



**Defence innovation and procurement reform: an empirical
evaluation of the US Defense Innovation Unit**

Ethan B. Kapstein,
Princeton University

Javier Ospital
Bruegel

Guntram Wolff
ECARES, Université libre de Bruxelles

May 2026

ECARES working paper 2026-17

Defence innovation and procurement reform: an empirical evaluation of the US Defense Innovation Unit

Ethan B. Kapstein, Javier Ospital and Guntram B. Wolff¹

Abstract

Battlefields are rapidly evolving, with new technologies reshaping military strategy and tactics. Many of the new technologies originate in firms outside the traditional defence sector, rather than with the existing prime contractors. In the United States, the Department of Defense (DoD) created the Defense Innovation Unit (DIU) in 2015 to incentivise commercial tech companies to work on national security challenges. How successful has this new unit been in achieving its mission? In this article, we provide the first causal evaluation of the DIU's effects on defence procurement. Using administrative procurement data and a firm-level panel covering 2017-2025, and employing propensity score matching, additional firm-level covariates and a staggered difference-in-differences design, we find that the DIU has expanded both the extensive and intensive margins of defence contracting. We find not only a significant increase in the likelihood of receiving a DoD contract because of DIU treatment, but also in the size of the contract. Our findings show that defence innovation organisations can broaden and deepen the defence-supplier base. Governments updating their defence acquisition strategies in response to the lessons from recent wars can benefit from reforms that facilitate firm entry into procurement, overcoming the transaction cost and information asymmetry problems typical in defence markets.

JEL Codes: O32, O38, H57

Keywords and phrases: defence innovation; information asymmetries; public procurement; transaction costs; technological innovation.

1 Introduction

The experience in Ukraine of real-time battlefield innovation and rapidly changing defence technologies is generating a stream of lessons for military forces. Ukraine has transformed its military procurement through a highly decentralised and digitalised system, sometimes described as an “*Amazon for weaponry*”², allowing the country to become a world leader in new defence tech with multiple companies producing for the armed forces and competing for contracts. The armament strategies of the United States, the United Kingdom and France, among others, are starting to reflect these lessons. The US Department of Defense (DoD), for example, issued in November 2025 an Acquisition Transformation Strategy to accelerate

¹ Kapstein: Ethan B. Kapstein is Executive Director of the Empirical Studies of Conflict Project, Princeton University; Ospital: Javier Ospital is research assistant at Bruegel; Guntram Wolff is Professor of economics and Euroclear Chair at Université libre de Bruxelles and senior fellow at Bruegel. Contact : Javier.ospital@bruegel.org, guntram.wolff@ulb.be ; **This working paper was first published as a Bruegel working paper.**

² Serhii Kuzan, ‘Pick and Mix: Ukraine Opens an Amazon for Weaponry’, *Insights & Analysis*, 3 September 2025, Center for European Policy Analysis, <https://cepa.org/article/pick-and-mix-ukraine-opens-an-amazon-for-weaponry/>.

delivery of high-tech weaponry³, while the UK Strategic Defence Review 2025 shows a substantive doctrinal shift characterised by the integration of conventional forces with digitally enabled and increasingly autonomous systems such as drones (Ministry of Defence, 2025).

In most countries, the defence ecosystem is dominated by a relatively small number of ‘prime’ contractors; yet Ukraine’s experience and the new defence strategy documents mentioned highlight the need to enlarge the defence industrial base and tap into new, highly innovative firms. How can governments reshape procurement processes to facilitate the entry of new, predominantly commercial firms that are at the forefront of defence-related innovations or civilian innovation with large dual-use potential? The answer to that question has major implications for security and possibly also economic growth (Goolsbee, 1998; Wolff and Reinthaler, 2008; Pallante *et al*, 2023; Santoleri *et al*, 2024; Gazzani *et al*, 2025; Antolin-Diaz and Surico, 2025).

Howell *et al* (2025) provide causal evidence that the design of procurement mechanisms matters for outcomes. They evaluate a 2018 reform of the US Air Force Small Business Innovation Research (SBIR) programme and show that ‘open’ bottom-up solicitations, which allow firms to propose technologies rather than respond to the requests of procurement officers, significantly increase subsequent defence contracts, venture capital funding and patenting, while conventional competitions have no statistically significant effects. Earlier work on the SBIR program documents positive long-run effects on firm growth and commercialisation (Lerner, 1999). In its review of SBIR at the DoD, the National Academy of Sciences (2026) similarly found open-topic competitions have expanded participation in SBIR solicitations.

This paper builds on that literature to study another major innovation in how defence procurement is organised. Specifically, it investigates whether a key US institutional reform – the creation of the Defence Innovation Unit (DIU) in 2015 – has had a measurable impact on defence procurement from new and highly innovative firms. In seeking to draw such firms into the weapon-acquisition ecosystem, the DIU’s mission was to minimise the high transaction costs for commercial tech firms that might have an interest in working with the Pentagon (Williamson, 1975, 1985) along with overcoming information asymmetries concerning the demand for defence-related products (Aghion and Tirole, 1997). These transaction costs and information asymmetries are important barriers to entry for new firms and bind procurement agencies and big incumbent contractors into a mutual, self-reinforcing

³ See US Department of Defense pamphlet, undated, ‘Acquisition Transformation Strategy’, <https://media.defense.gov/2025/Nov/10/2003819441/-1/-1/1/ACQUISITION-TRANSFORMATION-STRATEGY.PDF>.

lock-ups (Becht *et al*, 2026)⁴. Has the DIU reduced these barriers to entry? Are new entrants indeed receiving DoD contracts? This article seeks to answer these questions. Government reporting has highlighted the DIU's positive contributions (GAO, 2025a), but a systematic evaluation is still missing.

Studying the effectiveness of the DIU is also relevant for policymakers in Europe, India and Canada, where governments have begun opening dedicated defence-innovation agencies, such as France's *Agence de l'Innovation de Défense* opened in September 2018, Canada's *Defence Innovation Secure Hubs (DISHs)* opened in November 2025 and Germany's *Innovationszentrum Bundeswehr* opened in February 2026. The Indian Ministry of Defence (MoD) even launched its Defence Acceleration Ecosystem (INDUS-X) in collaboration with the US DIU. Whether such institutional reforms successfully broaden the defence industrial base on which governments can draw remains an open empirical question.

The DIU seeks to accelerate the adoption of commercial technologies by selecting firms to develop prototypes which are then featured in its Commercial Solutions Catalogue, used by DoD agencies for procurement. We analyse whether this DIU treatment leads to measurable changes in procurement. Our analysis compares treated and non-treated firms' procurement likelihood and amounts. First, we use administrative data from the US General Service Administration (GSA) to match firms appearing in the DIU catalogue to a large pool of DoD contractors that are comparable in their location, sector and products sold to the DoD. In a second matching step, we use firm-level data from the database Orbis and the US Patent and Trademark Office (USPTO) to account for differences in firm revenue, age, number of employees and patenting. This two-step approach allows us to estimate the effects of DIU treatment relative to a group of untreated firms that resemble treated firms across a broad set of covariates.

Our empirical results provide strong causal evidence of the benefits of DIU: prior to engaging with DIU, our treated and comparison groups of defence contractors have the same likelihood of receiving a DoD contract; after a listing in the DIU catalogue, the likelihood of winning a defence procurement contract increases significantly. We also show that the size of procurement contracts increases substantially thanks to DIU treatment. These results are robust to numerous robustness checks.

The remainder of the paper is structured as follows. Section 2 describes policy initiatives to reform US defence procurement and the workings of the DIU in detail. Section 3 describes the administrative data for this study. Section 4 explains the propensity-score matching procedure used to construct the comparison group and describes the additional firm-level data

⁴ These lock-ups occur when the contractor has committed workforce and client-specific weapon knowledge, while the government cannot easily switch to alternative suppliers due to years of accumulated programme knowledge. The mutual hold-up power is strong which hinders innovation and makes entry by new firms more difficult.

collected from Orbis and the USPTO. Section 5 presents the empirical strategy, a staggered difference-in-differences design that accounts for firm fixed effects, time effects, and variation in treatment timing. Sections 6 and 7 present our findings and robustness checks, while section 8 concludes.

2 How US military procurement has changed to improve access for new firms

Procurement by the US Department of Defense⁵ (now called the Department of War) is a complex and lengthy process, and as such favours those firms that can navigate it and have established relations with the DoD (Sorenson, 2008). This process, governed by the Federal Acquisition Regulation (FAR), is primarily designed to acquire goods and services at scale, which is a barrier to entry for smaller firms. Further, the DoD imposes numerous operating conditions on firms (eg regarding the handling of classified information) which are costly to incorporate, creating additional barriers. The National Research Council (2014) found that the standard US procurement process favours firms with government contracting experience, because of the need to comply with these extensive administrative and reporting requirements⁶.

To reduce high transaction costs and increase the participation of small firms and startups, the DoD has introduced several grant-based innovation programmes⁷. The Small Business Innovation Research (SBIR) is a government-wide programme, overseen by the US Small Business Administration, funds research at small businesses. Launched in 1982, the programme was augmented in 1992 by the Small Business Technology Transfer (STTR) programme, which requires small businesses to partner with research institutions⁸. The DoD has made over 13,400 awards via SBIR/STTR during the programme's lifetime (NAS, 2026). The DoD is its largest user, accounting for roughly for 80 percent of SBIR/STTR spending since 2014⁹. In 2018, the US Air Force reformed its implementation of the SBIR programme by introducing an 'open topic' solicitation mechanism (GAO, 2025b), which Howell *et al*

⁵ The US Department of Defense was renamed to the Department of War by executive order on 5 September 2025. However, the legal name remains the Department of Defense until it is changed by an act of Congress, and thus this paper uses Department of Defense. See White House (2025).

⁶ GAO (2025c) also found that this acquisition process leads to lengthy adoption timelines, with major defence acquisition programmes taking nearly twelve years on average to deliver a usable first version of a system.

⁷ See for example the US Small Business Administration website 'Scorecard Details', updated 14 July 2025, <https://www.sba.gov/federal-contracting/contracting-data/small-business-procurement-scorecard/scorecard-details?agency=DOD&year=2024>.

⁸ The STTR requires awardees to work with organisations that have tech-transfer expertise (eg universities or federally funded research centres) to facilitate commercialisation. STTR was created to complement SBIR (GAO, 2025b). We include STTR within SBIR contracting.

⁹ Since 2018, the DoD has awarded \$38 billion in SBIR/STTR contracts to small companies, amounting to 15 percent of total DoD research spending in 2025 (See Kapstein *et. al*, 2026).

(2025) showed to have been a successful policy reform¹⁰. Since the success of this reform, the SBIR and STTR Extension Act of 2022 mandated that the DoD expand its open topic program beyond the Air Force and apply it to the entire department.

The Defence Innovation Unit (DIU)

The Defence Innovation Unit (DIU) is an organisation created in 2015 to increase the uptake of commercial technologies at the DoD and to bridge the gap between defence procurement and Silicon Valley (Shah and Kirchhoff, 2024). Operating directly under the Secretary of Defense, the DIU's primary goal is to transition commercially available technologies into military use throughout the entire department. Following the creation of the DIU, the DoD has launched additional innovation organisations embedded within specific services and agencies. Howell *et al* (2025) study the effect of a policy change implemented by one of these organisations, AFWERX, at the Air force¹¹. Instead of commercially available technologies across the department, these organisations are focused on the development of new technologies fit for the operational needs of their particular unit.

Based in Silicon Valley, the DIU engages innovative companies in the US and allied countries to develop prototypes using commercially available technologies (Beck, 2024). The DIU uses a bottom-up solicitation mechanism by calling for proposals through so-called commercial solutions openings (CSO). CSOs set out a defence problem and invite companies to submit short proposals without stipulating technical requirements that restrict participation (GAO, 2025a). The DIU evaluates how viable these proposals are and then selects companies to present prototypes¹². Successful prototypes are then published as products by the DIU in its Commercial Solutions Catalogue, for procurement by DoD agencies¹³.


¹⁰ The purpose was to reduce the contracting and administrative burdens that typically make DoD agencies unattractive customers for new entrants (GAO, 2025a).

¹¹ In addition to AFWERX, other focused defence innovation organisations include the Army Applications Lab, Naval X, AFWERX, SOFWERX, DEFENSEWERX, the National Security Innovation Network and the Army Venture Capital Initiative.

¹² Grants for prototype development are awarded through Other Transaction Agreements (OTAs), which allow the DoD to enter into research, prototyping and production contracts outside the normal procurement system, and are therefore not featured in standard administrative procurement data. See DIU presentation of Winter 2025, 'Defense Innovation Unit: Other Transaction Contract and Success Memo Guidance', https://s3.us-gov-west-1.amazonaws.com/publicdocs.diu.mil/DIU_Success_Memo_Walkthrough_Feb25.pdf.

¹³ See DIU website, 'Commercial Solutions Catalogue', undated, <https://www.diu.mil/solutions/portfolio/catalog>.

Figure 1: Number of proposals received by the DIU



Source: Authors based on DIU. Note: ‘awarded’ refers to ‘other transaction agreements’ (OTAs) granted to commercial companies, while ‘rejected’ means that a proposal was received but an OTA was not granted.

In 2020, the DIU received 1016 proposals in response to CSOs, more than double the previous year (Figure 1). Even though few of these proposals became operational products¹⁴, GAO (2025a) argues that the DIU has demonstrated the military application of commercial technologies. The GAO report uses document reviews and interviews with DIU officials and end users to assess six DIU projects¹⁵. In the successful cases where DIU engagement led to procurement contracts with the DoD, officials highlighted three features important for success: (1) frequent interaction between the DIU, military end users, and firms; (2) DIU staff’s ‘dual fluency’ in both defence needs and commercial technologies; and (3) the use of broadly framed problem statements in CSOs, which allow firms to propose feasible commercial solutions rather than requiring the DoD to define a solution in advance.

This descriptive evidence suggests that the DIU can increase DoD procurement of commercial technologies by reducing information asymmetries and transaction costs between defence buyers and commercial technology suppliers.¹⁶ Anduril Industries, a company created in 2017 develops AI software for the autonomous control of defence platforms. Since having its first product published in the Commercial Solutions Catalogue in 2021, it grew as a DoD vendor, reaching \$300 million in sales to the DoD in 2025¹⁷. From 2017 to 2025, the

¹⁴ Of all proposals submitted since 2016, about 6 percent have been awarded a contract for prototype development (Figure 1).


¹⁵ While GAO (2025a) is perhaps the most comprehensive evaluation of the DIU to date, it remains a small sample exercise that is descriptive rather than causal.

¹⁶ At the same time, the report warns that the DIU lacks the performance measures needed to assess whether it is solving DoD’s most critical operational gaps at scale (GAO, 2025a).

¹⁷ See DoD announcement of 23 September 2020, ‘Contracts For Sep. 23, 2020’, <https://www.war.gov/News/Contracts/Contract/Article/2358917/>.

DIU has published products from 126 other companies in its catalogue¹⁸ across five different technological areas (Figure 2). While there has still been no empirical evaluation of the agency’s effectiveness in engaging with these companies, the US Defense Innovation Board has recommended additional congressional funding for the DIU to expand its activities and focus on scaling up the production of commercial solutions (DIB, 2025).

Figure 2: Number of products published in the DIU’s catalogue by technological area



Source: Authors based on DIU. Note: ‘AI/ML’: artificial intelligence and machine learning applications. ‘Cyber’: software for the protection of critical infrastructure and warfighting systems. ‘Autonomy’: small, unmanned aircraft systems, maritime systems and drones. ‘Space’: access to space, satellite capabilities and broadband. ‘Energy’: strengthening military installations and operations. ‘Human Systems’: capabilities that improve the war fighter’s readiness, survivability and lethality.

3 Data collection and descriptive statistics

To evaluate the effectiveness of the DIU, we first constructed a novel dataset of DoD procurement. Our dependent variable is whether a contract was awarded by the DoD (to capture the extensive margin) and how large the contracts are (to capture the intensive margin). We compile this data from the US government’s GSA database¹⁹. Specifically, we downloaded data on all the contracts awarded by the DoD from 2017 to 2025²⁰. We aggregated the value of these contracts in 2024 dollars at the level of the vendor’s parent

¹⁸ See DIU, ‘Commercial Solutions Catalogue’, undated, <https://www.diu.mil/solutions/portfolio/catalog>.

¹⁹ See the GSA Data Bank, undated, <https://sam.gov/reports/awards/adhoc>, and the GSA tutorial ‘Contract Data: How to use contract data ad hoc reports’, undated, https://www.fsd.gov/gsafsd_sp/en/video-tutorials-how-to-use-contract-data-ad-hoc-reports?id=kb_article_view&sysparm_article=KB0042536.

²⁰ We filtered for contract award data from the ‘Defense, Dept. of’ and selected the following fields: ‘Year’, ‘Ultimate Parent Unique Entity ID’, ‘Unique Entity ID’, ‘Legal Business Name’, ‘Vendor Zip Code’, ‘NAICS Code’, ‘PSC Code’, ‘SBIR/STTR Code’, ‘Contracting Officers Business Size Determination Code’ and ‘Dollars Obligated’. We downloaded the data separately for small businesses and for all other firms at two-year intervals. The resulting datasets were appended to create a consolidated dataset of 806086 contracts.

company and the year of the award to obtain the annual value of contracts with the DoD by firm.

Second, we obtained data on the DIU’s engagement with companies by downloading its Commercial Solutions Catalogue²¹. We then merged this information with our procurement dataset matching company names and years²². The resulting dataset is a large, balanced, firm-year panel comprising 118159 firms over eight years, covering all firms with procurement relationships with the DoD from 2017 to 2025. The large number of firms is evidence of the wide-ranging procurement relationships of the DoD: the dataset covers all kinds of purchases, including gardening and cleaning services for facilities.

In addition to annual contract values, the GSA database provides detailed information on vendor characteristics by contract, including their location, sector, business size and the types of products and services it sells. Each contract is classified using the six-digit North American Industry Classification System (NAICS) for industry and a four-character Product and Service Code (PSC) for the type of good or service being sold. We aggregate both classification systems at the two-digit level and, for each firm-year, count how many contracts fall into each NAICS and PSC category. Federal procurement records also contain a vendor ZIP code, which we map to their corresponding metropolitan areas (Core-Based Statistical Areas, CBSAs) and create indicators that capture firms’ location by metropolitan area.

The treatment variable is defined as the first year in which a firm has had a solution published in the DIU’s catalogue, while untreated firms are those that did not receive either DIU treatment or any SBIR/STTR award.²³ We face the problem that only 122 firms were treated with the DIU catalogue listing while 118037 firms were not listed: a large imbalance between the number of treated and untreated firms. We therefore need to match treated firms with untreated firms with similar characteristics. To assess the similarity between treated and untreated firms, we compute the standardised mean difference (SMD) for each covariate k ²⁴. To compute the SMD, let x_{itk} be covariate k for firm i in year t . We average this covariate across all observed years T_i to obtain a firm-level average:

$$\bar{x}_{ik} = \frac{1}{T_i} \sum_t x_{itk}. \quad (1)$$

Second, we average \bar{x}_{ik} separately for treated and untreated firms, resulting in treated (\bar{X}_{Tk}) and untreated (\bar{X}_{Uk}) group means, as shown in equation (2). N_T and N_U denote the numbers of treated and untreated firms:

²¹ See the DIU’s Commercial Solutions Catalogue, undated, <https://www.diu.mil/solutions/portfolio/catalog>.

²² The 30 companies that appear in the solutions catalogue but not in the procurement data were kept in the dataset, with their procurement values set to zero. Of the 127 firms included in the catalogue, five did not have a usable treatment year and were therefore excluded.

²³ Since we do not have data on which firms were selected for prototypes by the DIU but are not featured in the DIU catalogue, we cannot disentangle the effects of these two steps of engaging with the DIU.

²⁴ We use the SMD rather than a t-test because, as Imbens and Wooldridge (2009) note, t-tests are sensitive to differences in sample size and become less informative when there is a large difference in the number of treated and untreated firms. Instead, standardised differences are used in matching studies (Caliendo and Kopeinig, 2008; Hotchkiss *et al*, 2012)

$$\bar{X}_{Tk} = \frac{1}{N_T} \sum_{i \in T} \bar{x}_{ik}, \quad \bar{X}_{Uk} = \frac{1}{N_U} \sum_{i \in U} \bar{x}_{ik}. \quad (2)$$

The SMD for a covariate k , shown in equation (3), is then the difference between the treated (\bar{X}_{Tk}) and untreated (\bar{X}_{Uk}) group means, divided by the pooled standard deviation ($s_{p,k}$). Here, s_{Tk}^2 and s_{Uk}^2 denote the variances of the firm-level time averages for each group:

$$\text{SMD}_k = \frac{\bar{X}_{T,k} - \bar{X}_{U,k}}{s_{p,k}}, \quad s_{p,k} = \sqrt{\frac{(N_T - 1)s_{T,k}^2 + (N_U - 1)s_{U,k}^2}{N_T + N_U - 2}}. \quad (3)$$

We compute SMD_k for all twenty-seven covariates. Table 1 shows the six covariates that have the highest absolute SMD values for sector, products and services and location. Following standard practice in propensity-score matching, we interpret an absolute SMD below 0.1 as a negligible imbalance (Cutler *et al*, 2022; Gomez-Montero *et al*, 2026). Using this benchmark for significance, we observe that DIU-backed firms are more concentrated in scientific and technical sectors, sell IT and R&D services and products and are based in the San Francisco metropolitan area more often than other DoD vendors.

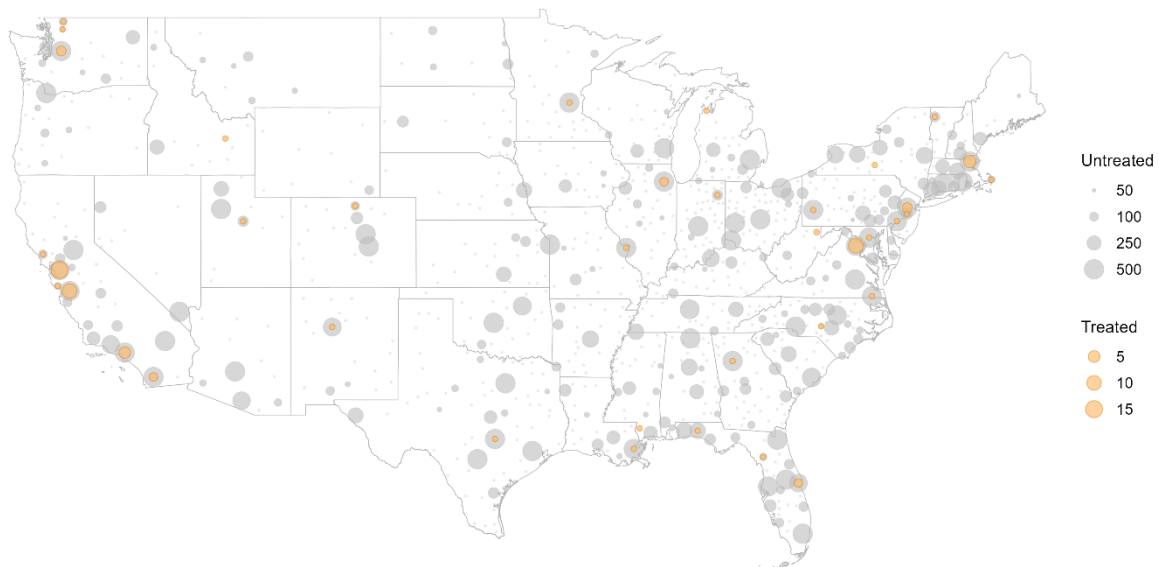
Table 1: Absolute SMD between treated and untreated firms

Category	Covariate (k)	$ SMD_k $
Sector (NAICS)	Scientific/Technical (54)	0.206
	Information (51)	0.170
	Retail trade (44)	0.136
	Machinery/Electronics (33)	0.105
	Other sectors	0.088
	Manufacturing (32)	0.077
Products and services (PSC)	IT R&D (AD)	0.280
	Communication equipment (58)	0.236
	Other products/services	0.204
	IT equipment (70)	0.188
	IT services (D3)	0.179
	Technical support (R4)	0.163
Location	San Francisco	0.146
	Boston	0.070
	Washington DC	0.056
	Miami	0.017
	Philadelphia	0.014
	Los Angeles	0.013

Source: Authors based on GSA Database.

Figure 3 documents that the location of treated and untreated firms is quite different. It follows that treated and untreated firms differ systematically on their sector, products, services and location. To address these differences, we leverage our large pool of untreated firms and detailed covariate data to match treated companies to a comparable group of non-treated firms.

Figure 3: Location of treated and untreated DoD vendors by US metropolitan area.



Source: Authors based on GSA Database. Note: ‘Treated’ refers to firms that had a solution published in the DIU’s Commercial Solutions Catalogue, while ‘Untreated’ refers to DoD vendors that did not have a solution published nor were awarded an SBIR/STTR contract.

4 Data matching

To ensure comparability between treated and untreated firms, we use propensity score matching to identify within the large pool of untreated companies those that are most similar to the treated ones. We proceed in two steps. First, we estimate each firm’s probability of receiving DIU treatment as a function of the covariates described above. Second, we use additional firm-level specific data from Orbis to further ensure comparability between the treated and untreated firms. These two steps are necessary, as collecting data in the Orbis dataset for the very large set of untreated firms would otherwise have been impractical.. As shown below, the first step already ensures substantial comparability between treated with untreated firms.


In the first step, we use nearest-neighbour matching with replacement and a caliper restriction: each treated firm was matched to a set of untreated firms with similar estimated propensity scores²⁵. The result is a sample of 115 treated firms and 456 matched untreated firms that are balanced on their sector, the products sold and location (See Figure 4). Due to our caliper restriction, seven treated firms do not have sufficiently comparable untreated matches and are therefore excluded from the sample. After matching, the absolute SMD in the estimated propensity score declines from 0.309 before matching to 0.000159 after matching²⁶. Balance in the estimated propensity score indicates that treated and matched

²⁵ Matching with replacement allows untreated firms to serve as controls for multiple treated firms, improving match quality. We also impose a caliper of 0.2, which restricts matches to treated and untreated firms with sufficiently similar propensity scores.

²⁶ ‘Propensity score’ is the probability that a firm receives DIU treatment, estimated on the observed covariates.

untreated firms are similar in their estimated likelihood of treatment given the covariates used in the first step.

Figure 4: Absolute SMD before and after propensity score matching



Source: Authors. Note: The figure reports absolute SMDs between treated and untreated firms, before and after propensity score matching. Sector variables are based on two-digit NAICS, product/service variables are based on PSC categories used in US federal procurement data. Metropolitan areas correspond to CBSAs defined by the US Office of Management and Budget. The San Francisco CBSA includes San Francisco, Oakland, Berkeley, and surrounding counties. An SMD below 0.10 indicates good covariate balance.

After constructing a matched sample that resembles the treated firms in terms of industry, location and products sold, we compare the two groups' other characteristics in a second step. From the Orbis database, we collect the incorporation date, number of employees and revenue from 2017 to 2025 for the 115 treated firms and 456 matched untreated firms. In addition, we match these firms to the USPTO Patents View database using standardised firm names and construct a firm-year panel of patenting activity. Orbis data were not used in the first matching stage because collecting firm-level information for 118149 firms would have been both impractical and unnecessary at that stage. We do not want to take into account untreated firms located far away from treated firms, which operate in totally different sectors or sell products (such as gardening services) that are unrelated to the DIU's focus. Instead, we base our matching procedure on the industry, product and location covariates that are available for the universe of DoD vendors and then collect data from Orbis for the much smaller sample of 571 firms.

Using this Orbis data, we construct an additional set of firm-level covariates. We calculate a firm's age as the number of years since its incorporation, while employment and revenue variables are computed as averages over the years in which Orbis reports non-missing values for that firm. We also construct an indicator for whether Orbis information is missing and then calculate the average value of each covariate for the treated and matched untreated groups. Table 1 compares the two groups on these covariates. Differences in missing data, firm age, employment and revenues are small in magnitude and not statistically significant. This indicates that treated firms and matched untreated firms are comparable on the additional covariates collected from Orbis. Indeed, average values of awarded contracts are similar across the two groups and not statistically different.

Table 2: Data on treated group and matched untreated

Covariates	Treated	Matched untreated	Difference	p-value
Averages for entire period				
Missing data	0.026	0.028	-0.002	0.923
Firm age (years)	14.365	15.689	-1.323	0.512
Employees (thousands)	0.782	0.820	-0.038	0.961
Revenue (\$ millions)	0.299	0.236	0.063	0.843
Pre-treatment averages				
Awarded contracts (log)	4.190	4.481	-0.291	0.617
Number of patents	27.611	0.942	26.669	0.262

Source: Authors based on Orbis and US Patent Office. Note: This table reports the mean values of Orbis covariates for firms in the treated group and firms in the matched untreated sample. N of treated firms = 115. N of matched untreated firms = 456. Table 2 also compares the average pre-treatment outcomes of treated firms to the average outcomes of the matched untreated group over the entire timespan²⁷. The reported p-values are the product of a difference-in-means t-test. Patent data was obtained from the USPTO Patents View database.

There is one notable difference: treated firms have a much higher average number of patents. This difference is not statistically significant because patenting is highly skewed, with one firm, Google, having had many patents before treatment²⁸ Google accounts for approximately 84 percent of the treated-group mean for pre-treatment patents²⁹. In the next section, we account for this difference between treated and control firms by introducing patenting as an adjustment variable in the estimator. We also examine whether the main results are driven by these patenting outliers through a leave-one-out exercise in Section 7, which shows that high-patenting treated firms do not explain the estimated effects on DoD contracting.

5 Regression methodology: difference-in-difference with staggered adoption time

After constructing the matched sample and assessing that observed characteristics of treated and non-treated firms are similar, we aim to identify the effect of the DIU treatment. This still presents three challenges. First, firms are potentially selected to participate in the program based on unobservable characteristics that may affect their contracting outcomes. Second, because defence procurement changes over time, comparing firms before and after treatment risks attributing general time trends to the policy rather than identifying its causal effect.

²⁷ Note that the treatment occurs in different periods for treated firms while untreated firms do not have a treatment time, so there is no pre-treatment time window to compute outcomes for untreated firms. Instead, we compute their mean value over the entire timespan of the data.

²⁸ The median number of patents of the treated group is 0, and excluding the two largest outliers lowers the mean from 27.6 to 0.83 patents.

²⁹ This contribution is calculated by dividing Google's firm-level pre-treatment patent average by the number of treated firms and then expressing that quantity as a share of the overall treated-group mean.

Third, since treatment occurs at different periods, treated firms need to be compared to contemporaneous controls.

Equation (3) summarises these key issues; a company's outcomes (Y_{it}) are given by firm-fixed effects (α_i), time-fixed effects (λ_t) and the effect of DIU treatment (τ_{it}) after firm i first appears in the Commercial Solutions Catalogue³⁰.

$$Y_{it} = \alpha_i + \lambda_t + \tau_{it} + \varepsilon_{it} \quad (3)$$

A staggered difference-in-differences estimator allows us to address the possible effects of unobserved characteristics, time trends and different treatment times (see Callaway and Sant'Anna, 2021). First, comparing outcomes for firms before and after treatment removes firm-fixed effects (α_i). Second, time-fixed effects (λ_t) are eliminated by comparing treated firms to control firms observed in the same year. Third, staggered treatment timing is handled by constructing group-time average treatment effects: firms first treated in year g are compared in year t only to firms that have not been treated in year t . This avoids the bias that arises in simple fixed effects models with time and firm dummies when already-treated firms are used as controls³¹. In this framework, the estimated coefficient of interest is the average treatment effect on the treated (ATT):

$$ATT(g, t) = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid G_i = g] \quad (4)$$

Here, $Y_{it}(1)$ denotes the outcome for firm i in year t under treatment, and $Y_{it}(0)$ denotes the corresponding untreated potential outcome. G_i is the first period in which firm i receives treatment. Because $Y_{it}(0)$ is unobservable for treated firms, it is identified using the outcomes of matched firms that have not been treated in the corresponding year. In the baseline specification (4), we estimate ATT on the matched sample without additional post-matching covariates, comparing firms that have been matched on observed metropolitan location, industry, and product.

This estimator also allows us to adjust the treated-control comparison by introducing additional firm-level characteristics at the estimation stage. These covariates enter as adjustment variables within the difference-in-differences estimator, so that treatment effects are identified by comparing treated firms to control firms that are also similar in these additional characteristics. This augmented specification allows us to assess whether the estimated ATT is sensitive to additional firm-level characteristics not fully accounted for in the initial matching stage. It particularly allows us to test whether the difference in patenting observed across treated and control firms in the previous section helps explain the differences in outcomes between the two groups. The corresponding treatment effect can be written as:

$$ATT(g, t; X) = E[Y_{it}(1) - Y_{it}(0) \mid G_i = g, X_i] \quad (5)$$

In this augmented specification, firm age, revenue, employment and innovative profile are introduced at the estimation stage through the covariate vector X_i . In our main results,

³⁰ Once treated, a firm remains treated in all subsequent years since their product remains in the catalogue. For firms with more than one product, we consider their first one as their treated timing.

³¹ Standard two-way fixed-effects estimators can be biased when treatment is staggered over time and treatment effects are heterogeneous (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021).

innovative profile is proxied by patenting activity, measured using pre-treatment patents for treated firms and patents observed over the sample period for controls. Section 7 shows that the results are robust to alternative definitions of this variable, including symmetric pre-treatment time-window measures applied equally to treated and control firms.

Because the DIU is a relatively recent institutional innovation and the number of treated firms is limited, the available data are better suited to identify short-run effects rather than longer-run dynamics. Therefore, our analysis focuses on a short-run event window that summarises the group-time treatment effects defined in equations (4) and (5). Let e denote event time relative to treatment; we define the short-run effect as:

$$ATT^{SR}(X) = \frac{1}{2} \sum_{e \in \{0,1\}} ATT(e; X). \quad (6)$$

That is, the short-run ATT is the average of the estimated treatment effects in the treatment year and the first post-treatment year. We restrict the analysis to this event-time window because it captures the relevant short-run pre-treatment and post-treatment periods, and because support remains broad across this range in terms of treated firms and cohorts (see annex Table 1). In the baseline specification, ATT^{SR} is estimated on the matched sample without additional post-matching covariates. In the augmented specification [$ATT^{SR}(X)$] the estimate is conditioned on the vector X_i of additional covariates.

Identification in this empirical design relies on a parallel-trends assumption: absent treatment, treated firms and their matched controls would have followed similar outcome paths. To assess the plausibility of this assumption, we estimate an event-study specification that reports group-time treatment effects by event time; we do not find statistically significant coefficients, which is consistent with comparable pre-treatment trends over the immediate pre-treatment window³². This assumption is made more plausible by the matching procedure, which produces treated and control firms with similar characteristics, and by the additional covariates showing the two groups are also broadly comparable in firm age, employment, revenue and data coverage.

6 Empirical results

In our baseline specification, DIU treatment is associated with a large and statistically significant increase in contracts with the DoD. We examine this effect on both the extensive and intensive margins. The extensive margin outcome captures whether a firm was awarded any contracts in a given year, while the intensive margin outcome captures the value of those contracts. On the extensive margin, we find that treatment raises the likelihood that a firm is awarded a contract by 11.7 percentage points, relative to the control. On the intensive margin, the baseline estimate shows a large increase in the value of procurement received by treated firms. The effect reported in Table 3 corresponds to an approximate increase of 447 percent in the value of contracts relative to matched control firms³³. Having a product published in

³² Annex table 4 and figure 5 report the dynamic effects by event-time period.

³³ On the intensive margin, we do not estimate percentage changes directly. Instead, the specification yields effects in log-points which cannot be interpreted as exact percentage effects (see Chen and Roth, 2024). For ease of interpretation, we translate these into approximate percentage terms by $\% \Delta Y = (e^{ATT} - 1) \times 100$ (van Garderen and Shah, 2002)

the DIU’s catalogue therefore makes firms more likely to secure at least one DoD contract and substantially increases the value of the contracts they receive.

Table 3 also shows augmented specifications that account for additional firm characteristics and still find large and statistically significant increases in DoD contracting. Across these specifications, the extensive-margin effect ranges from a 11.5 to 10.9 percentage-point increase in the likelihood of being awarded a contract after treatment. The intensive-margin effect remains large after we control for firm age, employment, revenue and innovative profile, corresponding to an increase of roughly 370 percent in the value of DoD contracts. While section 4 noted a large difference in the average number of patents between treated and control firms, this difference in firm’s innovation profile does not explain the estimated treatment effects on DoD contracting. The results thus remain robust after adjusting for these additional characteristics, indicating that DIU treatment is highly effective.

Table 3: Effects of the DIU, never-treated firms in control³⁴

Specification		Awarded contracts	
		Intensive	Extensive
Baseline [ATT ^{SR}]			
	No post-match conditioning on covariates	1.699***	0.117***
		(0.558)	(0.041)
Augmented [ATT ^{SR} (X)]			
	Innovative profile	1.594***	0.115***
		(0.550)	(0.044)
	Innovative profile and age	1.591***	0.111**
		(0.526)	(0.045)
	Innovative profile and employees	1.512***	0.112***
		(0.584)	(0.042)
	Innovative profile and revenue	1.563***	0.111**
		(0.555)	(0.047)
	Innovative profile, age, employees, and revenue	1.549***	0.109**
		(0.580)	(0.044)
	Number of treated	114	114
	Number of control	456	456


Source: Authors. Note: The intensive margin is measured as $\log(1 + Y_{it})$, where Y_{it} denotes awarded contracts in 2024 dollars or number of patents. The extensive margin outcome is an indicator equal to 1 if the firm had any contract awarded or if it had any patenting in that event-time period. Extensive-margin coefficients are

³⁴ Annex table 2 shows additional combinations of covariates in the augmented specification. Because a small number of firms in the matched sample have missing covariates, we re-estimate the specification after excluding these firms and the estimates remain close to the main results (see section 7 and annex table 6).

interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points. All estimates use the original matched sample; the additional covariates X are introduced at the estimation stage, we capture ‘innovative profile’ through the number of patents of each firm. Standard errors clustered at the firm level are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The one firm treated in 2017 is dropped from the estimator because it has no pre-treatment observations in the panel.

Figure 5 supports the validity of our specification when it comes to the treatment effect. Specifically, the chart shows that prior to treatment there is no discernible difference between treated and untreated firms, while the differences post treatment are statistically significant both at the intensive and extensive margins.

Figure 5: Dynamic effects of the augmented specification, awarded contracts



Source: Authors. Note: This figure shows the ATT(X) conditioned on missingness, revenues, employees and firm age, by event-time period. The left panel reports the intensive margin effects, and the right panel reports the extensive margin. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points. Vertical bars denote 95 percent confidence intervals, the dashed horizontal line marks zero effect, orange markers indicate pre-treatment periods and blue markers indicate post-treatment periods. Annex tables 3 and 4 show the estimated dynamic effects of the baseline and augmented specifications on awarded contracts and patenting.

7 Robustness checks

This section details various tests to assess the robustness of our findings. First, we conduct time placebo tests, in which treatment is assigned to treated firms two years before the actual treatment date. After re-estimating the short-run effects under this placebo timing, we find no statistically significant effects on either margin (see annex table 5). All four estimates are small relative to the main results and statistically indistinguishable from zero, which suggests that our findings are unlikely to be driven by pre-existing short-run trends before treatment.

Second, we consider unit placebo tests, in which ‘fake’ treatment is assigned at random to untreated firms. We repeat this procedure 500 times for both specifications and find placebo effects centred on zero, far from the large effects estimated for the actual treated firms (see annex figures 1 and 2). Taken together, these two placebo tests indicate that the large positive effects we identify are specific both to treated firms and to the timing of that treatment.

Third, we conduct a leave-one-out exercise, in which we re-estimate each specification repeatedly while omitting one treated firm at a time (see annex figures 3 and 4). The resulting estimates remain clustered around the original effects for both specifications and both outcomes. This indicates that our findings are not driven by any single treated firm but instead reflect a broader pattern across the treated sample. Although the treated-group average for pre-treatment patents is heavily skewed (Section 4), this leave-one-out robustness check shows that excluding patent-intensive firms does not change the estimated effects. In particular, omitting Google leaves both the intensive- and extensive-margin estimates close to their full-sample values and statistically significant.

Fourth, we modify our specifications to isolate the effect of DIU treatment. We re-estimate the baseline specification after excluding firms that received an SBIR/STTR grant from the DoD from the treated group³⁵. This reduces the estimating sample from 114 to 65 firms, but the estimated effects remain large, statistically significant and close to our baseline estimates (See annex table 6). This indicates that the DIU, and no other DoD policy interventions, are driving our main results.

Fifth, we assess whether missing Orbis covariates affect our results. We re-estimate excluding the three treated firms and 15 control firms that have missing Orbis data, and the estimated effects remain positive, statistically significant and close to the main estimates (see annex table 6). This suggests the small number of firms with missing firm-level data do not affect our estimates.

Sixth, we test the sensitivity of the results to using a control group composed of firms that receive DIU treatment in the future but have not yet been treated. This approach addresses concerns that remaining differences between treated firms and never-treated firms in our matched control group may be driving the results. Re-estimating the specification with this alternative control group, we find that the estimated effects remain positive and are close to those obtained in our augmented specifications (see annex table 6). Although these estimates should be interpreted with caution because this alternative design reduces the sample size, they provide evidence that our main results are not simply a consequence of how the matched control group is constructed.

Seventh, we show that our results are robust to alternative constructions of the patenting covariate. We re-estimate the augmented specification using patenting measures constructed over the same pre-treatment windows for both groups (see annex table 8)³⁶. We do this for three different windows, and in each case exclude firms already treated during the relevant period. The estimated effects remain positive, statistically significant, and close to our main augmented estimates, indicating that our findings are not driven by how the patenting covariate.

Finally, we consider whether DIU treatment causes patenting on the intensive and extensive margins to rise. In annex table 3, we re-estimate our baseline and augmented specifications with these alternative outcomes, we do not find statistically significant differences between treated and matched control firms before treatment, nor do we find evidence that they were on diverging patenting trajectories prior to treatment (see annex table 4). After treatment, we do not find evidence that DIU participation increases patenting on either margin. This

³⁵ Our administrative data allow us to identify DoD SBIR/STTR recipients, but not the type of solicitation through which the award was made, so we cannot distinguish between open and closed SBIR calls and exclude all firms that ever received a DoD SBIR/STTR award.

³⁶ In our main results, a firm's innovative profile is defined asymmetrically: it is based on pre-treatment patents for treated firms and on patents observed over the sample period for never-treated controls.

suggests that the positive effects of the DIU on contract awards are neither driven by pre-existing differences in patenting nor accompanied by measurable short-run increases in patenting.

8 Conclusions

This paper provides the first causal analysis of the effects of the US Defense Innovation Unit (DIU) on firm-level outcomes. We find that engagement with the DIU leads to economically large and statistically significant increases in contracting with the DoD, on both the intensive and extensive margins. By contrast, we find no measurable effects on patenting activity, suggesting that DIU's impact operates primarily through procurement and commercialisation rather than the development of new technologies.

Our results add to the previous work of Howell *et al* (2025) which examined the effects of SBIR on defence procurement. We focus on DIU, which is defined here as an institution that allows sellers of commercial technologies to gain easier access to the US DoD and thereby overcome important transaction-cost and information asymmetries. Our results show that making information on commercially available technologies more accessible to procurement agencies, for example through a catalogue of products, can substantially increase procurement from these firms. Facilitating access and reducing administrative burdens can increase the pool of innovative companies that want to work with the DoD. These lessons are also relevant for European and other countries around the world seeking to reform defence procurement processes in line with revised national defence strategies. This positive US experience suggests that similar dedicated defence-innovation organizations may also be effective in reforming defence procurement around the world.

Future research, once longer observation horizons are available, should aim to capture the longer-run impacts of DIU creation on contracting, innovation or integration into defence procurement.

References

- Aghion, P. and J. Tirole (1997) 'Formal and real authority in organizations', *Journal of Political Economy* 105: 1-29, available at <https://doi.org/10.1086/262063>
- Antolin-Diaz, J. and S. Paolo (2025) 'The Long-Run Effects of Government Spending', *American Economic Review* 115(7): 2376–2413, available at <https://doi.org/10.1257/aer.20231278>
- Becht, M., J. Mejino López and G.B. Wolff (2026) 'Who controls the defence industry?', *Working Paper*, Bruegel, available at <https://doi.org/10.64153/QGUR7847>
- Beck, D.A. (2024) *DIU 3.0: Scaling Defense Innovation for Strategic Impact*, US Department of Defense, available at <https://www.diu.mil/latest/diu-3-0-scaling-defense-innovation-for-strategic-impact>
- Callaway, B. and P.H.C. Sant'Anna (2021) 'Difference-in-Differences with Multiple Time Periods', *Journal of Econometrics* 225(2): 200-230, available at <https://doi.org/10.1016/j.jeconom.2020.12.001>
- Caliendo, M. and S. Kopeinig (2008) 'Some Practical Guidance for the Implementation of Propensity Score Matching', *Journal of Economic Surveys* 22(1): 31-72, available at <https://doi.org/10.1111/j.1467-6419.2007.00527.x>
- Chen, J. and J. Roth (2024) 'Logs with Zeros? Some Problems and Solutions', *Quarterly Journal of Economics* 139(2): 891-936, available at <https://doi.org/10.1093/qje/qjad054>
- Cutler, D., M.K. Ghosh, K.L. Messer, T. Raghunathan, A.B. Rosen and S.T. Stewart (2022) 'A Satellite Account for Health in the United States', *American Economic Review* 112(2): 494-533, available at <http://doi.org/10.1257/aer.20201480>
- DIB (2025) *Scaling Nontraditional Defense Innovation*, Report, Defense Innovation Board, 8 January, available at https://stib.cto.mil/wp-content/uploads/2026/01/2025-2_DIB-ScalingNontraditionalDefenseInnovation_250113PUBLISHED_9ee4ae.pdf
- Gazzani, A., J. Martinez, F. Natoli and P. Surico (2025) 'The Public Origins of American Innovation', CEPR Discussion Paper No. 20788, *Centre for Economic Policy Research*, available at: <https://cepr.org/publications/dp20788>
- Gomez-Montero, A., A. Garcia-Lopez, S. Cabas, A.A. Nieves-Rico and F. Giron-Luque (2026) 'Methodological Guidance on Implementing Propensity Score Matching in Observational Studies of Kidney Transplantation', *Scientific Reports* 16(1):1878, available at <https://doi.org/10.1038/s41598-025-31596-9>.
- Goodman-Bacon, A. (2021) 'Difference-in-differences with variation in treatment timing', *Journal of Econometrics* 225(2), 254-277, available at: <https://doi.org/10.1016/j.jeconom.2021.03.014>
- Goolsbee, A. (1998) 'Does Government R&D Policy Mainly Benefit Scientists and Engineers?', *American Economic Review* 88(2): 298-302, available at <https://www.jstor.org/stable/116937>
- GAO (2025a) *DEFENSE INNOVATION UNIT Actions Needed to Assess Progress and Further Enhance Collaboration*, Report to Congressional Committees GAO-25-106856,

United States Government Accountability Office, available at <https://www.gao.gov/products/gao-25-106856>

GAO (2025b) *SMALL BUSINESS RESEARCH PROGRAMS Clearer Guidance Could Improve Award Data to More Effectively Measure Outcomes*, Report to Congressional Committees GAO-25-107942, United States Government Accountability Office, available at <https://www.gao.gov/products/gao-25-107942>

GAO (2025c) *Defense Acquisition Reform: Persistent Challenges Require New Iterative Approaches to Delivering Capability with Speed*, Report to Congressional Committees GAO-25-108528, United States Government Accountability Office, available at <https://www.gao.gov/products/gao-25-108528>

Hotchkiss, J.L., M. Quispe-Agnoli and F. Rios-Avila (2012) ‘The wage impact of undocumented workers’, FRB Atlanta Working Paper Series 2012-4, available at <http://dx.doi.org/10.2139/ssrn.2060667>

Howell, S.T., J. Rathje, J.V. Reenen and J. Wong (2025) ‘Opening up military innovation: casual effects of “bottom-up” reforms to US defense research’, *Journal of Political Economy* 133(11), available at <http://doi.org/10.3386/w28700>

Imbens, G.W. and J.M. Wooldridge (2009) ‘Recent Developments in the Econometrics of Program Evaluation’, *Journal of Economic Literature* 47(1): 5-86, available at <http://doi.org/10.1257/jel.47.1.5>

Kapstein, E., J. Ospital and G. Wolff (2026) ‘Reforming European defence procurement to boost military innovation and startups’, Policy Brief 04/2026, Bruegel, available at <http://doi.org/10.64153/DRNO4343>

Lerner, J. (1999) ‘The Government as Venture Capitalist: The Long-Run Impact of the SBIR Program’, *The Journal of Business* 72(3): 285-318, available at <https://doi.org/10.1086/209616>

Ministry of Defence (2025) *The Strategic Defence Review 2025: Making Britain safer, secure at home, strong abroad*, Policy Paper, available at <https://www.gov.uk/government/publications/the-strategic-defence-review-2025-making-britain-safer-secure-at-home-strong-abroad>

NAS (2026) *Review of the SBIR and STTR Programs at the Department of Defense (2026)*, Consensus Study Report, National Academy of Sciences, available at <https://www.nationalacademies.org/projects/PGA-STEP-17-08>

National Research Council (2014) *SBIR at the Department of Defense*, National Research Council of the National Academy of Sciences, available at <https://www.nationalacademies.org/projects/PGAD-Q-06-01-A/publication/18821>

Pallante, G., E. Russo and A. Roventini (2023) ‘Does public R&D funding crowd-in private R&D investment? Evidence from military R&D expenditures for US states’, *Research Policy* 52(8), available at <http://doi.org/10.1016/j.respol.2023.104807>

Santoleri, P., A. Mina, A. Di Minin and I. Martelli (2024) ‘The Causal Effects of R&D Grants: Evidence from a Regression Discontinuity’, *The Review of Economics and Statistics* 106(6): 1495-1510, available at <https://dx.doi.org/10.2139/ssrn.3637867>

- Shah, R.M. and C. Kirchhoff, (2024) *Unit X: How the Pentagon and Silicon Valley are transforming the future of war*, Scribner
- Sorenson, D.S. (2008) *The Process and Politics of Defense Acquisition: A Reference Handbook*, Praeger
- van Garderen, K.J. and C. Shah (2002) 'Exact Interpretation of Dummy Variables in Semilogarithmic Equations', *The Econometrics Journal* 5(1): 149-159, available at <https://www.jstor.org/stable/23114086>
- Williamson, O.E. (1975) *Markets and Hierarchies: Analysis and Antitrust Implications*, Free Press
- Williamson, O.E. (1985) *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*, Free Press
- White House (2025) *Restoring the United States Department of War*, Executive order, available at <https://www.whitehouse.gov/presidential-actions/2025/09/restoring-the-united-states-department-of-war/>
- Wolff, G.B. and V. Reinthaler (2008) 'The effectiveness of subsidies revisited: Accounting for wage and employment effects in business R&D', *Research Policy* 37(8): 1403-1412, available at <https://doi.org/10.1016/j.respol.2008.04.023>

Annex

Table 1: Support by event time in the baseline specification

Event time	Number of firms	Number of cohorts
-5	81	4
-4	98	5
-3	106	6
-2	112	7
-1	113	7
0	111	7
1	93	6
2	60	6
3	47	4

Source: Authors. Note: counts are computed from the actual estimator sample used in the baseline dynamic-effects specification with never-treated controls. The table shows the number of treated firms and treatment cohorts contributing to each event period.

Table 2: Effects of the DIU, never treated firms in control

Augmented [ATT ^{SR} (X)]		Awarded contracts	
		Intensive	Extensive
	Pre-treatment patent history, age, and employees	1.551***	0.110***
		(0.573)	(0.043)
	Pre-treatment patent history, age, and revenue	1.576***	0.109**
		(0.585)	(0.045)
	Pre-treatment patent history, employees, and revenue	1.516***	0.108**
		(0.558)	(0.043)
	Number of treated	114	114
	Number of control	456	456

Source: Authors. Note: standard errors clustered at the firm level are in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table 3: Effects of the DIU, never-treated firms in control

Specification		Patenting	
		Intensive	Extensive
Baseline [ATT ^{SR}]			
	No post-match conditioning on covariates	-0.0186	-0.0334
		(0.0391)	(0.0326)
Augmented [ATT ^{SR} (X)]			
	Age	-0.0183	-0.0326
		(0.0421)	(0.0335)
	Employees	-0.0192	-0.0328
		(0.0422)	(0.0336)
	Revenue	-0.0185	-0.0333
		(0.0415)	(0.0298)
	Age, employees and revenue	-0.0200	-0.0329
		(0.0426)	(0.0326)
	Number of treated	114	114
	Number of control	456	456

Source: Authors. Note: this table reports the ATT in the corresponding event-time period. The intensive margin is measured as $\log(I + Y_{it})$, where Y_{it} denotes the number of patents corresponding to firm i in year t . The extensive margin outcome is an indicator equal to one if the firm filed for any patent in that year. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points.

Table 4: Dynamic effects of the DIU

Event time	Awarded contracts		Patenting	
	Intensive	Extensive	Intensive	Extensive
-2	0.274	0.014	0.054	0.056
	(0.598)	(0.047)	(0.056)	(0.043)
-1	0.073	-0.007	0.055	0.047
	(0.641)	(0.045)	(0.046)	(0.038)
0	2.079***	0.141***	-0.064	-0.051
	(0.505)	(0.039)	(0.042)	(0.035)
1	1.020	0.078	0.023	-0.009
	(0.842)	(0.058)	(0.061)	(0.042)
Number of treated	114	114	114	114
Number of control	456	456	456	456

Source: Authors. Note: this table reports the ATT (X) conditioned on innovative profile, age, revenues and employees in the corresponding event-time period. The intensive margin is measured as $\log(1 + Y_{it})$, where Y_{it} denotes either awarded contracts (in 2024 dollars) or the number of patents corresponding to firm i in year t . The extensive margin outcome is an indicator equal to one if the firm had any contracts with the DoD in that year, or if the firm filed for any patent in that year. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points.

Table 5: Time placebo tests for short-run effects on awarded contracts

Awarded contracts	Baseline		Augmented	
	ATT	p-value	ATT(X)	p-value
Intensive	0.470	0.395	0.271	0.637
Extensive	0.021	0.585	0.010	0.801

Source: Authors. Note: This table reports time placebo estimates for short-run effects on awarded contracts. Treatment is artificially reassigned to two years before the true treatment date for treated firms. None of the placebo estimates are statistically significant. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points.

Table 6: Alternative specifications

Alternative specification	Awarded contracts		Patenting		N. of treated	N. of control
	Intensive	Extensive	Intensive	Extensive		
(1) Exclude SBIR/STTR treated firms	1.747***	0.133***	-0.031	-0.049	65	456
	(0.641)	(0.050)	(0.048)	(0.037)		
(2) Exclude firms with missing Orbis data	1.602***	0.113**	-0.017	-0.031	111	441
	(0.606)	(0.048)	(0.044)	(0.036)		
(3) Not-yet-treated firms in control group	2.223**	0.161**	-0.152	-0.110	53	-
	(0.899)	(0.071)	(0.097)	(0.067)		

Source: Authors. Note: This table reports estimates from modified versions of the ATT^{SR} specification. In alternative specification (1), excluding firms that ever-received SBIR/STTR grants reduces the number of treated firms to 65. In specification (2), not-yet-treated firms are used as controls, which reduces the number of treated to 53 because the first and last treated cohorts cannot serve as controls. In specification (3), excluding firms with missing Orbis data reduces the sample from 114 treated and 456 control firms to 111 treated and 441 control firms. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points. Standard errors clustered at the firm level are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table 7: Effects of the DIU across periods excluding SBIR-treated firms

Event time	Awarded contracts		Patenting	
	Intensive	Extensive	Intensive	Extensive
-2	0.052 (0.632)	-0.002 (0.048)	0.060 (0.088)	0.035 (0.052)
-1	-0.355 (0.701)	-0.037 (0.055)	0.014 (0.064)	0.032 (0.046)
0	1.785*** (0.598)	0.134*** (0.045)	-0.068 (0.051)	-0.078* (0.040)
1	1.709** (0.790)	0.132** (0.061)	0.006 (0.058)	-0.021 (0.034)
N. of treated	65	65	65	65
N. of control	456	456	456	456

Source: Authors. Note: this table reports the ATT in the corresponding event-time period for alternative specification (1) in annex table 6. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points.


Table 8: Robustness to alternative patent-based measures of firms' innovative profiles

Alternative measures of innovation profile	Awarded contracts		N. of treated	N. of control
	Intensive	Extensive		
Pre-treatment window, 2017-2018	1.670*** (0.566)	0.114*** (0.042)	113	456
	1.582*** (0.607)	0.108** (0.045)		
Pre-treatment window, 2017-2019	1.517** (0.616)	0.112** (0.050)	99	456
	1.514*** (0.581)	0.106** (0.043)		
Pre-treatment window, 2017-2020	1.514*** (0.581)	0.106** (0.043)	114	456
	1.514*** (0.581)	0.106** (0.043)		
Full observed period, 2017-2025	1.514*** (0.581)	0.106** (0.043)	114	456
	1.514*** (0.581)	0.106** (0.043)		

Alternative measures of innovation profile	Awarded contracts		N. of treated	N. of control
	Intensive	Extensive		


Source: Authors. Note: estimates report the augmented $ATT(X)^{SR}$ for awarded contracts. In each specification, innovative profile is proxied by patenting activity measured over a common time window applied identically to treated and control firms. For the pre-treatment window specifications, treated firms whose treatment year falls within the relevant window are excluded to avoid contamination from post-treatment patenting. Extensive-margin coefficients are interpreted as percentage-point changes, and intensive-margin coefficients are reported in log points.

Figure 1: Distribution of unit-placebo tests, intensive margin of DoD awarded contracts



Source: Authors. Note: the figure reports the distribution of unit-placebo ATT^{SR} estimates on the intensive margin of awarded contracts. The left panel shows the baseline matched specification (ATT^{SR}), and the right panel shows the covariate-adjusted specification ($ATT^{SR}(X)$). The red vertical line marks our estimated effect for each specification. In the augmented specification, the distribution is computed from successful placebo replications only.

Figure 2: Distribution of unit-placebo tests, extensive margin of awarded contracts

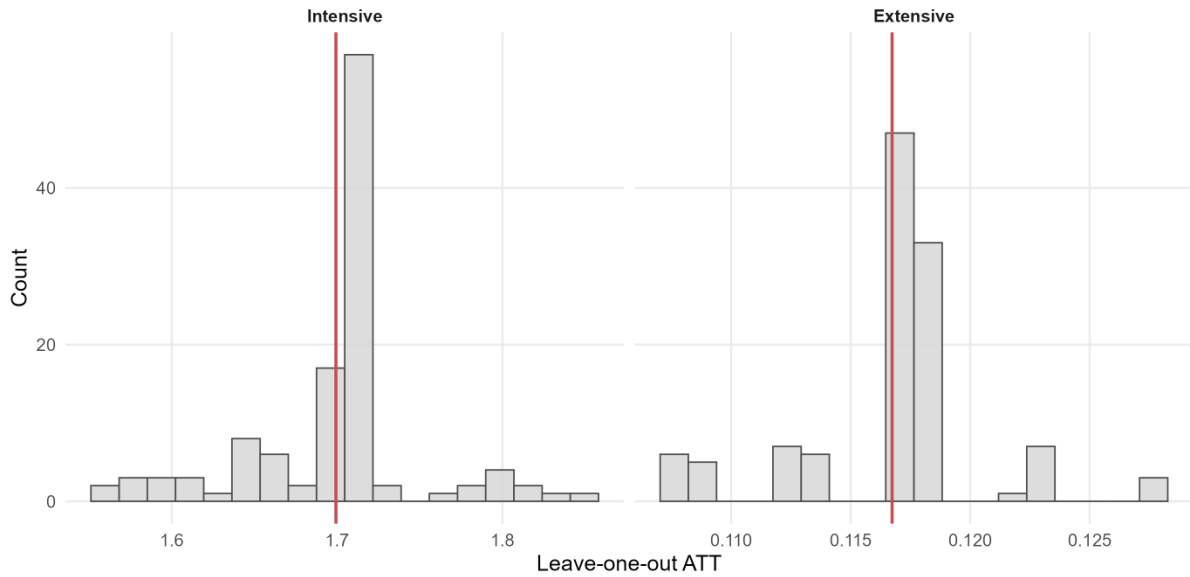


Source: Authors. Note: the figure reports the distribution unit-placebo ATT^{SR} estimates on the extensive margin of awarded contracts to the DoD. The left panel shows the baseline specification (ATT^{SR}) and the right panel shows the covariate-adjusted specification ($ATT^{SR}(X)$). The red vertical line marks our estimated effect for each specification.

Annex figures 1 and 2 show that the unit-placebo distributions are centred near zero in their central mass, while our actual short-run estimates, shown by the red vertical lines, lie in the right tail of each distribution. For the augmented specification on the intensive margin, the placebo density is based on 344 successful replications, since 156 of the 500 placebo replications did not converge³⁷. The estimation failed in these replications because, after random assignment to treatment and conditioning on covariates, some placebo-treated firms had too few comparable control observations for reliable estimation. This does not alter the substantive conclusion, since most placebo replications are successful and continue to produce estimates centred near zero and far from the observed effect.

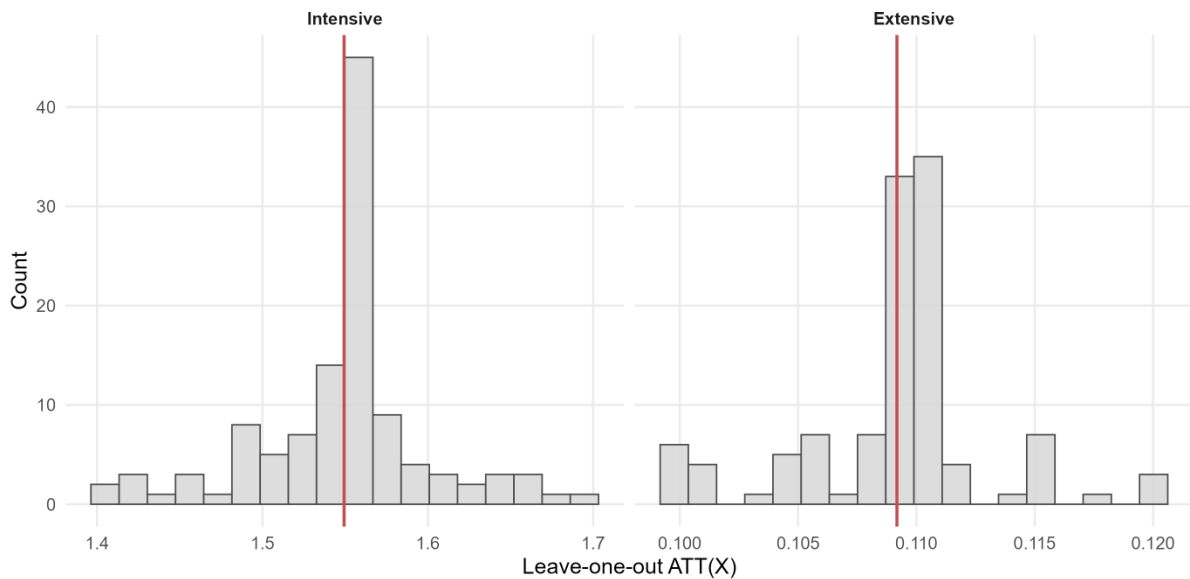
Figure 3: Leave-one-out replications, baseline specification

³⁷ All 156 failed placebo replications on the intensive margin were singular-matrix errors, indicating that after random placebo assignment and conditioning on covariates, some placebo samples became too sparse for reliable estimation. No failures occur on the extensive margin.



Source: Authors. Note: the figure shows the distribution of leave-one-out ATT^{SR} estimates for the baseline specification. In each replication, one treated firm is omitted, and the short-run effect is re-estimated without re-matching the sample. The red vertical line marks the corresponding estimate in the full sample.

Figure 4: Leave-one-out replications, augmented specification



Source: Authors. Note: the figure shows the distribution of leave-one-out estimates for the augmented specification, $ATT^{SR}(X)$. In each replication, one treated firm is omitted, and the short-run effect is re-estimated without re-matching the sample. The red vertical line marks the corresponding estimate in the full sample.

Figures 3 and 4 show that the resulting leave-one-out estimates are clustered around the original estimated effects for both specifications and outcomes. This indicates that no single firm is driving our estimated effects; instead, they reflect a broad pattern of increases in contracting across the treated sample. In the baseline specification, all leave-one-out replications remain statistically significant. In the augmented specification, which conditions on firm age, employment, revenue and patent history, all leave-one-out replications also remain positive and statistically significant on both margins. Taken together, these results indicate that the

estimated effects are broadly shared across treated firms rather than being driven by any single outlier.