliminary phase in 2021 and submitted research proposals. Among these, 116 firms were selected and received contracts; the remaining 163 applied but were not selected. Taking 2022

as the base year, I simulate alternative mechanisms offered to the same set of candidate firms and evaluate the resulting expected government payoffs.

5.1 Distribution of Success Probability

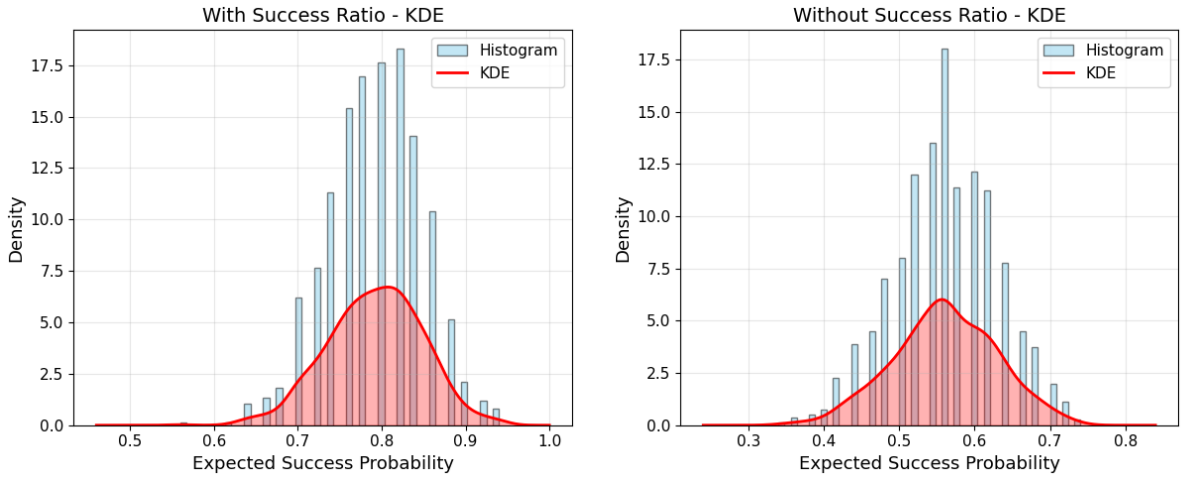
I estimate and calibrate the conditional posterior distribution $h(\lambda|\mathbf{s})$ using a nonparametric mapping from observable firm characteristics to success probability, and derive the unconditional distribution $f(\lambda)$ by integration.

The firm characteristic vector $\mathbf{s}_i$ includes 14 dimensions, as detailed in Section 2: patent intensity, past SBIR contract count and performance (pre-success rate), firm size, firm age, ownership structure (the branch status), and demographic classifications (women-owned, HUBZone-owned, and disadvantaged firms), along with location (West, South, Northeast, and Midwest) and timing controls. A notable feature of the data is that 42% of firms are new entrants without prior SBIR experience. For these firms, I use a 12-dimensional vector excluding past SBIR contract count and success rate. This partitioning accommodates firms with varying levels of available information.

To estimate the conditional distribution of success probability given firm characteristic $h(\lambda|\mathbf{s})$, I use a nonparametric approach that combines kernel weighting with bootstrapping. For a target firm with characteristic $\mathbf{s}$, I use weighted bootstrap resampling, where the weights are derived from kernel similarity measures between the target characteristics and observed firms. For each bootstrap sample, I compute the success rate. This process yields a collection of success rates from the bootstrap samples. Finally, I apply kernel density estimation to these bootstrap success rates to obtain a smooth estimate of $h(\lambda|\mathbf{s})$, the conditional distribution of success probability given observable firm characteristics. This approach allows the data to flexibly determine the relationship between observable firm characteristics and success probability without imposing strong parametric assumptions. The detailed estimation procedure is provided in Appendix A.1.

Figure 2 displays the estimated posterior distributions $h(\lambda|\mathbf{s})$ under example firm characteristics. Figure 2a shows the conditional distribution of success probability for an experienced, 10-year-old Southern firm with 100 employees, a log patent intensity of 0.5, five prior NASA SBIR contracts with 60% success rate, and not women-owned, HUBZone-owned, or disadvantaged. The probability that this firm will deliver an innovation has a mean of 80% and a standard deviation of 6%. Figure 2b represents the conditional distribution of success probability for a 3-year-old Northeastern firm with 50 employees, a log patent intensity of 0.3, without prior NASA SBIR contracts, women-

owned and disadvantaged but not HUBZone-owned. The probability that this firm will deliver an innovation has a mean of 56% and a standard deviation of 7%.



(a) With Past Performance

(b) Without Past Performance

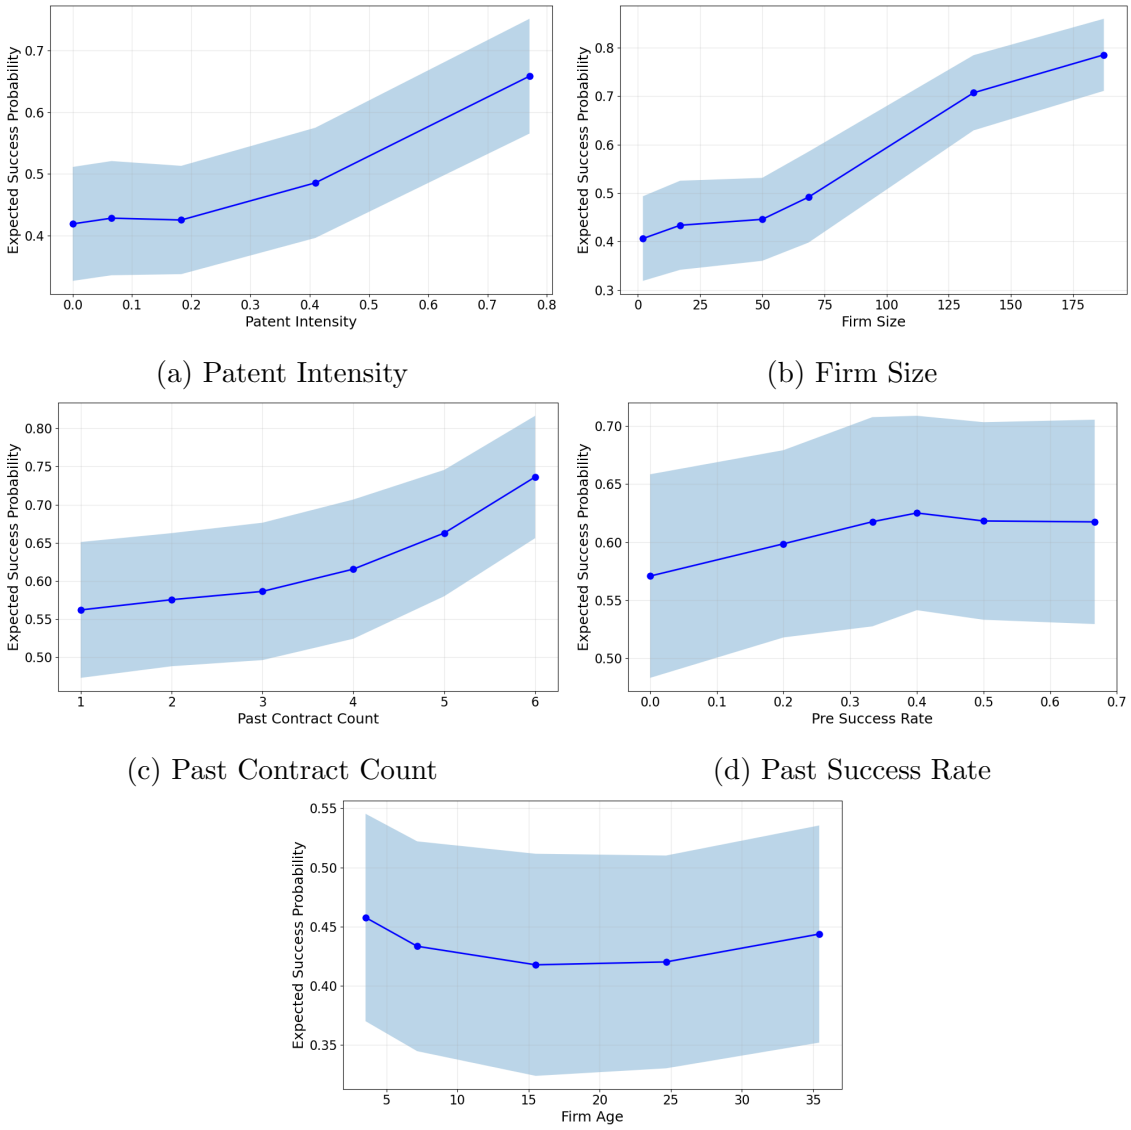
Figure 2: Estimated Posterior Distributions $h(\lambda|\mathbf{s})$

The estimated expected success probability $\mathbb{E}[\lambda|\mathbf{s}]$ increases monotonically with some favorable characteristics, as reported in Figure 3. Figure 3a illustrates that firms with higher patent intensity exhibit greater success probabilities, consistent with the intuition that more patents reflect stronger research ability. Figure 3b shows a steady increase in success probability as firm size grows. This may result from larger firms having more financial and operational resources to manage complex SBIR projects. Figure 3c indicates that past contract experience has a limited but positive impact, suggesting that familiarity with SBIR provides a slight advantage. As represented in Figure 3d, firms that performed well in past SBIR projects are more likely to succeed in future ones. In contrast, firm age exhibits no clear relationship with success probability, as shown in Figure 3e. Notably, while patent intensity and previous SBIR performance are strong predictors of success, Section 4 shows that the current contract places minimal weight on them.

To estimate the unconditional distribution $f(\lambda)$, I integrate the conditional density $h(\lambda|\mathbf{s})$ over the empirical distribution of observable firm characteristics $g(\mathbf{s})$ in the candidate pool:

$$f(\lambda) = \int h(\lambda|\mathbf{s})\, g(\mathbf{s})\, d\mathbf{s}.$$

Details are provided in Appendix A.2. Figure 4 presents the resulting unconditional distribution.



(a) Patent Intensity

(b) Firm Size

(c) Past Contract Count

(d) Past Success Rate

(e) Firm Age

Figure 3: Expected Success Probability by Observable Firm Characteristics

5.2 Parameter Calibration

This subsection calibrates three parameters—the budget B, the research cost θ, and the social value of each successful innovation γ—using observed data from the 2022 NASA SBIR program. Two fiscal quantities are central to this calibration: the total obligation of \$88.86 million, representing the government's committed funding to all participating firms, and the realized expenditure of \$61.29 million, representing actual payments to successful firms.

Budget. I use the realized expenditure of \$61.29 million as the budget constraint

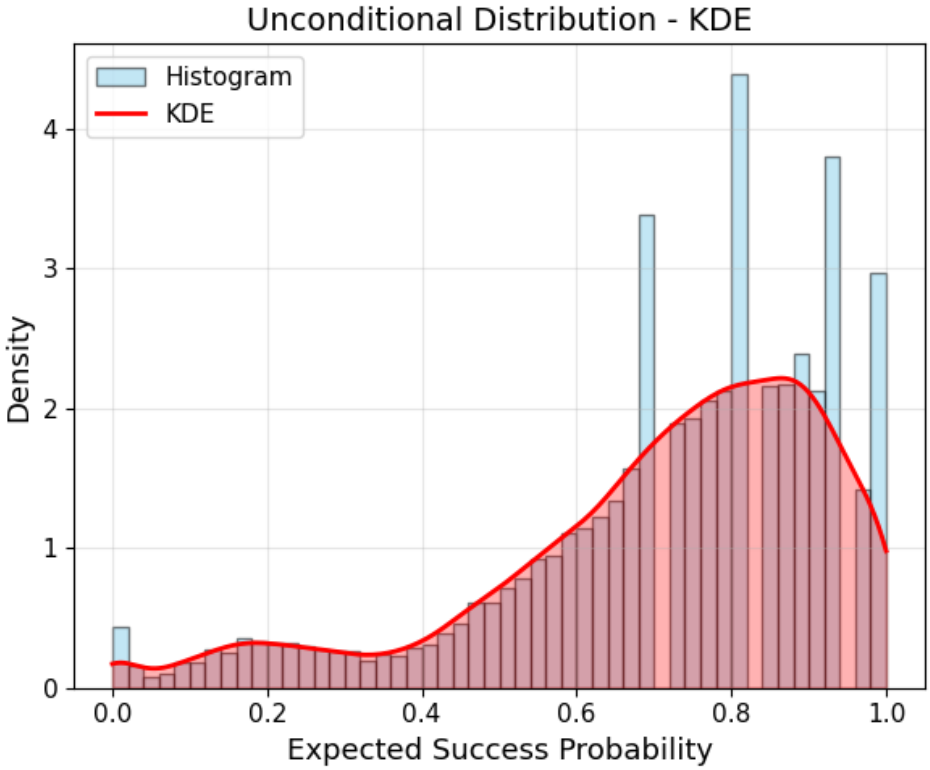


Figure 4: Unconditional Distribution of Success Probability

B in the counterfactual simulations. This reflects the government's actual fiscal outlay under the current contract.

Research Cost. I calibrate the research cost θ by matching the model's predicted budget to the observed realized expenditure of \$61.29 million, based on the simulation of current contract in Section 4. For each candidate firm j with predicted potential award A_j, payment occurs if the firm participates and successfully delivers the innovation. The expected payment to each candidate is $A_j \int_{\lambda_j \geq \theta/A_j} \lambda_j h(\lambda_j|\mathbf{s}) d\lambda_j$, where $\lambda_j \geq \theta/A_j$ is the participation constraint and the integral is the probability of receiving the payment. This calibration yields $\hat{\theta} = 0.65$ million dollars. Details are provided in Appendix A.3.

Social Value. I calibrate the social value per successful innovation γ using obligation data and NASA's economic impact assessment. According to NASA's 2017 SBIR report (NASA Ames Research Center, 2017), \$172.9 million in obligations generated approximately \$474 million in additional GDP. Applying this ratio to the 2022 obligation of \$88.86 million yields an estimated GDP contribution of \$252.13 million. With 80 completed contracts, each contract generated an average of \$3.05 million in additional GDP, implying $\gamma = 3.05$.

6 Results

This section evaluates the performance of alternative mechanisms through counterfactual simulations. Using the theoretical framework from Section 3 and the calibrated parameters from Section 5, I compare expected government payoffs across alternative mechanisms in Section 6.1. I then decompose the sources of improvement under the OFC mechanism in Section 6.2, and discuss policy implications in Section 6.3.

6.1 Government Payoff Improvement

To provide a comprehensive evaluation framework, the analysis compares four mechanisms, ranked by expected government payoff from low to high: (1) a pooling benchmark where the government, without any screening capability, allocates a uniform award to all firms, representing the lower bound of potential efficiency; (2) the current contract; (3) the OFC mechanism that optimally incorporates observable firm characteristics; and (4) the first-best benchmark with full observability of firms' research ability, representing the upper bound of potential efficiency. Each mechanism operates under the same budget constraint as the actual 2022 SBIR program ($B = 61.29$ million).

Table 3 reports the distribution of simulated awards under the current contract and the OFC mechanism, both weighted by firms' participation probability. The comparison reveals distinct allocation patterns between the two mechanisms. The OFC mechanism features smaller awards at the lower-to-middle range, with lower values at the minimum, 25th percentile, median, and mean. However, the pattern reverses in the upper portion of the distribution: the OFC mechanism shows higher values at the 75th percentile and maximum. These patterns indicate that the current contract produces a highly concentrated distribution with a narrow range of \$0.10 million, while the OFC mechanism generates greater dispersion with a range of \$0.20 million, reflecting its optimal use of observable firm characteristics to tailor awards to heterogeneous firms.

Table 3: Award Distribution: Current Contract vs. OFC Mechanism (Million Dollars)

Contract Type	Min	25%	Median	Mean	75%	Max	SD
Current Contract	0.70	0.76	0.76	0.77	0.78	0.80	0.01
OFC Mechanism	0.65	0.72	0.76	0.75	0.79	0.85	0.05

Table 4 presents the expected government payoffs under each mechanism. To measure the value of information, I use the difference between the first-best and pooling

payoffs as the maximum potential improvement, and evaluate each mechanism by the fraction of this gap it captures. The OFC mechanism achieves a government payoff of \$208.22 million, capturing 61.25% of the potential gains. In contrast, the current contract achieves only 4.79% of the potential gains with a payoff of \$181.52 million—barely \$2 million above the pooling benchmark of \$179.26 million. The results reveal substantial gains from optimally incorporating public signals. Table 5 shows that the OFC mechanism generates larger social surplus compared to the current contract (\$211.39 million to \$186.75 million) and lower firms' payoff (\$3.17 million to \$5.23 million), reflecting improved screening efficiency.[5] The robustness of these results is examined in Appendix B, including sensitivity to the estimation of $h(\lambda|\mathbf{s})$ and alternative calibrations of the social value parameter γ.

Table 4: Government Payoffs Under Alternative Mechanisms (Million Dollars)

	Pooling (Lower Bound)	Current Contract	OFC Mechanism	First-Best (Upper Bound)
Government Payoff	179.26	181.52	208.22	226.54
	(0.04)	(0.34)	(1.16)	(0.29)
Improvement over Pooling	0	2.26	28.96	47.28
Relative Improvement(%)	0	4.79	61.25	100
	–	(0.68)	(2.51)	–

Notes: All payoffs are expected values. Bootstrap standard errors in parentheses (100 replications). Standard errors account for estimation uncertainty in both the conditional success probability distribution $h(\lambda|\mathbf{s})$ and the research cost parameter θ.

The small improvement of the current contract over the pooling benchmark reflects its underutilization of available information. As shown in Section 4, the current contract places little weight on characteristics such as patent intensity and past success rate. Without conditioning on these signals, the current contract cannot effectively distinguish firms with different research abilities. The resulting award distribution is highly concentrated, with a standard deviation of only \$0.01 million, closely resembling the uniform allocation under pooling.

The OFC mechanism, by contrast, optimally uses observable firm characteristics to infer research ability and design contracts accordingly. This generates a more dispersed

[5]Social surplus generated from each contracted firm is defined as $\lambda\gamma - \theta$ and the firm payoff is defined as $\lambda A(\cdot) - \theta$.

Table 5: Innovation Outcomes Under Alternative Mechanisms (Million Dollars)

	Pooling (Lower Bound)	Current Contract	OFC Mechanism	First-Best (Upper Bound)
Social Surplus	184.37	186.75	211.39	226.54
	(0.04)	(0.34)	(1.17)	(0.29)
Firm Payoff	5.11	5.23	3.17	0.00
	(0.01)	(0.01)	(0.01)	–

award distribution that reflects underlying heterogeneity across firms. The following subsection decomposes this improvement to identify the channels through which the OFC mechanism outperforms the current contract.

6.2 Improvement Decomposition

To understand the sources of the improvement in Table 4, I decompose it through firm-level analysis and examine the underlying reallocation strategy.

I divide firms into two groups: Group 1 receives larger awards under the OFC mechanism than the current contract, while Group 2 receives smaller awards. Figure 5 reflects the distribution of award change.

Table 6 presents summary statistics for each group. 211 firms (76%) fall into Group 1 and 68 firms (24%) into Group 2. Comparing across groups, Group 1 and Group 2 exhibit distinct features. Group 1 firms have lower research ability than Group 2 firms. Under the current contract, Group 1 firms exhibit, on average, lower participation rates and generate lower government payoffs.

Comparing across mechanisms, the OFC mechanism treats these two groups differently. For Group 1, participation rates increase substantially from 0.12 to 0.21 (a 79.0% increase), and government payoffs rise from 0.23 to 0.40. For Group 2, participation declines modestly from 0.91 to 0.83 (a 9.1% decrease), with government payoffs falling from 1.94 to 1.81. This asymmetric response reflects heterogeneous participation elasticities: Group 1 firms respond more sensitively to award changes, while Group 2 firms largely continue to participate despite reduced awards. The OFC mechanism exploits this heterogeneity, expanding participation among elastic firms while retaining most inelastic participants.

Budget Reallocation. The asymmetric response above reveals the mechanism's

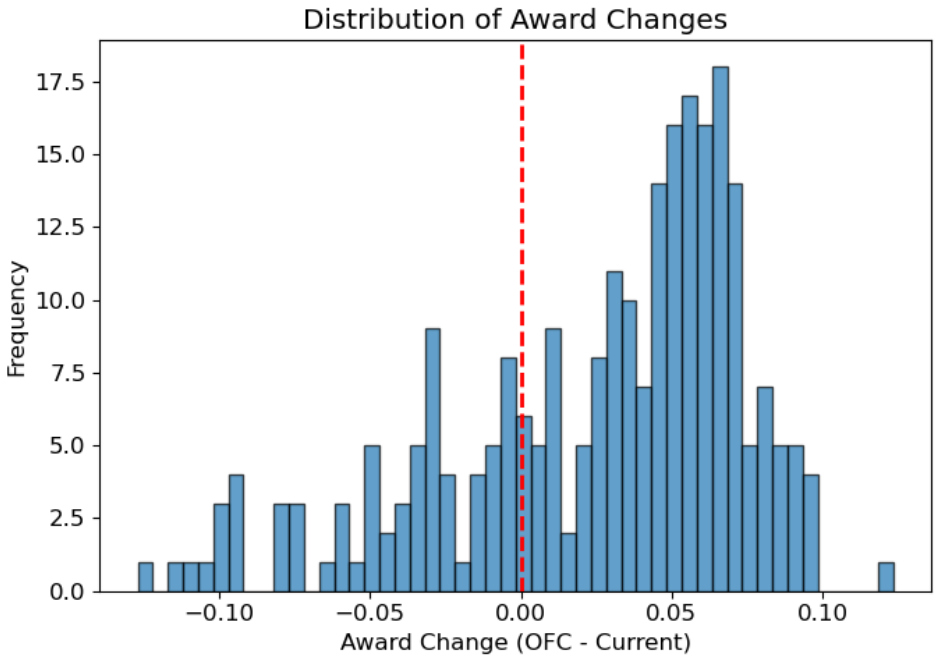


Figure 5: Distribution of Award Change

Table 6: Summary Statistics by Group

	N	Research Ability	Participation Rate		Expected payment		Gov. Payoff	
			Current	OFC	Current	OFC	Current	OFC
Group 1	211	0.70	0.12	0.21	0.08	0.14	0.23	0.40
Group 2	68	0.92	0.91	0.83	0.66	0.56	1.94	1.81

Notes: All values are group means. Expected payment and expected government payoff are in million dollars.

core budget reallocation strategy. For high-ability firms in Group 2, the current contract provides excessive awards beyond their participation requirements. The OFC mechanism reduces these awards while government payoffs decline only slightly (from \$1.94 to \$1.81 million per firm), as these inelastic firms largely continue to participate. This saved budget is reallocated to low-ability firms in Group 1, where increased awards attract substantially more participants and raise government payoffs from \$0.23 to \$0.40 million per firm.

Effect Decomposition. The government payoff from each firm is $\int_{\lambda \geq \lambda^*} \lambda(\gamma - A(\cdot))h(\lambda)d\lambda$. For firms in Group 1, the larger award in the OFC mechanism shifts the participation threshold $\lambda^* = \theta/A(\cdot)$ to the left ($\lambda^*_{OFC} < \lambda^*_{current}$), increasing partici-

pation probability. This shift generates two effects. First, the "entry effect": the OFC mechanism gains from the region between the two thresholds ($\lambda^*_{OFC} \leq \lambda < \lambda^*_{current}$) because of larger participation probability. Second, the "allocation effect": for the overlapping participation region ($\lambda \geq \lambda^*_{current}$), the government pays more per successful innovation. For Group 2, the logic reverses: smaller awards shift thresholds rightward ($\lambda^*_{OFC} > \lambda^*_{current}$), generating an "exit effect" and an "allocation gain."

Table 7: Decomposition of Government Payoff Improvement (Million $)

	Group 1	Group 2
Entry/Exit Effect	+35.78	−11.27
Allocation Effect	−0.41	+2.60
Subtotal	+35.37	−8.67

Table 7 presents the decomposition results. The "entry effect" dominates: Group 1 firms generate $35.78 million from expanded participation, more than offsetting the $11.27 million "exit effect" from Group 2. Allocation effects are relatively small (−0.41 and +2.60 million), indicating that improvement primarily stems from changing participant composition rather than optimizing payments to existing participants.

Heterogeneous Effects. Table 8 examines how the OFC mechanism's improvement varies across firm subgroups. Panel A shows that experienced firms benefit more than new entrants (19.2% vs. 12.1%). The current contract ignores patent intensity for all firms and additionally ignores past success rates for experienced firms. The OFC mechanism exploits both sources of information, and the additional information available for experienced firms generates larger gains.

Panels B–D decompose effects by observable firm characteristics. For both patent intensity and firm size, firms in the high group experience the smallest improvements (−1.8% and 6.5%, respectively), while those with lower values benefit more. This heterogeneity reflects the OFC mechanism's reallocation of budget: it offers higher awards to low-ability firms to induce participation, while reducing awards to high-ability firms who would participate regardless. These patterns align with Figure 3, where patent intensity and firm size are positively associated with expected research ability. In contrast, firm age exhibits no clear pattern (Panel D), consistent with its weak relationship with research ability documented in Figure 3.

Panels E and F report results for experienced firms by past contract count and past

Table 8: Heterogeneous Effects of OFC Mechanism

		Award		Participation		Gov Payoff		Improvement	
Group	N	Current	OFC	Current	OFC	Current	OFC	$M	%
Panel A: By SBIR Experience									
All Firms	279	0.77	0.79	0.31	0.36	181.52	208.22	26.69	14.71
New Entrants	146	0.77	0.79	0.37	0.42	113.68	127.38	13.70	12.05
Experienced	133	0.77	0.79	0.24	0.29	67.84	80.84	13.00	19.16
Panel B: By Patent Intensity									
Low	184	0.77	0.79	0.31	0.37	118.32	138.72	20.40	17.24
Medium	48	0.77	0.80	0.22	0.29	21.90	28.94	7.04	32.15
High	47	0.77	0.77	0.41	0.40	41.30	40.56	−0.74	−1.80
Panel C: By Firm Size									
Low	98	0.77	0.80	0.30	0.36	60.46	70.57	10.10	16.71
Medium	88	0.77	0.80	0.25	0.33	46.27	58.03	11.76	25.41
High	93	0.76	0.77	0.37	0.39	74.79	79.62	4.84	6.47
Panel D: By Firm Age									
Low	93	0.77	0.79	0.37	0.41	71.58	79.17	7.59	10.60
Medium	95	0.76	0.79	0.31	0.36	60.45	70.92	10.47	17.32
High	91	0.77	0.79	0.26	0.30	49.49	58.13	8.64	17.45
Panel E: By Past Contract Count (Exp.)									
Low	54	0.77	0.81	0.15	0.21	16.96	23.46	6.51	38.36
Medium	56	0.77	0.79	0.22	0.29	26.32	34.00	7.68	29.16
High	23	0.76	0.76	0.50	0.47	24.56	23.38	−1.18	−4.82
Panel F: By Past Success Rate (Exp.)									
Low	56	0.77	0.81	0.09	0.19	10.78	20.51	9.73	90.28
Medium	39	0.76	0.77	0.35	0.38	29.80	32.16	2.36	7.92
High	38	0.77	0.79	0.34	0.35	27.26	28.17	0.91	3.33

Notes: Award and Gov Payoff in million dollars. Low/Medium/High based on terciles. Panels B–D report all firms; Panels E and F report experienced firms only.

success rate. The most notable finding appears in Panel F: firms with low past success rates see a 90.3% improvement, compared to only 3.3% for those with high success rates. This disparity arises because the current contract assigns zero weight to previous performance, treating high and low performers identically. The OFC mechanism corrects this misallocation by conditioning awards on firms' previous performance. Similarly, Panel E shows that firms with fewer prior contracts benefit more (38.4%) than those with more prior contracts (-4.8%).

Additional Congressional Objectives. Beyond stimulating technological innovation, Congress established distributional objectives for the SBIR program, namely fostering participation by small businesses and underrepresented firms,[6] as well as encouraging new entrants.[7]

The budget reallocation and heterogeneous effects documented above suggest that the OFC mechanism is broadly consistent with these goals. First, the OFC mechanism favors smaller firms. Panel C of Table 8 shows that small and medium firms experience bigger participation gains than large firms, shifting the participants toward smaller businesses. Second, the OFC mechanism expands access for new entrants: Panel A shows that participation increases from 0.37 to 0.42 among new entrants. Third, because underrepresented groups—including women-owned firms, HUBZone firms, and socially disadvantaged firms—tend to be lower-ability firms with higher participation elasticities, the OFC mechanism naturally directs more resources toward them. Appendix C reports detailed firm size distributions and participation by firm demographics, confirming that the OFC mechanism does not compromise these additional objectives.

6.3 Policy Implications

Effect of Budget Expansion. Policymakers may consider increasing funding for R&D programs such as SBIR. A natural question arises: if the government expands its budget, would the OFC mechanism continue to outperform the pooling benchmark, or would its advantages diminish? I examine the performance of the OFC mechanism under hypothetical budget expansions, holding the candidate pool fixed, to provide guidance

[6] "The goals of the SBIR program include: Use small business to meet Federal R/R&D needs; Foster and encourage participation by the socially and economically disadvantaged small businesses and those that are 51 percent owned and controlled by women, in technological innovation", see U.S. Small Business Administration (2023).

[7] The INNOVATE Act of 2025 (S.853, 119th Congress) introduced "encouraging new SBIR entrants" as a program objective. See `https://www.congress.gov/bill/119th-congress/senate-bill/853`.

for the potential adoption of the OFC mechanism.

Table 9: Government Payoffs Under Alternative Budget Levels

	Government Payoff ($M)			
Budget ($M)	Pooling	OFC	First-Best	OFC Rel. Imp.
61.29 (baseline)	179.26	208.22	226.54	61.25%
71.29	201.59	230.71	248.75	61.74%
81.29	222.14	256.07	269.90	71.05%
91.29	240.94	277.91	289.92	75.48%

Table 9 reports government payoffs under the pooling benchmark, the OFC mechanism, and the first-best outcome for four budget levels: the baseline budget of $61.29 million and three expansions of $10, $20, and $30 million. The current contract is excluded because its allocation rule is estimated from observed data and cannot be extrapolated to alternative budget levels. I find that all three mechanisms generate higher government payoffs as the budget increases. More importantly, the OFC mechanism's relative improvement increases with the budget. At the baseline budget, the OFC mechanism captures 61.25% of the gap between the pooling benchmark and the first-best outcome. This share rises to 61.74% at $71.29 million, 71.05% at $81.29 million, and 75.48% at $91.29 million. This pattern suggests that the value of optimally incorporating observable firm characteristics into contract design becomes more pronounced when the budget is less constrained.

The intuition is as follows. As the budget grows, the pooling benchmark can only raise its uniform award, attracting additional participants without differentiating among them. In contrast, the OFC mechanism strategically allocates additional resources to firms with high participation elasticity, expanding participation more efficiently than the pooling benchmark. This pattern can be understood through the shadow price framework in Section 3.4. As the budget expands, the shadow price μ falls under all mechanisms. However, under the pooling benchmark, each additional dollar must be spread uniformly, generating a large information rent R: raising the uniform award to attract one additional marginal firm simultaneously increases payments to all inframarginal firms. The OFC mechanism, by conditioning awards on signals $\mathbf{s}$, reduces $R(\mathbf{s})$ within each signal group, so that each dollar of additional budget translates more efficiently into participation gains. Consequently, the gap between the two mechanisms widens as the budget grows.

These results have direct policy implications. If NASA considers expanding SBIR funding, the benefits of adopting the OFC mechanism would be amplified. The OFC mechanism not only improves performance at current funding levels but becomes increasingly valuable as budgets grow.

External Validity. Within the SBIR program, the Department of Defense (DoD) is the largest participating agency with a higher budget than NASA. Given that the OFC mechanism's advantage increases with budget size, the potential efficiency gains at DoD could be even larger than those documented here for NASA. Several institutional and empirical features support this extension. First, the DoD SBIR program operates under the same Congressional mandate as NASA, requiring agencies to allocate at least 3.2% of their extramural R&D budgets to small businesses. The two-phase structure, application process, eligibility requirements, and program goals are uniform across agencies. The two agencies also share technology interests in areas such as aerospace and autonomous systems.

Second, NASA and DoD effectively draw from a similar pool of innovative firms: of the 1,974 firms receiving NASA Phase II contracts since 1984, 57.1% also received DoD Phase II contracts within a five-year window. Moreover, the characteristics identified as predictive—patent intensity, past contract performance, and firm size—reflect general research ability rather than agency-specific factors.

Therefore, the OFC mechanism's approach extends naturally to the DoD SBIR program. More broadly, it may also apply to public R&D procurement in areas such as clean energy, artificial intelligence, and biotechnology. Extending this paper's quantitative analysis to the DoD SBIR program would require detailed payment data, which are not currently available.

7 Conclusion

This paper addresses a fundamental challenge in government research contracting: how to allocate funding efficiently when firms' research ability is not directly observable. I develop an OFC mechanism that optimally designs contracts based on observable signals, and evaluate its performance using data from NASA's SBIR program.

The empirical analysis yields two main findings. First, observable firm characteristics contain substantial information about research ability. Patent intensity and prior SBIR success rates are strong predictors of future performance, yet the current contract assigns minimal weight to these signals. Instead, the current contract depends primarily on

calendar year, firm age, and firm size.

Second, the OFC mechanism generates large efficiency gains. Through counterfactual simulations holding budget constant at the 2022 level ($61.29 million), I find the OFC mechanism achieves a government payoff of $208.22 million, capturing 61.25% of the potential gains between a pooling benchmark and the first-best outcome. The current contract captures only 4.79% of these gains, performing barely better than the pooling benchmark. Decomposing this improvement reveals that the gains stem primarily from the budget reallocation: the OFC mechanism reallocates budget from high-ability firms who would participate regardless, toward low-ability firms with higher participation elasticities.

These findings have direct policy implications for SBIR reform. The current contract underweights innovation-related signals such as patent intensity and past success rates, explaining its poor performance. Shifting weight toward these characteristics could substantially increase government payoff while maintaining additional Congressional objectives. More broadly, the framework extends beyond NASA to other SBIR agencies and public R&D procurement in general, with potentially greater gains at agencies such as DoD, where budgets are substantially larger.

References

Aboody, D. and B. Lev (2000). Information asymmetry, R&D, and insider gains. *The Journal of Finance 55*(6), 2747–2766.

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018). Innovation, reallocation, and growth. *American Economic Review 108*(11), 3450–3491.

Acs, Z. J. and D. B. Audretsch (1988). Innovation in large and small firms: An empirical analysis. *American Economic Review*, 678–690.

Acs, Z. J. and D. B. Audretsch (1990). *Innovation and small firms*. MIT Press.

Aghion, P., J. Van Reenen, and L. Zingales (2013). Innovation and institutional ownership. *American Economic Review 103*(1), 277–304.

Akcigit, U., D. Hanley, and S. Stantcheva (2022). Optimal taxation and R&D policies. *Econometrica 90*(2), 645–684.

Auranen, O. and M. Nieminen (2010). University research funding and publication performance-an international comparison. *Research Policy 39*(6), 822–834.

Bhattacharya, V. (2021). An empirical model of R&D procurement contests: An analysis of the DOD SBIR program. *Econometrica 89*(5), 2189–2224.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying technology spillovers and product market rivalry. *Econometrica 81*(4), 1347–1393.

Bol, T., M. De Vaan, and A. Van De Rijt (2018). The matthew effect in science funding. *Proceedings of the National Academy of Sciences 115*(19), 4887–4890.

David, P. A., B. H. Hall, and A. A. Toole (2000). Is public R&D a complement or substitute for private R&D? A review of the econometric evidence. *Research Policy 29*(4-5), 497–529.

Fang, F. C., A. Bowen, and A. Casadevall (2016). NIH peer review percentile scores are poorly predictive of grant productivity. *Elife 5*, e13323.

Fieldhouse, A. and K. Mertens (2023). The returns to government R&D: Evidence from US appropriations shocks. Working Papers 2305, Federal Reserve Bank of Dallas, revised 21 Nov 2024.

Fischer, E. and A. R. Reuber (2003). Support for rapid-growth firms: A comparison of the views of founders, government policymakers, and private sector resource providers. *Journal of Small Business Management 41*(4), 346–365.

Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *The RAND Journal of Economics 36*(1), 16–38.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). The NBER patent citation data file: Lessons, insights and methodological tools. Technical report. 656-669.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies 66*(1), 169–182.

Howell, S. T. (2017). Financing innovation: Evidence from R&D grants. *American Economic Review 107*(4), 1136–1164.

Jacob, B. A. and L. Lefgren (2011). The impact of research grant funding on scientific productivity. *Journal of Public Economics 95*(9-10), 1168–1177.

Jones, B. F. and L. H. Summers (2020). *A calculation of the social returns to innovation*, Volume 27863. National Bureau of Economic Research.

Laffont, J.-J. and J. Tirole (1986). Using cost observation to regulate firms. *Journal of Political Economy 94*(3, Part 1), 614–641.

Li, D. and L. Agha (2015). Big names or big ideas: Do peer-review panels select the best science proposals? *Science 348*(6233), 434–438.

Li, Q. and J. S. Racine (2007). *Nonparametric econometrics: Theory and practice*. Princeton University Press.

Link, A. N. and J. T. Scott (2010). Government as entrepreneur: Evaluating the commercialization success of SBIR projects. *Research Policy 39*(5), 589–601.

NASA Ames Research Center (2017). 2017 NASA SBIR/STTR economic impact report. `https://www.sbir.gov/impact/impact-reports`.

National Academies of Sciences, Engineering, and Medicine (2016). *SBIR at NASA*. Washington, DC: The National Academies Press.

Pier, E. L., M. Brauer, A. Filut, A. Kaatz, J. Raclaw, M. J. Nathan, C. E. Ford, and M. Carnes (2018). Low agreement among reviewers evaluating the same NIH grant applications. *Proceedings of the National Academy of Sciences 115*(12), 2952–2957.

Stam, E. and K. Wennberg (2009). The roles of R&D in new firm growth. *Small Business Economics 33*(1), 77–89.

Szucs, F. (2024). Discretion and favoritism in public procurement. *Journal of the European Economic Association 22*(1), 117–160.

U.S. Small Business Administration (2023). Small business innovation research (SBIR) and small business technology transfer (STTR) program policy directive. `https://www.sbir.gov/about/policies`.

A Estimation Procedures

A.1 Conditional Distribution Estimation

The estimation of $h(\lambda|\mathbf{s})$ for a target firm with characteristics $\mathbf{s}$ proceeds as follows:

Step 1: Select the appropriate data group. I first identify firms as experienced or new entrants based on whether the firm has previous SBIR records. Therefore, I select the relevant data group: Group 1 for experienced firms with historical SBIR count and performance information included in the observable firm characteristics (14-dimensional), or Group 2 for new entrants with only 12-dimensional observable firm characteristics. The selected group provides the data $\{(\mathbf{s}_j, y_j)\}_{j=1}^{N_g}$, where $\mathbf{s}_j$ is the characteristic vector for firm j (14-dimensional for Group 1, 12-dimensional for Group 2), $y_j \in \{0, 1\}$ is a binary contract success indicator (1 for completion, 0 for failure), and N_g is the number of firms in the group.

Step 2: Compute similarity weights. For the selected group, I compute similarity weights between the target firm's characteristics $\mathbf{s}$ and those of each firm $\mathbf{s}_j$ in the group. The weights are given by: $w_j(\mathbf{s}) = K\left(\frac{\mathbf{s}_j - \mathbf{s}}{\mathbf{h}}\right) \Big/ \sum_{k=1}^{N_g} K\left(\frac{\mathbf{s}_k - \mathbf{s}}{\mathbf{h}}\right)$, where $K(\cdot)$ is a Gaussian kernel, and $\mathbf{h}$ is a vector of bandwidths tailored to each dimension of $\mathbf{s}$.

Step 3: Perform bootstrap with replacement. Using the similarity weights $\{w_j(\mathbf{s})\}$ as sampling probabilities, I draw $B = 1000$ bootstrap samples with replacement, each of size $n = 50$, from the selected group's data. Firms with characteristics more similar to the target have higher weights and are therefore more likely to be drawn. For each bootstrap sample b, I compute the average success rate across the sampled firms as $\hat{\lambda}_b = n^{-1} \sum_{j \in b} y_j$. This process yields a collection of success rates $\{\hat{\lambda}_b\}_{b=1}^{1000}$ for the given target firm characteristics.

Step 4: Estimate the conditional distribution. Following Li and Racine (2007), I apply kernel density estimation (KDE) to the collection of bootstrap success rates to obtain a smooth estimate of $h(\lambda|\mathbf{s})$, the estimated conditional posterior distribution of the success probability given the firm's characteristics $\mathbf{s}$.

A.2 Unconditional Distribution Estimation

The estimation of the unconditional distribution $f(\lambda)$ is conceptually based on integration over the empirical distribution of observable firm characteristics. For efficient implementation, this is achieved through a two-stage Monte Carlo procedure.

First, I stratify the candidate pool based on data availability: firms with historical success rates and those without. For each candidate firm i in the pool, I estimate the

conditional distribution $h(\lambda|\mathbf{s}_i)$ and draw 100 samples from this distribution.

I then perform stratified sampling that preserves the natural proportions of firms with and without historical data, using Monte Carlo $M = 3000$ draws. For each of the Monte Carlo draws, I: (1) randomly select a candidate firm from the pool, maintaining the empirical distribution of observable firm characteristics, and (2) randomly draw one λ value from the pre-computed conditional distribution samples for that firm. This generates a collection $\{\lambda^{(m)}\}_{m=1}^{3000}$ that reflects both the heterogeneity in success probability across firms in the candidate pool and the estimation uncertainty in each firm's conditional distribution $h(\lambda|\mathbf{s})$.

Finally, I apply kernel density estimation (KDE) with a Gaussian kernel to this collection to obtain the smooth unconditional distribution $f(\lambda)$.

A.3 Research Cost Estimation

I calibrate the research cost θ by matching the model's predicted expected expenditure to the observed realized expenditure of \$61.29 million in the 2022 NASA SBIR program.

Specifically, for a given θ, each firm's participation decision depends on whether its success probability λ_j exceeds the threshold θ/A_j, where A_j is the predicted award from Section 4. The expected payment to firm j is:

$$E[\text{Payment}_j] = A_j \int_{\lambda_j \geq \theta/A_j} \lambda_j \, h(\lambda_j|\mathbf{s}_j) \, d\lambda_j,$$

where the participation constraint $\lambda_j \geq \theta/A_j$ determines whether the firm accepts the contract, and the integral represents the expected success probability conditional on participation.

Since firms' success probabilities λ_j are unobservable, I employ kernel density estimation to approximate the conditional distribution $h(\lambda_j|\mathbf{s}_j)$ based on observed outcomes among participants. For each candidate j, I draw $B = 500$ samples $\{\lambda_j^{(b)}\}_{b=1}^{B}$ from this distribution. Given research cost θ, firm j's expected payment is approximated by:

$$\widehat{E[\text{Payment}_j]}(\theta) = A_j \cdot \frac{1}{B} \sum_{b=1}^{B} \lambda_j^{(b)} \cdot \mathbf{1}[\lambda_j^{(b)} \geq \theta/A_j].$$

The predicted total expenditure under parameter θ is:

$$\hat{B}(\theta) = \sum_{j=1}^{279} \widehat{E[\text{Payment}_j]}(\theta).$$

I identify $\hat{\theta}$ by matching the predicted and observed expenditure. This calibration yields $\hat{\theta} = 0.65$ million dollars.

B Robustness Checks

B.1 Robustness to the Estimation of $h(\lambda|\mathbf{s})$

The conditional distribution $h(\lambda|\mathbf{s})$ is estimated using kernel-weighted bootstrap sampling, which involves several methodological choices. I examine sensitivity to: (1) the bandwidth selection method for kernel density estimation (Scott's rule versus cross-validation), and (2) the bootstrap parameters, including sample size ($n \in \{30, 50, 75, 100\}$) and the number of replications ($B \in \{500, 1000, 2000\}$).

Table 10: Robustness of Results to Estimation Methods

			Gov. Payoff ($M)				Rel. Impr. (%)	
Specification	$\mathbb{E}[\lambda]$	θ	Pooling	Current	OFC	First-Best	Current	OFC
Panel A: Bandwidth Selection								
Scott's Rule (Baseline)	0.71	0.65	179.26	181.52	208.22	226.54	4.79	61.25
Cross-Validation	0.68	0.59	180.21	181.31	204.58	226.17	2.40	53.03
Panel B: Bootstrap Sample Size (n)								
$n = 30$	0.71	0.65	177.85	182.14	212.78	223.98	9.30	75.72
$n = 75$	0.71	0.64	179.64	182.23	203.87	212.08	7.99	74.70
$n = 100$	0.71	0.64	179.95	181.93	206.83	216.84	5.36	72.86
Panel C: Bootstrap Replications (B)								
$B = 500$	0.71	0.65	179.08	182.19	205.90	226.35	6.58	56.74
$B = 2000$	0.71	0.65	179.41	181.77	206.84	226.98	4.97	57.65

Notes: This table reports government payoffs under alternative specifications for estimating the conditional distribution $h(\lambda|\mathbf{s})$. Panel A varies the bandwidth selection method for kernel density estimation. Panel B varies the bootstrap sample size while holding $B = 1000$. Panel C varies the number of bootstrap replications while holding $n = 50$. The baseline specification (Scott's rule, $n = 50$, $B = 1000$) corresponds to the main analysis.

Table 10 reports the results across seven specifications. The baseline specification uses Scott's rule with $n = 50$ and $B = 1000$, which corresponds to the main analysis. Across all specifications, the OFC mechanism consistently outperforms both the current

contract and the pooling benchmark. The relative improvement of the OFC mechanism ranges from 53.03% to 75.72%, with a mean of 64.57% and a standard deviation of 8.86 percentage points, consistent with the baseline estimate of 61.25%.

The cross-validation yields a smaller OFC improvement (53.03% vs. 61.25%). This can be explained as follows: cross-validation selects a larger bandwidth multiplier on average (1.48 vs. 1.0), which smooths the conditional distributions $h(\lambda|\mathbf{s})$ across firms, reducing the estimated heterogeneity that the OFC mechanism can exploit. Since the smoother $h(\lambda|\mathbf{s})$ may understate the true firm heterogeneity, the cross-validation result represents a conservative estimate.

In contrast, the current contract shows small improvements of only 2.40% to 9.30% across specifications. These results confirm that the main finding—that optimally leveraging observable firm characteristics generates substantial government payoff gains—is robust to the estimation approach for $h(\lambda|\mathbf{s})$.

B.2 Robustness to the Social Value γ

The main analysis calibrates the social value parameter γ using NASA's economic impact assessment. This subsection examines whether the efficiency gains from the OFC mechanism are robust to alternative calibrations drawn from both NASA reports and the broader R&D economics literature.

Table 11 presents government payoffs under alternative values of γ. Panel A extends the baseline calibration using NASA's 2017 SBIR economic impact assessment for FY2016. This report captures the economic impact generated within that fiscal year. While the main analysis uses additional total output to calibrate γ, I also consider two alternative measures: value added and labor income. The social value parameter is calculated using the same methodology as in Section 5.2.

Panel B draws on estimates from the broader R&D economics literature. Jones and Summers (2020) calculate that the average annual gross social return to innovation ranges from 20% to 67% in the USA, depending on assumptions about capital deepening and net domestic investment effects. They also provide an alternative measure, the permanent benefit-cost ratio (BCR), finding that "$1 of R&D investment today produces, on average, a $4 to $13.3 benefit in today's dollars," calculated as BCR $= \rho/r$, where ρ is the annual social return and $r = 5\%$ is the social discount rate. Following the same approach, I convert estimates from two additional studies. Bloom et al. (2013) estimate an annual social return of 55%, capturing both private returns and technology spillovers, which implies BCR $= 0.55/0.05 = 11.0$. Fieldhouse and Mertens (2023)

Table 11: Robustness: Results under Alternative Social Value γ

Panel A: NASA SBIR Economic Impact Assessment (2017)			
	Output (Baseline)	Value Added	Labor Income
Additional Impact ($M)	474.00	262.90	169.40
γ ($M)	3.05	1.69	1.09
Government Payoff ($M):			
Pooling	179.26	72.21	24.82
Current Contract	181.52	73.54	25.73
OFC Mechanism	208.22	85.36	30.89
First-Best	226.54	94.78	36.45
Relative Improvement (%):			
Current Contract	4.79	5.87	7.83
OFC Mechanism	61.25	58.27	52.23

Panel B: Literature Estimates ($r = 5\%$)						
	Jones & Summers (2020)			Bloom et al. (2013)	Fieldhouse & Mertens (2023)	
	Cap. Deep.	Net Dom.	Baseline	—	Low	High
Social Return (%)	20.00	27.00	67.00	55.00	140.00	210.00
BCR	4.00	5.30	13.30	11.00	28.00	42.00
γ ($M)	3.00	3.98	9.98	8.25	21.00	31.50
Government Payoff ($M):						
Pooling	175.69	253.11	727.07	590.41	1597.58	2427.02
Current Contract	177.93	256.02	734.14	596.28	1612.27	2448.98
OFC Mechanism	204.24	292.68	837.50	680.21	1837.33	2788.59
First-Best	222.15	317.43	900.78	732.58	1972.21	2993.08
Relative Improvement (%):						
Current Contract	4.81	4.53	4.07	4.13	3.92	3.88
OFC Mechanism	61.44	61.51	63.57	63.16	64.00	63.88

estimate substantially higher annual social returns to public nondefense R&D of 140%–210%, reflecting greater spillover effects from publicly funded basic research; these imply BCRs of 28.0–42.0. For all literature-based estimates, I calculate the social value γ as the product of the benefit-cost ratio and the average government investment per contract (i.e., the average OFC award): $\gamma = \text{BCR} \times \bar{A}_{\text{OFC}}$, where $\bar{A}_{\text{OFC}} = \$0.75$ million.

Table 12 further examines sensitivity to the discount rate used to convert annual social returns to benefit-cost ratios, using alternative rates of 3%, 7%, and 10%. Across all specifications in Tables 11 and 12, the OFC mechanism achieves 52–64% of potential government payoff gains, with most estimates clustered between 61% and 64%, while the current contract consistently achieves only 4–8%. These results confirm that the main conclusion is robust to substantial variation in both the calibrated social value of each successful innovation and the choice of discount rate.

C Additional Congressional Objectives

Table 13 reports firm size distributions and expected participation by firm demographics under the current contract and the OFC mechanism. Panel A shows that the firm size distribution shifts downward under the OFC mechanism, with lower values at the 25th percentile, median, mean, and 75th percentile. Panel B reports expected participation across demographic groups. The number of women-owned firms and new entrants increases under the OFC mechanism. HUBZone participation remains essentially unchanged, while the number of socially disadvantaged firms declines slightly from 12 to 11.

Table 12: Robustness: Government Payoff under Alternative Discount Rates

Panel A: $r = 3\%$						
	Jones & Summers (2020)			Bloom et al. (2013)	Fieldhouse & Mertens (2023)	
	Cap. Deep.	Net Dom.	Baseline	—	Low	High
Social Return (%)	20.00	27.00	66.67	55.00	140.00	210.00
γ ($M)	5.00	6.75	16.75	13.75	35.00	52.50
Relative Improvement (%):						
Current Contract	4.36	4.21	3.95	3.99	3.87	3.85
OFC Mechanism	61.98	62.42	63.13	63.02	63.90	63.97

Panel B: $r = 7\%$						
	Jones & Summers (2020)			Bloom et al. (2013)	Fieldhouse & Mertens (2023)	
	Cap. Deep.	Net Dom.	Baseline	—	Low	High
Social Return (%)	20.00	27.00	66.67	55.00	140.00	210.00
γ ($M)	2.14	2.89	7.18	5.89	15.00	22.50
Relative Improvement (%):						
Current Contract	5.32	4.85	4.18	4.27	3.97	3.91
OFC Mechanism	59.28	61.07	62.49	62.98	63.07	64.02

Panel C: $r = 10\%$						
	Jones & Summers (2020)			Bloom et al. (2013)	Fieldhouse & Mertens (2023)	
	Cap. Deep.	Net Dom.	Baseline	—	Low	High
Social Return (%)	20.00	27.00	66.67	55.00	140.00	210.00
γ ($M)	1.50	2.03	5.03	4.13	10.50	15.75
Relative Improvement (%):						
Current Contract	6.25	5.43	4.36	4.50	4.05	3.97
OFC Mechanism	56.65	58.96	62.71	62.09	62.85	63.87

Table 13: Distributional Effects of Contract Mechanisms

	Current Contract	OFC Mechanism
Panel A. Firm Size Distribution		
Min	1	1
25th Percentile	6	5
Median	15	14
Mean	58	50
75th Percentile	61	43
Max	406	406
Panel B. Firm Demographics		
Women-Owned Firms	9	11
HUBZone Firms	1	1
Disadvantaged Firms	12	11
New Entrants	55	62