

Giuffrida, Leonardo M.; Raiteri, Emilio

**Working Paper**

## Bureaucratic Frictions and Innovation Procurement

CESifo Working Paper, No. 10775

**Provided in Cooperation with:**

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

*Suggested Citation:* Giuffrida, Leonardo M.; Raiteri, Emilio (2023) : Bureaucratic Frictions and Innovation Procurement, CESifo Working Paper, No. 10775, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/282463>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Bureaucratic Frictions and Innovation Procurement

*Leonardo M. Giuffrida, Emilio Raiteri*

## **Impressum:**

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email [office@cesifo.de](mailto:office@cesifo.de)

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: [www.SSRN.com](http://www.SSRN.com)
- from the RePEc website: [www.RePEc.org](http://www.RePEc.org)
- from the CESifo website: <https://www.cesifo.org/en/wp>

# Bureaucratic Frictions and Innovation Procurement

## Abstract

Is work overload a friction to public agencies? Using data on R&D procurements, patents, and contracting units from a US federal agency, we investigate how officer workload impacts innovation procurement outcomes. Unanticipated retirement shifts provide an exogenous source of variation that we exploit as an instrument for workload. When workload declines, we find a significant increase in patent rates. One additional officer leads to a 28 percent increase in the probability that a contract will generate a patent. Our findings suggest that officers burdened with excessive workloads may not provide adequate guidance to R&D suppliers when it is most needed.

JEL-Codes: D230, D730, H570, J240, O310.

Keywords: workload, procurement, bureaucrats, R&D, patents, instrumental variable.

*Leonardo M. Giuffrida*  
ZEW Mannheim / Germany  
[leonardo.giuffrida@zew.de](mailto:leonardo.giuffrida@zew.de)

*Emilio Raiteri*  
School of Innovation Sciences, Eindhoven  
University of Technology / The Netherlands  
[e.raiteri@tue.nl](mailto:e.raiteri@tue.nl)

First version: July 29, 2021

This version: November 7, 2023

An earlier version of this paper circulated under the title “Buyers’ Workload and R&D Procurement Outcomes: Evidence from the US Air Force Research Lab.” For their helpful comments and suggestions, the authors would like to thank Zareh Asatryan, Gian Caspari, Adriano De Leverano, Francesco Decarolis, Marian Moszoro (discussant), and Janne Tukiainen, as well as participants at departmental seminars (ZEW Mannheim, University of Turku, Utrecht University, ITAM), workshops (EPPE 2020 at KU Leuven, MaCCI IO Day 2020 at Uni Mannheim, Berkeley-Padova-Paris 7th Organizational Economics Workshop at University of Padova) and conferences (EARIE 2021 Meeting in Bergen, EPIP 2021 Meeting in Madrid) where earlier drafts of this study were presented. Corrado Carniato provided excellent research assistance. Giuffrida is thankful to the Leibniz Association for its support through the project “Market Design by Public Authorities.” The usual disclaimers apply.

# 1 Introduction

Organizational frictions in hierarchical structures affect transaction costs and, by extension, impact decision-making processes (Williamson 1975). When it comes to government organizations, such frictions may alter the course and outcomes of public policy (Niskanen 1968). In this paper, we examine workload as a source of bureaucratic friction and use an instrumental variable design to identify its effect on the provision of public goods. Using data on US federal procurement, this paper quantifies the causal impact of buyer workload variation on patents as an innovation contract outcome.

Governments use their extensive purchasing power to pursue an array of policy objectives. Indeed, recent studies have highlighted that technical and scientific progress depends crucially on public spending (Fleming et al. 2019; Moretti et al. 2019). In addition to granting intellectual property rights (Jaffe 2000) and adopting supply-side policies (Dechezleprêtre et al. Forthcoming; Azoulay et al. 2019; Santoleri et al. 2022), governments can effectively spur innovation by directly procuring R&D from businesses and institutes of higher education. In the US, R&D procurement accounts for about one-third of the \$130 billion annual federal R&D budget (de Rassenfosse et al. 2019)—with the remainder distributed among other types of extra-mural research (e.g., grants) and in-house research (Bruce and de Figueiredo 2020). Even though the primary objective of R&D procurement is to acquire innovative supplies or services for the direct benefit of federal agencies (Federal Acquisition Regulation [FAR] § 35.003), policymakers and scholars increasingly view this type of public outlay as a crucial tool of innovation policy. Anecdotal and historical evidence indicates that US federal procurement has been critical to developing some of the most influential technologies of the 20th and 21st centuries (Ruttan 2006). Recent empirical evidence supports such a link (Raiteri 2018).<sup>1</sup>

Innovation scholars argue that buyers play a key role by communicating needs and specifying the functional requirements that procured activities or goods must fulfill (Edquist 2015). Despite empirical evidence confirming that contracting personnel have a direct impact on innovation outcomes (Bruce et al. 2019; Decarolis et al. 2021), the underlying causal mechanisms remain a black box. Unpacking this issue promises to shed light on how bureaucracies may promote (or stifle) innovation (Kelman 2021). In light of the pressing issue of under-staffing in government procurement offices (Rendon et al. 2012; AAP 2007; GAO 2017; Rau and Stammersky 2009), this paper examines whether buyer workload is a source of capacity constraints in innovation contracting.

We spotlight the role of the contracting officer (CO)—i.e., the leading official in the acquisition process—as a buyer. COs are responsible for the elaboration of the acquisition strategy and have the authority to enter into and terminate contracts on behalf of the federal government (FAR § 1.602). Higher CO workloads may be particularly disruptive in the context of R&D procurement as they may prevent COs from adequately addressing the technical complexities associated with the granting of an R&D contract. Particularly in the context of increasing R&D budgets, such capacity constraints are cause for concern, as they could potentially prevent the attainment of innovation policy goals.

Despite the key importance of COs in the acquisition process, no prior work has analyzed how

---

<sup>1</sup>See Chiappinelli et al. (2023) for a recent survey of the innovation procurement literature.

CO workload may impact the government’s ability to procure innovative solutions to its needs. We seek to address this question by combining several different data sources for the very first time. First, we draw on information on the quasi-universe of contracts awarded by US federal agencies. Second, we rely on patented inventions associated with a federal contract; in line with the literature, we use patents as a measure of R&D contract performance (Bruce et al. 2019). Third, we collect extensive information on the solicitation process including information concerning the CO identity. To the best of our knowledge, we are the first to incorporate data specifically related to the identity of the CO. The inclusion of this information lends a unique dimension to our analysis and provides greater insight into the R&D procurement process. This information is more readily available for contracts awarded by the US Department of the Air Force, one of the major subdivisions of the US Department of Defense (DoD). As nearly half of all R&D contracts awarded by the Air Force come from the Air Force Research Laboratory (AFRL) and given its high degree of specialization in R&D procurement, we focus on contracts awarded by the AFRL contracting offices. Controlling for the number of R&D awards and budget of the office, we use the headcount of COs actively procuring R&D contracts in a given year as an inverse workload indicator. After this data selection and merging process, we end up with a sample of 1,970 R&D contracts, which represent the quasi-universe of AFRL’s R&D procurement processes, including CO information, for 2005 through 2012.

Identifying the effect of buyer workload on R&D contract performance presents multiple empirical challenges. For example, the quality of the CO assigned to a particular contract is known by the office manager (also referred to as the program manager) when planning task allocation, and CO quality may correlate both with workload and the likelihood a project will generate patents. The average complexity of the office’s yearly procurement activity, which is also anticipated by the program manager, may also introduce further omitted variable bias. For identification purposes, we use an instrumental variable (IV) approach. We use a fourth data source which reports statistics on the federal workforce. In particular, we construct an instrument that builds on the fact that retirement decisions among federal employees are strongly influenced by i) the attainment of the threshold years of service that qualify workers for immediate pension benefits and ii) idiosyncratic motives (Asch et al. 2005; Feldman 1994; Gustman and Steinmeier 1993; Lewis and Pitts 2018). Therefore, we consider unexpected retirement-postponing decisions—defined as the difference between the number of contracting employees eligible for retirement (anticipated by the office manager) and actual retirees (unanticipated by the manager)—as a good potential source of exogenous variation in workload. Specifically, as managers’ hiring decisions are based on expected retirements, the larger the gap between managerial expectations and reality, the larger the short-run positive shock to contracting employment in an office. We show that our IV is strong and does not correlate with any relevant dynamic characteristics of the work unit. We augment our IV approach with CO fixed effects, which allows us to hold unobserved CO features and interpret the coefficients as the impact of *individual* workload on R&D award outcomes.

Our results show that the same CO exposed to decreased workload induces an increase in the average probability of an award delivering a patented invention. Specifically, one additional CO colleague in the procurement unit (corresponding to 3 percent of the average number of COs in the office–year pair in our sample) leads to a 2.5 percentage point increase in the probability of

an R&D contract generating a patent. The effect corresponds to 28 percent of the average patent rate in the sample. Our IV estimates are more than one order of magnitude larger than estimated by the corresponding endogenous model. To provide a more transparent economic interpretation of the estimates, we consider what would happen if we used them to infer the effect of raising the workload of all AFRL’s office-year combinations to the level of the office-year with the largest workload in our sample. In this case, patents would be 50 percent less likely (i.e., the number of patents per contract would be 13 percent lower).

Our results are robust to several modifications in the empirical specification. In particular, our findings are robust to the metric of contract performance. Based on the logic of relational contracts (Calzolari and Spagnolo 2009)—which are well-suited for R&D contracts, given the difficulty of stipulating quality in the contractual terms—we expect firms that fulfill the scope of the contract to be rewarded subsequently by the same buyer with non-compete awards (Che et al. 2021), even in absence of filed patents. We propose follow-on contracts as an alternative metric for contract success. Consistently with our focal analysis, we find that a reduction in CO workload triggers a significant increase in the probability of the contractor being awarded a follow-on R&D contract without competition.

We provide suggestive evidence regarding the importance of the CO for a clear work statement when drafting solicitations. We enrich our dataset by collecting (for the first time to our knowledge) records of pre-solicitation information on R&D procurement enabling us to establish a novel correlation between CO workload and pre-award schedules. First, we demonstrate that increased workload directly results in more stringent pre-award time constraints. This is evident from the reduced time frame afforded to firms for proposal submission, a factor solely under the CO’s discretion. Second, less time available to a CO induces tighter time constraints, which in turn results in poorer contract outcomes. The workload effect is more disruptive at the margin when the time available to conclude the transaction is lower—i.e., toward the end of the fiscal year or when the average workload level in the unit is high. Third, we contend that as time constraints tighten, the quality of guidance provided by the CO in the pre-award phase diminishes. Our findings show that suppliers most in need of this guidance—specifically smaller and less experienced contractors—are more severely affected by the overloaded COs. In contrast, we find that workload does not impact the innovative output of large and experienced firms. Although we are aware that these mechanism results do not provide definitive evidence of a direct causal link, they strongly suggest that overloaded COs cannot devote sufficient time to tender and contract specifications, thus inducing a worsening of innovation outcomes, and an effect that is particularly pronounced when guidance is necessary.

**Related literature** This work relates to three strands in the literature. First, we add to the extensive empirical literature on the drivers of procurement efficiency. Specifically, past studies have examined awarding mechanisms (Decarolis 2014, 2018), the end-of-year spending rush (Liebman and Mahoney 2017), external audits (Gerardino et al. 2017), industry consolidation (Carril and Duggan 2020), performance-based insurance schemes (Giuffrida and Rovigatti 2022), and centralized purchasing (Bandiera et al. 2009). Our results advance this scholarship by emphasizing bureaucratic workload variation as a driver of contract performance. Accordingly, our findings

expand in particular on the literature concerned with how buyers as agents affect procurement outcomes (Decarolis et al. 2020; Buccioli et al. 2020; Baltrunaite et al. 2021; Spenkuch et al. 2023; Best et al. 2023).

To our knowledge, the only other study on buyer workload and contract performance is Warren (2014), which is concerned with regular procurement. The author argues that as workloads increase, the agency will have less time to spend on each task, leading to sub-optimal contract specifications and poorer execution outcomes. Specifically, contracts awarded by busier agencies undergo more modifications to the original contract and obligate greater spending. We diverge from Warren (2014) in two main ways. First, we focus on R&D contracting. While in traditional procurement the CO usually seeks existing goods or services to fulfill a well-defined requirement or need of the public agency, in R&D procurement the CO looks for novel solutions to unresolved problems or areas of interest without prescribing specific solutions. When procuring R&D, buyers have greater flexibility, can engage in discussion with potential bidders more freely, and evaluate proposals based on technical merit rather than cost and time to delivery. In essence, while both fall under the umbrella of procurement, regular procurement and R&D procurement operate in fundamentally different realms, each tailored to serve its unique objectives. Given the higher complexity and discretion inherent to procuring R&D services, spikes in buyer workload are more likely to have a dramatic effect on the acquisition process. Second, in Warren (2014), buyer workload have neutral effects (i.e., limited competition) or positive effects (i.e., cost overruns with cost-plus contracts) for suppliers; on the contrary, we show that higher buyer workload induces less innovative R&D activity, which is a negative outcome for both sides.

In innovation procurement, a first analysis of the buyer relevance is highlighted by Bruce et al. (2019). The paper emphasizes that federally funded R&D grants in the US have poorer innovation output than comparable cooperative agreements (i.e., a contract-like arrangement) precisely because the buyer has greater discretion over the latter. When discretion is more relevant, such as for R&D contracts, any friction to officers may plausibly imply even worse outcomes. Our results indicate that the government should pay particular attention to the workload of contracting officers in R&D contracting and ensure adequate unit staffing. Unlike Bruce et al. (2019), our results show that more oversight in the acquisition process results in better contract performance. Decarolis et al. (2021) provide the first empirical quantification of the role of public buyers in the success of R&D procurement. In particular, the authors exploit variation in manager deaths across agencies and years to estimate the probability that contracts deliver patented inventions. Even a small increase in the share of manager deaths causes a sizeable decline in patents per contract. Deaths occurring in the six months before the contract award drive this effect, indicating the relevance of the design stage relative to the execution stage. We add to Decarolis et al. (2021) in two respects. First, we consider the friction from workload as a mechanism for the buyer role. The data and methods of Decarolis et al. (2021) do not provide conclusive evidence on the mechanisms by which a sudden loss of specialized human capital affects single contract performance. Our evidence on the workload effect applies to the award year and validates the hypothesis that the award phase is pivotal in R&D procurement. Second, we show that the decreasing supervision of the process induced by a higher workload removes surplus from the contractor in the form of less patenting, yet without shedding further light on the buyer–supplier relationship, which is beyond the scope



of their paper.

Furthermore, in contrast to both Warren (2014) and Decarolis et al. (2021), we focus on single officers rather than units, leveraging individual workload variation. Our approach, therefore, sheds first light on how bureaucrats' features may affect their job performance beyond the self-selection argument posited by the relevant literature (e.g., Friebel et al. 2019; Ashraf et al. 2020; Hanna and Wang 2017).

Second, we contribute to the broad literature on the effectiveness of innovation policies. See Bloom et al. 2019 and Bryan and Williams 2021 for overviews. Our results show that capacity constraints for front-line officials could limit the ability of the government to translate funding allocated to R&D work into valuable knowledge and innovation. This, in turn, could dampen the positive spillover effects that publicly funded R&D has been shown to have on the economy (Fleming et al. 2019; Moretti et al. 2019). Our research has parallels to Furman and Stern (2011), as we are interested in understanding how institutions shape the formation of new knowledge and the frictions they encounter in this regard.

Third, our results expand on the studies concerned with organizational frictions. Transaction cost economics, as exemplified by Williamson (1975), emphasize how internal organizational frictions inherent in structural arrangements can profoundly influence economic behaviors and decision-making processes. The operations management literature stresses how workload is a source of friction that affects outcome of activities undertaken not only by organizations such as hospitals (Berry Jaeker and Tucker 2017) and commercial banks (Xu et al. 2022), but also by individual workers such as like midwives (Freeman et al. 2017) and physicians (Shurtz et al. 2022). We add to this strand of research by documenting how high CO workload can impact procurement outcomes. In doing so, we contribute to the long-standing research in economics, public administration, and management on the sources of organizational frictions and their implications for the nature, scope, and effectiveness of government activity (Niskanen 1968; Peters 2018; Hood 1995).

The remainder of the paper is organized as follows. In Section 2, we outline the institutional setup, identification problem, and our research design. Data sources and sample selection are described in Section 3. Section 4 presents our baseline results. In Section 5, we discuss the mechanisms underlying our findings. Section 6 concludes.

## **2 Institutional and empirical setting**

In this section, we introduce the necessary institutional background of US federal procurement concerning the role of COs and the most used solicitation mechanisms. We will then document the trends in understaffing for agencies involved in R&D procurement. Finally, we present our solution to circumvent the identification concerns for our analysis.

### **2.1 The role of the contracting officers in R&D procurement**

As briefly discussed in the introduction, one of the main elements of the novelty of our work relies on focusing on the role of COs in the specific context of federal R&D procurement. According to FAR § 1.602, a CO is a government official who is legally authorized to enter into, administer,

and/or terminate contracts on behalf of the federal government. The CO identifies the requirements that the procured good or service needs to satisfy, conducts primary and secondary market research to assess the market availability of solutions satisfying the agency’s needs, drafts and issues the formal request for proposals, and, eventually, selects the best source. The CO is hence responsible for the whole acquisition process. Depending on the agency, CO might also be in charge of contract administration, but other contracting personnel often take the leading role in the contract management phase.<sup>2</sup>

Even in the context of regular procurement, the CO has wide discretionary power in performing her tasks. The FAR states that CO “should be allowed wide latitude to exercise business judgment” (FAR § 1.602-2). In the acquisition of R&D work, the importance of the CO’s judgment is significantly heightened. As reported in FAR § 35.001, the objective of an R&D contract is to advance scientific and technical knowledge. Unlike contracts for supplies and services, R&D contracts are directed toward objectives for which the work or methods cannot be precisely described in advance. The main task of the CO is to translate a rather abstract idea into contractual and work requirements that are clear to the potential contractor (US Air Force 1967). To succeed in such a complex task, the CO must develop an ad-hoc approach for each new procurement activity and allow contractors the freedom to exercise innovation and creativity (FAR § 35.005). Therefore, the acquisition strategy needs to be individually tailored to attain the desired objectives of a specific contract.

A substantial workload for an office’s acquisition team can significantly affect the ability of COs to allocate adequate time for tailoring solicitations to the individual R&D needs of the agency and for effective communication with potential vendors. This, in turn, may profoundly impact the performance of the R&D contracts they award.

## **2.2 Soliciting R&D: Requests for Proposal and Broad Agency Announcements**

The COs use requests for proposals when the agency has a specific procurement need, with clear criteria and expectations for deliverables (FAR § 15.203). The solutions provided by suppliers will result in a contract. Requests for proposals can also be used to solicit R&D. In our dataset, a request for proposal predates each R&D contract and corresponds to its solicitation phase.

To increase the likelihood of success in R&D contracting, FAR permits the CO to diverge from the stringent requisites of requests for proposals. A Broad Agency Announcement (BAA) is the most common alternative mechanism to solicit proposals for federal R&D.<sup>3</sup> BAAs are broader in their subject matter than requests for proposals. COs use them to fulfill their requirements for scientific study and experimentation directed toward advancing the state-of-the-art or increasing knowledge rather than focusing on a specific solution.

BAAs provide a particularly flexible and interactive avenue for R&D procurement compared to standard requests for proposals. While agencies are encouraged to promote early exchanges of

---

<sup>2</sup>Monitoring is carried out in collaboration with the CO representative. FAR § 1.604 specifies that a CO representative “assists [the CO] in the technical monitoring or administration of a contract.” In practice, in a memorandum on contracting practices audited at the AFRL facilities in 2007, COs are reported to “[...] designate qualified personnel as their authorized representatives to assist in either technical monitoring or administration of a contract”. See <https://media.defense.gov/2007/Sep/28/2001712655/-1/-1/1/07-130.pdf>.

<sup>3</sup>See FAR § 35.016. See an extract of a BAA in Appendix E.

information about future acquisitions with potential vendors in procuring all kinds of goods and services (FAR § 15.201), the interactions after a formal solicitation is issued are heavily regulated for regular procurement processes. The BAAs instead permit continuous communication between the CO and potential vendors, allowing technical discussions to understand the feasibility of proposed R&D efforts even after the publication of the announcement. Relying on BAAs fosters more open dialogue with potential vendors, enabling a deeper understanding and discussion of innovative proposals. The CO is instrumental in this process, clarifying the BAA's guidelines and overarching expectations and resolving ambiguities for potential contractors before they submit their proposals (see Section 5).

### 2.3 Trends in the US contracting staff's workload

We briefly argued above that CO are particularly sensitive to increased workload in the complex realm of R&D contracting due to the many tasks involved. Yet problems arising from overloaded contracting personnel would be a second-order problem if spikes in agency workloads were sporadic and temporary. However, over the past two decades, federal institutions and scholars have expressed concern about an increasing trend in procurement outlay that has not been accompanied by adequate growth in contract personnel.

In 2007, the Acquisition Advisory Panel reported to the Office of Federal Procurement Policy and the US Congress that between 2000 and 2005, total government purchasing volume had increased by nearly 75 percent, from \$219 billion to more than \$380 billion (AAP 2007), while the federal procurement workforce had remained stable over the course of the same period and shrank significantly vis-à-vis the 1990s. The Panel reported a significant mismatch between the demands of the acquisition workforce and the personnel available to meet them. It recommended that an improved human capital planning process be implemented. In 2010, procurement volume reached \$534 billion, and although it declined to \$430 billion in 2015, the upward trajectory has continued over the past five years, reaching \$579 billion in 2019.<sup>4</sup> Federal spending on R&D followed a similar trend, peaking at \$57 billion in 2010 with a subsequent decline to \$38 billion in 2015, followed by another increase to \$52 billion in 2020. As a result, several agencies still lament an acquisition workforce shortage. A recent Government Accountability Office (GAO) report highlighted that, although the DoD made important changes to its workforce planning to address the AAP 2007's recommendation and increased its procurement workforce by 24 percent between 2006 and 2016, in 2017, it still fell far short of its workforce growth goal, particularly in areas such as contracting and auditing (GAO 2017). In 2017, the GAO highlighted persistent problems with the acquisition workforce in its High-Risk list and emphasized that DoD agencies still faced challenges maintaining sufficient staffing levels and overseeing their acquisition workforce.<sup>5</sup> In 2020, the Federal Acquisition Institute claimed a shortage of qualified contracting professionals in its professional field brochure.<sup>6</sup>

The increase in the federal contracting budget and institutional concerns are not the only in-

---

<sup>4</sup>Source: [www.usaspending.gov](http://www.usaspending.gov).

<sup>5</sup>In 1990, the GAO began a program to report on government operations identified as *high risk*. The list is used to identify and address serious vulnerabilities in areas where significant resources are expended and critical services are provided to the public. See <https://www.gao.gov/highrisk/overview>.

<sup>6</sup>Source: <https://www.fai.gov/sites/default/files/1102-Career-Field-Brochure.pdf>.

dications that contracting personnel are struggling with problems stemming from excessive workloads. For example, studies based on survey data confirm that federal procurement personnel cite under-staffing as one of the primary problems within their work unit (Rau and Stammersky 2009). Specifically, Rendon et al. (2012) show that in two of the DoD’s largest sub-agency, the Department of the Army and, notably for this paper, the Department of the Air Force, the vast majority of procurement personnel responsible for service acquisition disagree that the size of the procurement workforce is adequate to meet objectives and also disagree that vacant positions are adequately filled.

## 2.4 Prima facie evidence and research design

As discussed above, peaks in workload can strongly influence the discretionary decisions of COs. The primary empirical objective of this paper is to assess whether a variation in workload within contracting offices procuring R&D affects contract performance.<sup>7</sup> Properly identifying such an effect presents us with several empirical challenges. First, we need to define a satisfactory method to measure both the workload of the procurement office and the performance of a particular R&D contract. Second, we need to consider potential factors that might challenge the causal interpretation of a negative relationship between workload and contract performance.

To address the measurement issues, we build on the recent literature. Warren (2014) discusses the complexities associated with constructing a robust measure of workload. In the paper, the author opts for a relatively agnostic approach and uses the size of contracting personnel (i.e., including but not limited to COs) in a federal agency while controlling for the number of contracts (*purchases*, from now on) as an inverse proxy for office workload. We follow a similar approach and use the COs actively working in the R&D process for a contracting office in a given fiscal year. As in Warren (2014), we control for the number of purchases; in addition, we use the total annual amount of obligations (*budget*, from now on) of the contracting office. Another difference is that our focus is exclusively on R&D contracting. Therefore, our workload measure counts only COs responsible for R&D procurement. In the same vein, we control only for the purchases and budget in a given office for the procured R&D activities. In contrast to Warren (2014), our data and setting allow us to pinpoint the exact procurement unit of a federal agency. Other details about the office and the definition of the CO are presented in Section 3.1.

Regarding measuring R&D contract performance, we follow Decarolis et al. (2021) and Bruce et al. (2019). As reported in FAR Part 35, most R&D contracts are focused on goals for which work or methods cannot be precisely described in advance and for which it is not easy to assess the probabilities of success ex-ante. Because of the uncertainty that characterizes this process and the resulting high degree of incompleteness, it is not easy to estimate costs accurately. For this reason, FAR recommends the use of cost-reimbursement contracts for R&D procurements rather than fixed-price contracts, which are typically preferred for off-the-shelf procurements and more standardized services. Instead, the primary goal of R&D contracts is to advance scientific and technical knowledge and apply that knowledge to the extent necessary to achieve agency

---

<sup>7</sup>A contracting office is an entity that executes procurement transactions—goods, services, constructions, and R&D—on behalf of the government. A contracting office belongs only to a subagency, that is, the bureau responsible for the transaction.

goals, not to deliver a product in the most cost- and time-effective manner. Given the unique characteristics associated with R&D contracting, standard measures of contract performance—such as unit price comparisons, cost overruns, time delays, and the number of contract renegotiations—are not well suited to assessing the performance of an R&D contract. As in Bruce et al. (2019) and Decarolis et al. (2021), we consider an R&D contract successful if associated with subsequently patented inventions. Although using patents as a proxy for (relevant) innovation is not without drawbacks (Boldrin and Levine 2013), the issuance of a patent application associated with an R&D contract ensures that the contract has generated new knowledge that can be used to solve a particular technical problem.<sup>8</sup> In addition, Peña et al. (2017) confirm that most DoD research offices themselves use metrics such as patent applications and grants to assess the success of their early-stage research and technology projects.

Once we have defined a valid measure for an office’s workload and one for the outcome of an R&D contract, we need to consider how to identify the impact of the former on the latter. As a first step, we can match trends in R&D contracting personnel in those federal agencies that regularly purchase R&D from the private sector with the dynamics of R&D procurement spending and the number of patents associated with those contracts. Figure (1) shows these trajectories. Like general contracting, there were substantial increases in R&D expenditures (solid line) between 2005 and 2008, while the size of the acquisition forces in these offices (dashed line) remained stable. In 2009-2012, the data show a reversal of trends: the size of R&D contracts consistently declined, while the size of the CO pool steadily increased between 2009 and 2012. Therefore, the workload of COs for R&D contracts seems to have increased until 2008 and decreased ever since. Interestingly, the number of patents associated with contracts awarded by the offices (dotted line) shows a negative correlation with workload: it decreased significantly between 2005 and 2008. It reversed the trend after 2008 and especially since 2010. This could be considered as *prima facie* evidence of a negative impact of COs’ workload on the performance of the R&D contracts they award. However, providing evidence that would bolster a causal interpretation of this relationship requires much more thought.


Ideally, we would look at all federal contracting offices that award contracts for procuring R&D and randomly divide them into different groups. We would then assign a different number of additional COs to each group; finally, we would evaluate whether the contracts awarded by the offices whose CO employment grows (shrinks), i.e., offices with a lower (higher) workload per officer, subsequently experience a higher (lower) likelihood of being associated with a patented invention. In such an ideal experiment, randomization of the treatment would ensure that the relationship between workload and contract performance is causal. While conducting such an experiment is not feasible, US federal agencies produce a wealth of observational data that, if properly used, would allow us to examine the existence of a causal relationship (if any) between COs’ workload and contract performance. Section 3.1 describes the data in detail, but for now, it is sufficient to stress that our data provide enough information to estimate the following linear probability model:

$$Patent_{i,k,t} = c + \beta \text{Log} \# \text{CO}_{k,t} + \gamma X_i + \delta Z_{k,t} + \zeta_k + \zeta_t + \epsilon_{i,k,t}, \quad (1)$$

---

<sup>8</sup>A discussion of the drawbacks of measuring R&D activity via generated patents and how we deal with them is covered in Section 3.1.

Figure 1: Workload and Patent Trends



Notes: Average annual R&D obligated budget (solid line), contracting force (dashed line), and the average number of patents per contract (dotted line). The source is the FPDS database, and all agencies awarding at least one R&D contract in the sample are considered.

where the variable  $Patent_{i,k,t}$  reports whether R&D contract  $i$  awarded by the office  $k$  in the fiscal year  $t$  leads at least to a patented invention and the  $Log \# CO_{k,t}$  reports the  $CO$  employment in log-terms. Let  $X_i$  be a vector of controls for contract characteristics, whereas  $Z_{k,t}$  is a vector that includes control variables at the office-fiscal-year level that is, the budget and purchases. The vectors  $\zeta_k$  and  $\zeta_t$  include a set of office- and fiscal-year fixed effects. The former take into account the unit's time-invariant characteristics (e.g., location, procedures, and work practices). The latter account for the government budget cycle affecting all offices simultaneously, with the resulting time-varying sources of bias. A detailed description of the control variables is provided in Section 3.1.

Although the wealth of information available for each contract and awarding office allows us to control for several potential confounding factors that may affect the relationship between workload and contract performance, observational data still presents some fundamental challenges for identification. First, unlike in the ideal experiment,  $CO$  employment is not randomly decided. In year  $t - 1$ , the program manager of each federal contracting office plans the budget for purchases to be made for the fiscal year  $t$ . Specifically, in January of fiscal year  $t - 1$ , federal agencies update the budget plan for  $t$  and submit it (as part of the DoD budget) to Congress, which approves it before the beginning of the new fiscal year (i.e., October of the same calendar year).<sup>9</sup> Thus, after considering the human resources available at  $t - 1$ , the incoming workload of an office at time  $t$  is known in advance to the contracting office manager. Therefore, the manager can already decide at

<sup>9</sup>The DoD budget is requested along with the budgets of all other departments and constitutes the budget of the US government, which the President submits in early February at  $t - 1$ . Typically, agencies prepare their proposals about 18 months before the budget takes effect. The budget must be passed no later than  $t - 1$  on September 30, just before the start of the new fiscal year. Otherwise, the government will have no budget, it will shut down, and many functions will cease. Then Congress must pass a continuing resolution to temporarily fund the government. Therefore, each office must know its budget before the fiscal year begins. The Department of the Air Force, which is the focus of our empirical analysis, sets a department-level budget years in advance while also conducting detailed budget estimates for each office that plans for the near future. Source: <https://www.usa.gov/budget>.

$t - 1$  to hire additional contracting staff (including COs and other procurement staff) in case of an expected growing workload. If the office budget could perfectly measure workload, this would not pose a critical problem for identification as long as we can control for trends in the office budget. However, measuring workload by simply considering the budget would overestimate the workload of officials who draft contracts for the simplest tasks. This problem is particularly relevant in the context of R&D procurement. The contract size could depend heavily on the fixed costs of the R&D to be performed. Still, these fixed costs do not necessarily correlate with the complexity of the task, a contract characteristic that we cannot measure perfectly in our cross-R&D-category data. Although we can control for other contract-specific characteristics that might partially capture the technical complexity of a task, if the complexity level of the average contract awarded by a given office changes over time, the program manager can adjust the workforce accordingly. If the manager knows at  $t - 1$  that the office will need to award contracts to perform more complex tasks at  $t$  and knows that awarding more contracts for more complex tasks will require more clerk time, the office manager can plan to hire additional COs in year  $t - 1$ . The possibility of anticipating the increase in workload due to the increase in average task complexity of the contract implies a positive correlation between our main explanatory variable,  $\text{Log \# } CO_{k,t}$ , and the omitted factor, i.e., the average complexity of the contracts awarded by an office over time. At the same time, the complexity of the task for which a contract is awarded—which in turn is omitted—could be strongly correlated with our outcome variable. More complex tasks are almost by definition more uncertain, and their probability of success is lower than simpler tasks. In the context of R&D contracting, this means that a more complex research project is more likely to lead to inconclusive and, thus, unpatentable results. As omitted contract complexity is negatively correlated with the variable  $\text{Patent}_{i,k,t}$  and positively correlated with our variable of interest  $\text{Log \# } CO_{k,t}$ , we would expect the estimate of our coefficient of interest  $\beta$  in Equation (1) to be downward biased.

A second unobserved factor that may introduce bias into our estimates involves the quality of the CO assigned to the contract. Even if we could rule out the possibility that the program manager could anticipate at  $t - 1$  the workload and adjust the CO employment to such needs, the manager would still be able to make some adjustments at  $t$  to minimize the impact of an increased workload. For example, even in the event of a sudden and unpredictable jump in workload, the manager could still assign the more complex projects that require extra time and expertise to the higher quality (or more experienced) COs with above-average productivity. Therefore, offices with a greater proportion of high-quality COs would be better able to respond to unforeseen shocks in workload. In such cases, we would then underestimate the potential effect of a shock in workload for offices staffed with more high-quality COs and overestimate the shock for offices with fewer high-quality COs.

To tackle the two threats to identification described above, we implement a two-step strategy. First, we focus on a specific set of offices for which we can obtain reliable information about the identity of the CO who awards a specific contract. In particular, we use data from the FedBizOpps platform (described in detail in the next subsection). The website reports contact information about the CO in charge of a procurement process to facilitate communication between potentially interested contractors and the awarding agency. Unfortunately, for most federal agencies that award R&D contracts, contact information about the CO is sparingly and incompletely reported.

However, for a few agencies, in particular, the Department of the Air Force, the name and contact information of the relevant COs are reported quite systematically. Taking advantage of this, we focus on the Air Force contracting offices that award the vast majority of R&D contracts, and in particular those offices that are part of the AFRL. Then, for R&D contracts awarded by these offices, we can pinpoint the CO responsible for awarding the contract. This fact allows us to rewrite the Equation (1) as

$$Patent_{i,k,t,o} = c + \beta \text{Log} \# \text{CO}_{k,t} + \gamma X_i + \delta Z_{k,t} + \zeta_k + \zeta_t + \zeta_o + \epsilon_{i,k,t,o}, \quad (2)$$

where  $\zeta_o$  is an additional vector of CO fixed effects, allowing us to account for the intrinsic quality—plus other time-invariant idiosyncrasies, such as integrity, creativity, motivation, risk-aptitude, and alignment to agency mission—of the CO awarding the contract.

To address the residual source of endogeneity provided by the omitted complexity, we adopt an IV approach. In our setting, the instrument must identify a shock in the workload that the program manager cannot predict. Warren (2014) considers a similar problem and proposes the number of retiring contracting employees as an IV. The rationale for this focus is grounded in the US civil service retirement system, where eligibility for immediate pension benefits is predicated on reaching a specified number of service years. However, research indicates that *actual* retirement decisions are influenced more by individual circumstances than by societal standards or pension regulations (Asch et al. 2005).<sup>10</sup> Furthermore, there are no specific incentives for federal employees to retire, as they retain the same health insurance post-retirement that they had during their active service (Asch et al. 2005).

At time  $t - 1$ , the head of a given contracting office knows how many employees are eligible for retirement benefits at  $t$ . As both the office’s budget and purchases at  $t$  are established at  $t - 1$ , the manager would anticipate fluctuations in both the workload and the workforce and make hiring decisions accordingly. If many COs are retirement eligible at  $t - 1$ , the manager will likely hire more staff at  $t - 1$  to compensate for future retirement. In addition, hiring at  $t - 1$  is especially important if the size of the likely retirees at  $t$  requires it, as hiring a new CO could be a lengthy process. Despite the goal of 80 days set by the Office of Personnel Management in 2008, federal agencies took an average of 106 days to hire a new employee in 2017 (127 days for the DoD), with little change from 2004 when the average was 103 days (GAO 2019). This pattern is confirmed in our data: the correlation of retirement-eligible at  $t - 1$  with hiring at  $t - 1$  and hiring at  $t$  is 0.72 and 0.73, respectively.

Our approach exploits the plausibly exogenous variation provided by retirement postponements. At  $t$ , actual retirement is realized. The number of actual retirees could be higher or lower than predicted by the program manager and, according to empirical research, managers have little influence on idiosyncratic retirement decisions (Lewis and Pitts 2018). Therefore, we consider the difference between the counts of eligible retirement and actual retirement to be a good candidate measure of an unanticipated workload shock (controlling for budget, purchases, and officer quality). The larger the difference between the number of officers eligible to retire in the service at  $t - 1$  and

<sup>10</sup>Notably, factors such as a spouse’s retirement and health status (Feldman 1994; Gustman and Steinmeier 1993), along with other demographic variables (e.g., race, ethnicity, family structure) and personal idiosyncrasies (e.g., hobbies, health conditions) (Lewis and Pitts 2018), play a significant role. These personal attributes are recognized as having a profound impact on job turnover overall (Black et al. 1990).



the actual retirement in the same office at  $t$ , the larger the number of COs who (unexpectedly) decide to postpone retirement. As managers’ hiring decisions are based on the number of expected retirees, the larger the difference, the larger the short-run positive shock to the number of COs active in an office at  $t$ .<sup>11</sup>

### 3 Data

In this section, we document our data sources and define the variables of interest. Then we describe our working sample: the merging of these datasets, the agency we focus on, and a descriptive analysis of the working sample.

#### 3.1 Data sources and description of variables

The dataset developed for this study combines four data sources for the first time.

**Procurement data** To retrieve contract-specific information, we rely on the Federal Procurement Data System (FPDS), the source of procurement data of the US government used extensively in recent research, including studies by Liebman and Mahoney (2017), Warren (2014), Kang and Miller (2021), Giuffrida and Rovigatti (2022), Decarolis et al. (2020). Since the fiscal year 2000, federal agencies have been required to complete procurement action reports, which in turn feed into the FPDS. The FPDS covers all federal contracting agency transactions related to an award above the federal micro-purchase threshold.<sup>12</sup> Like Decarolis et al. (2021), we focus only on the pool of R&D contracts from FPDS. The R&D code (and stage) specified for each award comes from the FPDS variable “Product or Service Code.”<sup>13</sup>

Moreover, most of the other information we use in our empirical analysis to build covariates or fixed effects comes from FPDS. We observe the *Fiscal Year* of the project award; the expected cost at the award stage (*Award Amount*) and the final cost, computed as the cumulative sum of the *Award Amount* with all subsequent price renegotiations; the expected and actual duration of a project (*Expected Duration* and *Final Duration*, respectively). *Last Week* identifies a project if

<sup>11</sup>In Section C.1, we show that our results hold when the actual retirement counts among the contracting staff are used as an IV as in Warren (2014).

<sup>12</sup>The value was \$3,000 during the period under analysis. In 2015, it was revised to \$3,500; in 2020, it was revised again and increased from \$3,500 to \$10,000. The amounts for public R&D contracts are typically very high, and we can confidently state that we observe the universe of these projects over the period under analysis.

<sup>13</sup>The variable consists of two alphabetic and two numeric figures. The first figure is always the letter “A” to identify R&D; the second figure is the letter “A” through “Z” to identify the major category of R&D; the third figure is numeric, 1 through 9, to identify a subdivision of the major category of R&D. The categorical variable *R&D Category* is defined according to the combination of the first three figures. The fourth figure is numeric, 1 to 7, to identify the corresponding level of R&D with: (1) Basic Research; (2) Applied Research and Exploratory Development; (3) Advanced Development; (4) Engineering Development; (5) Operational Systems Development; (6) Management and Support; (7) Commercialization. The categorical variable *R&D Stage* is generated accordingly. The R&D usually includes the first six categories. According to the FPDS Product or Service Code Manual, the first stage (i.e., basic research) includes all scientific endeavors and experiments aimed at expanding the body of knowledge and forms part of the basis for subsequent applied research and exploratory and advanced development in the various disciplines, as well as new or improved functional capabilities. The second stage includes all efforts toward solving specific problems except major development projects. Advanced development includes all efforts directed at projects that have transitioned to hardware development for testing, for example. The primary outcome of this type of effort is proof of the design concept and/or prototype.

it is awarded in the last week of the fiscal year (i.e., the last seven days of September).<sup>14</sup>

**Patent data** The 3PFL Database of Federally Funded Patents (3PFL) database, developed by de Rassenfosse et al. (2019), gathers USPTO patents granted between 2005 and 2018 directly related to federal contracts. More specifically, we use the information contained in the 3PFL database to construct our performance measures for our sample of R&D contracts with the extensive margin of patents, i.e., *Patent*, which is a dummy indicating that the project is associated with at least one registered patent. We also observe the number of patented inventions associated with a given federal R&D contract.

Two main concerns might cast doubt on the suitability of patents as a proxy for the innovative output of an R&D contract in our context. First, a contractor might favor secrecy over patenting to protect its invention. However, FAR § 27.3 states that a contractor must timely file a patent application and disclose it to the government to retain title to an invention made under a government contract. If the contractor fails to do so, it risks losing the title to the invention, as the government has the right to file a patent application on its behalf. Thus, there are strong incentives for the contractor to file a patent application when an invention materializes. Second, the government itself may recommend the contractors keep the invention secret in the interest of national security. In such cases, even if the contractor has duly filed a patent application, the Patent Office imposes a secrecy order that halts (at least temporarily) the patent prosecution process.<sup>15</sup> Nevertheless, as discussed in de Rassenfosse et al. (2020), the actual number of secrecy orders issued each year is quite small and only a limited number of these orders appear to target the output of federal R&D contracts.

**Solicitation data** If the CO determines that the appropriate method for procuring the goods or services is a contract, and the expected value is greater than \$25,000, then the contracting authority is required by the FAR to solicit a request for proposal on the Federal Business Opportunities (FedBizOpps) platform.<sup>16</sup> The FedBizOpps can be thought of as the government’s call for tenders’ point-of-entry, and its purpose is to collect, maintain, and disseminate information to the public about federal solicitations. System information is used to administer and manage access by federal buyers, maintain lists of interested vendors, and notify vendors of federal solicitations of business interest. Government contractors use FedBizOpps as a search engine to find immediate solicitations or bid opportunities as well as archived records.

A subset of solicitations on FedBizOpps reports an additional piece of information: the identity of the bureaucrat responsible for the solicitation process, i.e., the CO (and the associated contract identifier once awarded and tracked by FPDS). This point of contact is located at the bottom of the solicitation documents and includes the first and last name, title, phone number, and email address of the CO. For R&D procurement solicitations, this information is particularly rich for the activity of the AFRL, and we refer the reader to Section 3.2 for more details on the sample

---

<sup>14</sup>This control variable is similar to Liebman and Mahoney (2017), who highlight how the federal budget expiring at the end of the fiscal year creates incentives for government buyers to rush to spend resources on low-quality projects.

<sup>15</sup>The Invention Secrecy Act of 1951 governs this process.

<sup>16</sup>FBO.gov moved to the beta version SAM.gov in November 2019—after our data collection—and is now known as Contract Opportunities. See also Carril (2021) for a discussion of the FedBizOpps’ reporting threshold.

of contracting offices under study. We use this valuable information to calculate the total number of COs actively working on the procurement of R&D in AFRL purchasing units in a given fiscal year. We have already defined this variable as  $\# COs$  and referred to it as CO employment. As discussed in Section 3.2, this variable is the main explanatory variable for the project innovation outcome under study. In addition, we use the information about the identity of the CO retrieved from FedBizOpps to control for the quality of the officer responsible for the award of a specific contract through CO fixed effects. Finally, we define *Specialist* as a dummy variable indicating whether the CO is assisted by a contract specialist for a particular R&D process.<sup>17</sup>

**Office personnel data** The Office of Personnel Management—an independent federal agency that acts as the central human resources department for the executive branch—collects, maintains, and publishes data on approximately 96 percent of civilian federal employees. These data are published through the federal Human Resource database (FedScope), which is the most comprehensive source available on the size and scope of the US federal workforce.<sup>18</sup> Fedscope is the fourth data source we use. The data are divided into five subject categories (called “cubes”), of which we consider only the Employment cube and the Separations cube. The Employment cube contains various demographic characteristics and information about appointments and assignments, such as length of service, job category, pay grade, pay level, type of appointment, work schedule, and location of each employee. The Separations cube contains all separation events (inflows and outflows), that is, employees who are transferred to other offices or agencies, resigned voluntarily, retired, experienced a reduction in force, terminated, or died during employment. In both cubes, we focus on GS-1102 employees, which is the government’s job classification series for contracting and acquisition personnel, including COs plus support positions. We will label the sum of GS-1102 employees as  $\# GS-1102$ . This is an alternative and less conservative measure of contracting employment—used by (Warren 2014)—we will use to check the robustness of our results.

Importantly, the FedScope dataset allows us to build our IV as outlined in Section 2.4. The retirement eligibility of federal civilian employees is determined by age and the number of years of creditable service. People who have reached the minimum retirement age qualify for immediate retirement benefits provided they have at least 10 years of creditable service. The minimum retirement age at the beginning of the period we consider (i.e., 2005) was 55 years and 6 months, and it was 56 at the end of the period (i.e., 2012).<sup>19</sup> To identify the number of retirement-eligible COs in a given office—which we define as *Retirement Eligible*—we exploit the information about the age group and the years of creditable service of the employees in the GS-1102 category as reported in the Fedscope database employment cube.<sup>20</sup> Instead, *Retirement Actual* indicates the

---

<sup>17</sup>COs and contract specialists fulfill different roles. A CO is an individual who can bind the US government to a contract and has signature authority as a government contract agent. The contract specialist is a lower-grade contracting bureaucrat: they act as consultants and assist the CO in planning acquisitions. Only the CO is authorized to sign the contract. A specialist is not always necessary—they appear in 81 percent of contracts in our sample, see Table (A2)—while a CO is always required for a purchasing process to be initiated.

<sup>18</sup>The database has exclusions that affect, for example, some national security and intelligence agencies and the US Postal Service. The data are already used by Decarolis et al. (2020), who also provide a detailed description of FedScope.

<sup>19</sup>For details, see: <https://www.opm.gov/retirement-services/fers-information/eligibility/>.

<sup>20</sup>We count the GS-1102 that qualify for immediate retirement benefits, i.e., contracting employees older than 55 years of age and with more than 10 years of creditable service, or older than 62 years of age and five years of service, or older than 65 years of age. Data are available at: <https://www.fedscope.opm.gov/employment.asp>.

GS-1102 who retire over the course of a given fiscal year. Finally, *Non-retirement* is defined as the difference between lag Retirement Eligible and Retirement Actual and is our IV. Before showing the results, we describe in detail our sample of AFRL’s bureaus and present some relevant descriptive facts about the data used to connect FedScope and the other data sources.

### 3.2 The working sample

**The Air Force Research Lab (AFRL)** We emphasize once more that, in our study, we focus on different contracting offices all belonging to one subagency, namely the AFRL, which is a sub-unit of the Department of the Air Force. This is motivated by the richness of data of this federal subagency. Relying on the AFRL has two additional merits. First, it provides a diverse science and technology portfolio. Second, it is big. According to the FPDS data, procurement of R&D activities at the AFRL equaled \$2.22 billion per year from 2005 to 2012. To benchmark this budget, we report that the SBIR program obligated a similar average yearly amount during the same period (Bhattacharya 2021). Moreover, this amount of spending represents one-third of the total NSF’s budget for FY 2009.<sup>21</sup> In Appendix D, we provide details on the AFRL’s mission and procurement activity.

**Data combination and sample selection** In this paper, we require a metric for workload within the contracting office. FedScope data are available at the subagency level—referred to as “AGYSUB”—but the geographic information in FedScope allows us to determine the location (i.e., state) of each federal employee. Since FedScope information does not allow us to pinpoint offices below the level of the subagency, we must use geographic information to link FedScope to the other data sources.

The merging process is as follows. Our starting point is FPDS R&D data. We start by splitting the raw transaction records—i.e., all transactions between government procurement offices and private suppliers—into two main groups: base contracts records and amendment records. The former refer to the first transaction between a procurement office associated with a given contracting process and a supplier and correspond to our unit of observation for this study, the reported characteristics of which constitute the base agreement information. The latter capture all revisions, modifications, or corrections to the contract. All transactions associated with a contract are identified by a unique procurement instrument identifier (“PIID”) that marks a signed contract and all its future modifications; therefore, we can track the contracts’ entire history from award to completion (or close-out) and link each contract to its revisions. Second, FPDS is combined with 3PFL and FedBizOpps at the contract level. This is straightforward as both 3PFL and FedBizOpps report the contract PIID. Finally, the intermediate dataset is merged with FedScope. As the level of observation of FedScope is the subagency-state-year, the data are to be merged at that level. The nomenclature of FedScope bureaus differs from that of FPDS. Still, we have relied on an external dictionary that maps the variable “Contracting Office Agency ID” in FPDS to the variable “AGYSUB” of FedScope.<sup>22</sup> Following the discussion in the Appendix D, we must limit our focus for a correct dataset matching to the combined contract-level information associated

<sup>21</sup>See [https://www.nsf.gov/news/speeches/bement/09/a1b090514\\_budget.jsp](https://www.nsf.gov/news/speeches/bement/09/a1b090514_budget.jsp).

<sup>22</sup>FedScope releases are monthly. To ensure temporal consistency with FPDS and FedScope, we employ the September snapshot of the FedScope cubes as a reference for the closing fiscal year.

with AFRL awarding agencies FA9453, FA8650, FA8718, and FA8750, which represent 88 and 83 percent, respectively, of the spending and contract counts in the AFRL raw sample.<sup>23</sup>

We further restrict the sample according to the following rules: R&D activities conducted within US borders; award amount greater than \$25,000; expected contract end date before the end of the sample; no SBIR contracts; 2005-2012; R&D preceding the commercialization phase.<sup>24</sup> This ultimately leaves us with a sample of 1,970 R&D contracts, with a total value of \$9.6 billion, 12,020 bids submitted, and 579 unique winners (of which 87 were universities or other higher education institutions). The final sample includes 275 distinct COs, whose associated contracts yielded a total of 522 patents (5 percent of contracts with one patent, 4 percent with two or more).

**Descriptive analysis** Table (A1) shows the details of the R&D activities included in the working sample. Each cell reports the number of contracts for each combination of procurement category and R&D stage and the total number of associated patents in parenthesis. Most contracts and patents are observed for the first three stages of the R&D process, i.e., basic research, applied research (and exploratory development), and advanced development. This cross-tabulation highlights how diverse is the science and technology portfolio of the AFRL, going beyond military R&D and spanning basic research to advanced development stages.

Contract amounts are relatively large and highly skewed: 50 percent of contracts have an award price below \$998,000, while 10 percent of contract spending is on contracts worth more than \$7.4 million. The average award amount is approximately \$4 million, but the average total cost, including all subsequent modifications, averages \$4.9 million. Correspondingly, the average expected and final contract durations, including any delays, are 1,000 and 1,113 days, respectively. The significant cost increase is typical of R&D activity, which is compounded by the cost-plus nature of most contracts in our sample (95 percent). The prevalence of cost-plus contracts in DoD procurement is well documented (Carril and Duggan 2020). It is explained by the DoD’s interest in achieving timely completion of contracts whose cost is highly uncertain at the time of bidding.<sup>25</sup> However, we find that most tendering processes are characterized by full and open competition (81 percent), consistent with the statistics on the entire population (Decarolis et al. 2021). The main characteristics of these contracts—depending on whether they are associated with at least one patent—are presented in Table (A2). R&D contracts that lead to one or more patents are on average larger, last longer, and receive more bids. This is consistent with de Rassenfosse et al.

---

<sup>23</sup>Combining FPDS and FedScope data, we define *Same State* as a binary variable—which we include as a covariate in our empirical model—that signals a job performed in the same state where the AFRL contracting office is located.

<sup>24</sup>The \$25,000 threshold is the lowest contract value associated with a contract publicized in FedBizOpps, as described above. R&D contracts are usually very large and this selection prompts the loss of very few observations. Regarding the exclusion of SBIR contracts, they are intended to assist certain small businesses in conducting innovative activities aimed at their eventual commercialization, not their patentability (Howell 2017; Howell et al. 2021; Bhattacharya 2021). Contracts awarded before the fiscal year 2005 (i.e., October 1, 2004) are very few and of poorer quality, according to Liebman and Mahoney (2017). Those awarded after 2012 (September 30, 2012) are excluded because public R&D activity in our data lasts more than three years on average and, once completed, potentially produces a patent 18 months later on average. Since 3PFL tracks patents registered through 2018, contracts awarded from the fiscal year 2013 onward may not have a patent due to the limited time horizon and not as a poor contract outcome. Finally, contracts in the commercialization phase are excluded from the analysis because they do not consist of an R&D process, only the commercialization of the output. However, their share of the raw transaction data is only approximately 3 percent.

<sup>25</sup>See Bajari and Tadelis (2001) for a detailed study of the trade-off between time and cost to completion induced by contract pricing format.

(2019), who show that the size and duration of a contract are positively associated with the total number of patents associated with the contract.

Table (A3) shows the characteristics of the contracting units. Each office spends an average of \$0.54 billion per fiscal year on 65 different R&D contracts. In the median office, the GS-1102 are 111, 27 of which are COs and 2 have managerial responsibilities.<sup>26</sup> We refer to the latter category as *Top GS-1102*.<sup>27</sup> Retirement Eligible represents 23 percent of the median contracting workforce. Actual retirement counts are low during the period (4 percent of the contracting workforce and 19 percent of the Retirement Eligible). This results in high non-retiree counts, i.e., 21 per office.<sup>28</sup>

Finally, we define contractor-level variables. First, we build a metric for winner size by assigning the firm to a quartile of the empirical distribution of the FPDS variable “Annual Revenues”, reporting the average annual firm revenues in the previous three years. A contractor is labeled in a data-driven fashion as *Small* if associated with the first two quartiles—resulting in the cross-section to turnover lower than \$ 10 million. Second, we define *Inexperienced* as a dummy variable indicating that the winner has been awarded no R&D contract overall (i.e., looking at the entire FPDS dataset) in the three years before. Third, we split suppliers depending on their business nature as retrieved from their string name. Accordingly, we define *University* as a binary variable indicating a higher-education or research institute—in contrast to a private supplier.

## 4 Results

This section presents the results from our IV estimates of workload on contract innovativeness.

### 4.1 Baseline results: workload on patenting

We begin the presentation of our results with Table (1) row (1), which displays the estimates corresponding to the binary version of Equation (2), that is,

$$Pr(Patent_{i,k,t,o} = 1 | \text{Log\#CO}_{k,t}, X_i, Z_{k,t}, \zeta_t, \zeta_o) = \Phi(c + \beta \text{Log\#CO}_{k,t} + \gamma X_i + \delta Z_{k,t} + \zeta_t + \zeta_o), \quad (3)$$

where  $\Phi$  is the cumulative standard normal distribution function. Thus, we estimate the binomial response to our variable of interest  $\beta$  via a probit model.<sup>29</sup> We note that the COs in our panel are grouped in contracting units and never switch. Accordingly, in contrast to Equation (2), we exclude the fixed effects for the office to avoid perfect collinearity with the fixed effects for the CO.

Moving from left to right expands the set of controls included. Column 1 reports the most parsimonious model specification and only includes budget and purchases as covariates controlling for office size. Holding the office’s budget and purchases in a given fiscal year helps us interpret

<sup>26</sup>We refer the reader to the definition of GS-1102 contracting employee in Section 3. Recall that COs are a subset of the GS-1102s, which represent the entire contracting workforce in the office.

<sup>27</sup>GS-1102 having managerial responsibilities are those having at least pay grade 14, that is, the salary scale used to set salaries for most government employees according to the General Schedule. Pay grade 14 is reserved for top positions such as supervisors, high-level technical specialists, and top advanced degree holders.

<sup>28</sup>Since the workforce is planned in the previous fiscal year, the fair comparison between Retirement Eligible at  $t - 1$  is with Retirement Actual at  $t$ . Non-retirement follows this scheme and that is the reason why the average of Non-retirement in Table (A3) differs from the mean difference between Retirement Eligible and Retirement Actual.

<sup>29</sup>We do not consider the linear probability model appropriate in our context as the underlying distribution of patents per contract is quite sparse, with many zeros for the contracts without patents, leading to a mean of 0.14 for the binary patent metric, quite far from the interval 0.4-0.6 that would well accommodate both methodologies.

each additional CO colleague as a reduction in the office’s total procurement workload. Column 2 contains fixed effects for the CO and the fiscal year. The former are key to our identification, as discussed in Section 2.4; they also bring fixed effects for the offices. Column 3 includes controls for project value and duration to capture observable and unobserved features of the underlying R&D activity associated with the project scale, which may predict project success and potentially correlate with the office-level workload. In fact, Section 3.2 shows that contracts that yield patents are quite different from those that do not: their award amounts and total cost are approximately three and four times as large, respectively, and their expected and actual duration is almost 50 percent longer. The construction of these variables proceeds as follows. Using the universe of R&D contracts sourced from FPDS, we evaluate within each of the R&D categories cells the empirical distribution of final costs and final duration of contracts. We then assign the final cost and duration of the contract in our sample to the respective decile of the category-specific distribution. We include this classification with fixed effects for both dimensions. To control for another shared layer of unobserved characteristics, column 4 also includes fixed effects for the procurement category and procurement phase of R&D. Controlling for procurement typology is useful for controlling for time-invariant unobserved characteristics related to the probability of generating a patent. In addition, controlling for the stage of R&D activity is particularly important because contracts awarded to conduct basic research may be characterized by a higher degree of uncertainty and have a different probability of being associated with patents than contracts for subsequent stages. Finally, column 5 contains the dummies *Last Week*, *Same State*, and *Specialist*, which capture possible dimensions that may simultaneously correlate with outcome and treatment. Also, we enrich this model specification by adding *Small*, *Inexperienced*, and *University* as contractor-level controls covariates accounting for the different frictions and incentives to channel innovation through patent filing. This is our favored specification.<sup>30</sup> To facilitate the interpretation of the estimates, we report as coefficients the average marginal effects with robust standard errors. In Appendix C, we display how our results are virtually unchanged when standard errors are assumed to be homoscedastic or clustered at different levels.

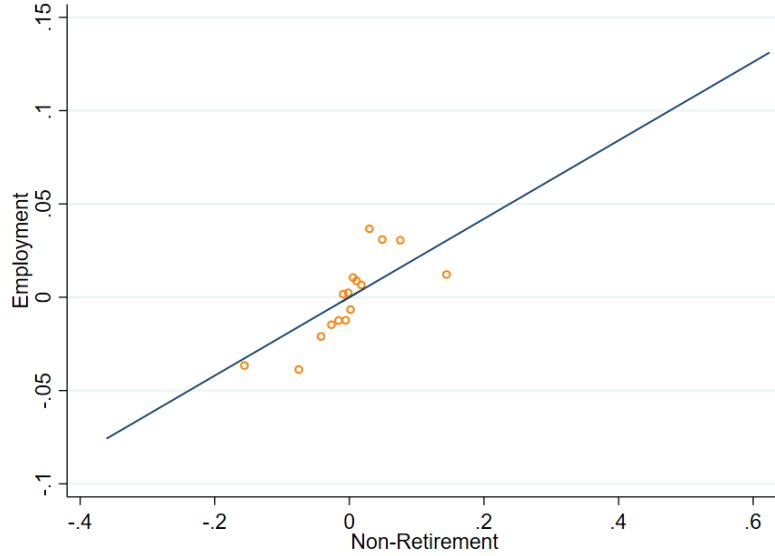
In line with the descriptive evidence, a naive association between workload and patents (column 1) leads to a positive and statistically significant estimate; however, the coefficient loses significance once additional controls are included. In particular, this already happens in column 2, where we add officer and fiscal year fixed effects. Finally, adding more controls increases the magnitude of the estimates but not their precision. Despite the inclusion of these controls, the problem of potential downward bias in the estimates of the structural probit workload remains, as discussed in Section 2.4. To address these concerns, we implement an IV strategy based on non-retirement as the instrument.

**Validity of the IV** For this approach to be valid, we are to satisfy two conditions. First, the instrument must cause the variation in the treatment variable. Second, the instrument must not affect the outcome variable directly but only indirectly through the treatment variable.

---

<sup>30</sup>Please note that the working sample amounts to 1,173 observations instead of the 1,970 presented in Table (A2). This sample reduction is mostly due to the CO fixed effects predicting success or failure perfectly. For comparability, all model specifications are executed on the working sample. Table (C5) in Appendix C displays the robustness of our IV probit’s results to using a 2SLS using the full sample of contracts.

Figure 2: Visual Representation of the First Stage



*Notes:* Graphical representation of the relationships between # COs and Non-retirement. The variables are residualized, including as controls those from column 5 of Table (A5). Each graph is a binned scatterplot. This means that each point represents the mean statistic of the residualized IV and the residualized endogenous variable inside each bin. The selected number of bins is 122, which optimizes the (asymptotic) integrated mean-squared error following Cattaneo et al. (2019).

As for the strength of the instrument, Table (A5) in Appendix A shows the first-stage results. Non-retirement is expressed in log terms and enters with a positive and significant term. Sudden non-separations trigger a positive employment effect, as expected. In numbers, a 1 percentage point increase in non-retirement induces a 0.2 percentage point increase in CO employment. The elasticity of the effect is less than one, likely because we regress CO employment on the non-retirement counts for all contracting employees—not only the non-retirement for COs.<sup>31</sup> Standard statistical tests of the performance of our instrument—reported at the bottom of Table (A5)—reject weak- and under-identification and advocate a strong first-stage relation. Figure (2) provides visual evidence inconsistent with the null hypothesis in which Non-retirement does not affect # COs. The variables are residualized by including the variables from column 5 of Table (1) as controls and grouping them in binner scatterplots as in Cattaneo et al. (2019). Specifically, each dot represents the residualized IV’s mean statistic and the residualized endogenous variable within each bin. This graphical evidence further stresses the positive effect of our IV on project innovativeness.

In terms of the exclusion restriction of our instrument, the sudden non-separation refers to an existing experienced employee who has already reached full productivity and not to a newly hired officer who has not yet reached full productivity. As a result, we expect this exogenous labor surplus as to be good for the R&D office outcomes, but only through a variation in workload. This is confirmed by the reduced-form relationship between the patent dummy and the instrument, as the coefficients on the instrument tend to enter with a positive and significant effect on our outcome variable (see Table A6 in Appendix A). More specifically regarding the exclusion restriction, we

<sup>31</sup>We emphasize again that the instrument is constructed using FedScope data and the GS-1102 classification, which does not distinguish between COs and other procurement bureaucrats. Instead, the endogenous metric for employment is constructed using FedBizOpps records that assign a responsible CO to each contract. However, we control for project-level assistance from a contract specialist in our base specification.



need to consider that non-separation could determine a surplus of skills in the form of knowledge and timely managerial decisions, which can also positively impact the quality of work. Once eligible for retirement, the CO could postpone the decision to retire for some time. As long as the primary determinant of retiring now versus later is idiosyncratic and depends on personal circumstances, in addition to the unobserved office- and year-level circumstances shared by the other employees, we can include the fixed effects, and the instrument will be valid.

Thus, the validity of the instrument depends on the unobserved office features being as good as random. Of course, time-varying office-level characteristics may also influence retirement decisions. If the decision to postpone retirement changes due to office tasks and characteristics changes, the exogeneity assumption would not be satisfied, and the instrument would not be valid. We use variation in the set of employees eligible for retirement and need to test its orthogonality to their workplace features. A crippling condition for us would be that workplace characteristics affect the individual decision to postpone retirement. This would create a potential reverse causality problem, particularly when large changes in workload somehow induce people to stay at work even though they are eligible for retirement. To test this, we collapse the data at the office-year level and run an auxiliary regression analysis—presented in Table (A4) in Appendix A—to detect possible observable determinants of our instruments and provide evidence for our exogeneity argument. Based on how we construct the instrument, we find that non-retirement log-counts are mechanically associated with contracting employment metrics (i.e., # GS-1102 and # Top GS-1102). There is no clear pattern of association between non-retirement counts and any of the other potential office-level predictors we include through our data and that appear in Table (A4): none reach statistical significance across model specifications. Some unobserved change in a qualitative factor of contracts may still drive retirement, undermining identification—in addition to scale variables—but we cannot detect much from observable factors. A further discussion on the exclusion restriction of our IV is presented in Appendix B.

**IV estimates** We can now turn to the presentation of the second-stage relationship between patent and workload. The structural relationship from the probit model depicted by Table (1) shows that the estimated effects of decreasing workload in the patentability of R&D contracts are positive but insignificant. This result would suggest that additional CO colleagues in the office do not affect the innovativeness of their purchases. Offices with more COs are equally likely to generate patents, given the same budget and purchases. However, the IV probit results shown in Table (1) row (2) suggest that the structural probit results are misleading when exogenous changes in the number of COs are considered. The set of controls is identical in all columns and is the same as those of Table (1). In our sample, an additional CO in the office during the award year corresponds to a 3 percent increase in average CO employment. According to the baseline IV probit estimates from column 5, this leads to an increase in the probability of generating patents by about 2.5 percentage points. The effect corresponds to approximately 28 percent of the average project patentability in the sample.<sup>32</sup> Compared with the probit estimates, the magnitude of the

<sup>32</sup>From Table (A2), the baseline probability of patenting is 0.089. The interpretation of the baseline result of our probit model is as follows: A one-unit increase in the log number of COs is associated with a 142 percentage point increase in the probability of patent equal to 1 (or equivalently, an increase in the expected probability of  $0.089 + 1.42 = 1.509$ ). Such an implausible number results from the massive increase in the underlying predictor. The CO employment variable itself, # COs, being natural-log transformed, is multiplied by approximately 2.718, which

IV probit estimates is more than one order of magnitude larger and exceeds the upperbound of the 95 percent confidence interval of the structural relationship. Moreover, the Wald chi-squared test of exogeneity of the IV is rejected at the 95% level, highlighting the endogeneity in the structural model, the inappropriateness of the regular probit regression, and the need for the IV Probit.<sup>33</sup> For a detailed discussion of the robustness of our results to the definition of workload and the IV as well as additional endogeneity concerns, we refer the reader to Appendix A.

**Economic interpretation of the magnitudes** To provide a more transparent interpretation of the estimates, we can consider what would happen if we used them to infer the effect of raising the workload of all offices to the level of the office with the largest workload in our sample. We first collapse our information at the office-year level. Then, we regress employment on budget and purchases. Finally, we rank the office year in terms of residualized employment—interpretable as an inverse workload proxy. In our data, the office-year pair with the highest workload (i.e., lowest residualized employment) is the purchasing unit Rome Laboratory in 2012. If we bring all office-year pairs up to its workload level, this implies a reduction in the number of patents by 13 percentage points on average per contract (i.e., patenting about 50 percent less likely), or about 33 total across all contracts in the dataset every year. This is an economically large effect. For instance, Breitinger et al. (2020) estimate the average long-term effect of one technological patent to an additional \$0.084 of US per capita income, corresponding to an increase in the GDP (constant 2010) of about \$25 million. Removing the excessive workload in key R&D contracting offices can accordingly have a large impact on unleashing growth potential.

All in all, these results suggest that the innovativeness of procurement outlay by offices with labor shortages may be significantly limited. In particular, we show that the performance of the average R&D contract could decline if an increase in the budget allocated to innovative purchases is not matched by a corresponding increase in the budget for contracting personnel. Thus, even in the context of growing R&D budgets, the capacity of federal agencies to match pressing technical needs with innovative solutions could be impaired if procurement departments are not adequately staffed. The existence of this type of bottleneck is particularly troubling when we compare the annual gross salary of a typical CO, which is between \$55,756 and \$72,487—according to the General Schedule pay-scale<sup>34</sup>—to the size of the median R&D contract awarded by a DoD research laboratory—such

---

corresponds to an approximate 172 percent increase in average CO employment. To make the interpretation more tractable and realistic, we may want to see the impact on project innovation of adding one additional CO. To do this, we need only divide the marginal effect of 142 percentage points by one log-unit of CO employment. Since the average number of COs employed in our sample before the log transformation is 32.11, the marginal increase by one log unit of employment equals  $32.11 \times (2.718 - 1)$ , that is, an absolute increase of about 55 employees. Proportionally, one additional CO in the office—implying an average increase of  $1/32.11 \approx 3.11$  percent of # COs—corresponds to an increase in the probability of patenting of  $1.42/55.16 \approx 2.5$  percentage points.

<sup>33</sup>We recognize that the supply side—firms, universities, or research institutes—matters for explaining the variability of the R&D outcome through the idiosyncratic tendency toward secrecy and patenting activity, which is not captured by contractor-level controls. In an auxiliary exercise, we include fixed effects for the supplier and remove contractor-level controls. The results are qualitatively and quantitatively in line, but the effect is not significant at conventional levels due to the reduced sample and predictive power of the model. We circumvent this issue with a two-step strategy. We run an OLS regression of the patent outcome against supplier fixed effects. We then fit the residuals to a 2SLS as an alternative linear outcome. The results are comparable to our baseline estimates and statistically indistinguishable. Further, in Section 5, we discuss how contractor characteristics interplay with our results in a battery of sample splits.

<sup>34</sup>Source: <https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/2020/general-schedule/>

Table 1: Baseline Results

	$\mathbb{1}(\# \text{ Patents} > 0)$				
	(1)	(2)	(3)	(4)	(5)
(1) Log-# COs (Probit)	0.16 (0.048)	0.049 (0.11)	0.085 (0.11)	0.087 (0.12)	0.092 (0.12)
(2) Log-# COs (IV Probit)	0.14 (0.071)	1.14 (0.64)	1.15 (0.60)	1.37 (0.61)	1.42 (0.61)
Wald $\chi^2$	0.19	3.39	3.72	5.18	5.57
CO FEs	No	Yes	Yes	Yes	Yes
Fiscal Year FE	No	Yes	Yes	Yes	Yes
Cost and Duration FEs	No	No	Yes	Yes	Yes
R&D Category FEs	No	No	No	Yes	Yes
R&D Stage FEs	No	No	No	Yes	Yes
Controls	No	No	No	No	Yes
N	1173	1173	1173	1173	1173

*Notes:* Coefficients report average marginal effects of the probit and the IV probit regression. We report the Wald chi-squared test of exogeneity of the IV. In addition to the reported fixed effects, controls for budget and purchases in the unit-year are included in all specifications. Controls include indicators for the last week of the fiscal year, seller and buyer located in the same state, assistance from a contract specialist, and binary indicators of seller size, experience, and higher-education or research institute. Robust standard errors are in parentheses.

as the Air Force Research Lab—which is approximately \$1 million over 3 years. Hiring COs results in being extremely cost-effective in terms of publicly induced patenting. For instance, Azoulay et al. (2019) reckon a net increase of 2.3 patents induced by a \$10 million boost in NIH funding.

## 4.2 Other results: follow-on contracts as a further innovation metric

As reported in the previous section, the CO workload appears to have a strong negative impact on the performance of R&D contracts awarded by the AFRL. As discussed, we have so far measured contract performance based on whether a contract leads to patented inventions or not. However, one might well argue that the existence of a patent constitutes an imperfect measure of contract performance. Indeed, the limitations of using patent data as a proxy for innovation are widely acknowledged in the extant literature. First and foremost, not every valuable innovation is patented and not everything that is patent-protected is valuable. Second, a company’s propensity to patent may heavily depend on the industry in which it operates or on the IP strategy it adopts to appropriate the returns to its R&D investment. To mitigate these concerns, we develop an alternative measure of contract performance based on repeated contracting (Che et al. 2021). The main idea behind this alternative outcome is rather simple as we expect firms able to fulfill the goals of an R&D contract to have a higher likelihood of winning follow-on contracts with the same agency than companies that fail to do so.

We define a follow-on contract as a non-competitive procurement placed with the incumbent contractor for the continued development or production of a major system or highly specialized equipment, or as the provision of highly specialized services (FAR § 6.302). A typical example of such a situation would be the case of a successful R&D contract identifying a solution to a

specific problem posed by a DoD Research lab, which leads to a contract for the development or the production of the actual product for the same agency. The latter contract is not open for competition on the grounds that there is a unique source that has the required knowledge to fulfill the agency requirement. Unfortunately, FPDS data do not allow us to unambiguously pinpoint a link between an initial R&D contract and a potential follow-on contract. The main advantage of using patents as an outcome comes precisely from the fact that the invention disclosure required by the acquisition regulation provides an explicit link between the patented invention and the contract that is connected to it. Nevertheless, FPDS data allows us to check whether a procurement contract is awarded when deviating from the full and open competition procedure and, if so, on what basis. Leveraging this feature of the data, we create the variable  $follow - on_i$  as taking the value 1 if the company involved in the procurement of our focal R&D contract  $i$  wins a sole source, non-compete contract with the Department of Air Force in the year of completion of the focal contract  $i$ , and 0 otherwise. In addition, we consider an alternative time window from the contract completion date and include contracts awarded up to one year after the completion date of the focal contract to construct our follow-on variable. We also exploit the richness of the FPDS data and distinguish follow-on contracts awarded for the performance of R&D work from those awarded for the performance of other types of services or products supplied.

Panel a of Table (A7) reports the results of the estimates obtained when using our measure(s) of follow-on contract as the outcome variable(s) and adopting the same estimation strategy as in the focal analysis (column 5 of Table 1). Column 1 and column 2 show that a reduction in the workload of the CO results in a significant increase in the probability of the focal contractor winning a non-compete contract in the year of completion of the focal R&D contract or in the following year. Columns 3-6 display the results when the follow-on variable is split between R&D and non-R&D contracts. The results appear to be entirely driven by non-compete contracts awarded for the performance of R&D work. In short, R&D contracts awarded by CO with a lower workload are substantially more likely to lead to a subsequent R&D contract for the winning contractor.

We interpret this result as evidence that a reduction in the CO's workload leads to better contract performance that in turn leads to a continuation of the R&D work procured via the original contract. To further corroborate this interpretation, we estimate the model by focusing on R&D follow-on contracts exclusively and constructing three additional follow-on variables that account for three separate R&D stages: basic, applied, and developmental R&D. Clearly, if our interpretation is correct and non-compete, successive contracts are indeed a good proxy to measure the performance of our focal contracts, we should expect the results presented above to be driven by non-compete contracts awarded for more advanced R&D stages that build on the one carried out in previous phases. Panel b in Table (A7) reports the results of this exercise and confirms our hypothesis. The positive effect of a reduction in the CO's workload on the likelihood of winning a subsequent non-compete contract is entirely driven by follow-on contracts awarded for developmental R&D. Once again, this result corroborates the interpretation that the lower the workload of the CO handling the assignment of an R&D contract, the better the contract's performance.

## 5 CO workload and contract performance: A first look at the mechanism

Our findings suggest that decreased CO workload enhances R&D contract performance. We present four arguments for which focusing on the pre-award phase (instead of the post-award phase) is necessary to investigate the mechanism for this effect. First, the effect manifests in the award year when the variation of our IV occurs. Second, R&D contracts are multi-year (see Section 3.2), implying that the average share of execution time over awarding time at  $t$  is low. Third, the role of CO is pivotal until the contract signing and decreases during execution (see Section 2.1). Fourth, existing evidence indicates that disruptions in the contracting office level six months pre-award affect contract patent rates significantly more than disruptions six months post-award, hence already during contract execution (Decarolis et al. 2021). Thus, to rationalize the channel for our IV estimates, exploring the variation in tender characteristics is necessary, based on these arguments.

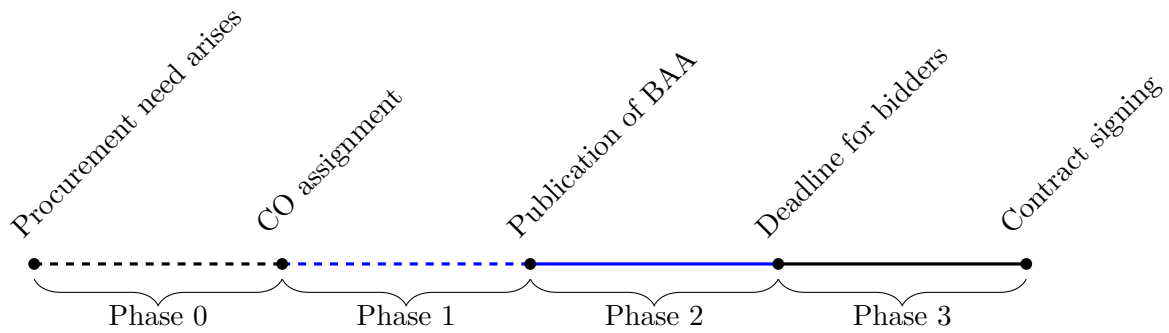
We argue that a marginal decrease in workload reduces time pressure on COs, which allows them to devote more time to pre-award tasks. As discussed in Section 2, R&D procurement is marked by considerable flexibility, with opportunities for significant interaction and knowledge exchange between COs and suppliers. We propose that less time-constrained COs are better able to develop productive relationships with potential contractors. In such cases, the additional guidance that CO can provide is likely to be most beneficial when suppliers have less experience with government contracting.

In what follows, we provide additional empirical evidence supporting the idea that less overloaded COs do in fact experience less tight time constraints (Claim 1) and that such lower time pressure influences contract outcomes (Claim 2) through an improvement of the guidance embedded in contract specifications (Claim 3).

**Higher workload, tighter time constraints** To substantiate Claim 1, it is important to delineate the various stages of the acquisition process, as seen from the CO's perspective. This timeline is graphically depicted in Figure (3). The process commences for the CO upon their assignment as the lead officer for the procurement activity aimed at fulfilling an agency's R&D need (Phase 0). In Phase 1, the assigned CO undertakes market research to gather insights regarding market capabilities and to pinpoint eligible vendors. This stage culminates in the issuance of an official solicitation. As elaborated in Section 2.2, on top of articulating the agency's needs like a request for proposals, a BAA solicits potential solutions and delineates a time-frame for interested vendors to submit their proposals. This information is crucial for the scope of this section and we will focus on BAA as solicitations in the rest of the discussion.

Following the release of the solicitation, Phase 2 unfolds, lasting until the deadline for submitting proposals. In this phase, the CO is tasked not only with ensuring widespread dissemination of the BAA to attract a diverse pool of potential bidders but also with maintaining open communication with potential suppliers to ensure they clearly understand the specifications contained in the BAA. Finally, in Phase 3, the CO collects the proposals from vendors, evaluates them based on the criteria defined in the BAA, and awards the contract to the selected vendor according to the requirements of the BAA. Oftentimes, the evaluation process involves a peer review system or

Figure 3: Timeline of the pre-award phase and the role of the CO



*Notes:* The dotted and solid line segments represent information unavailable and available, respectively, through the BAA. Blue segments indicate phases where CO exercises maximum discretion, while black segments imply minimal to no CO involvement.

a scientific review to guarantee the quality of the selected proposal. We will focus on Phases 1 to 3, as these are the phases in which one CO guides the process.

The time pressure perceived by the CO is likely to influence the duration of each of the three acquisition phases. Specifically, COs under tight time constraints might tend to minimize the duration of each phase, potentially cutting corners. A lower workload may alleviate perceived time pressure, possibly leading to more extended and perhaps more thorough acquisition processes. Yet viewing workload as a proxy for perceived time pressure during Phases 1 to 3 might not be reasonable given variation in CO productivity as a function of workload levels. In fact, a lower workload, while easing time pressure, could enable the CO to focus on fewer contracts, reducing the average time needed to draft a solicitation and award a contract.

However, in Phase 2, the initiative shifts to the potential vendors drafting their proposals. The duration of this phase solely hinges on the response time designated by the CO in the BAA. Various factors may impact the response time, including the urgency of the agency’s need, but the FAR mandates agencies to set a response time that provides potential contractors a reasonable opportunity to respond to a contract action. In the R&D procurement realm, a CO must allot at least a 45-day response time for receipt of proposals from the date of the BAA’s publication (FAR § 35.016). As noted above, during this phase, the CO can continue assisting potential vendors in developing their proposals. Extending Phase 2 beyond the stipulated 45 days affords potential suppliers more time to refine their proposals and, notably, more opportunities to obtain additional guidance.<sup>35</sup> All else being equal, COs facing less stringent time constraints might be in a position to provide longer response times. Consequently, if workload truly affects perceived time pressure during the acquisition process, our inverse measure of workload is expected to correlate positively with the response time set in a BAA.

To validate the hypothesized relation between workload and time pressure, we collect additional data on the pre-award period not directly available from the contract award notice available at FedBizOpp. Specifically, we try to match our contract data with their respective BAA, if any.<sup>36</sup> While identifying this potential link is quite challenging (given no systematic recording of solicitation identifiers in the award notice data), we were able to match 20 percent of the contracts to the

<sup>35</sup>In our data, the average Phase 2 duration is 63 days.

<sup>36</sup>Source: FedBizOpp.

associated BAA and its text.<sup>37</sup> By parsing the text of the BAA, we can extract the publication date of the solicitation from FedBizOpp, and the response date before which the bidders should submit their proposal. In short, we are able to create a variable capturing the duration of Phase 2 for a given contract and to study its relation with our inverse measure of workload.

In Table (A8), we show the result of an auxiliary exercise with the linear version of our model specification using the duration of Phase 2 as an outcome. When we introduce controls for the technical complexity of the BAA text, the CO employment (i.e., lower workload) positively correlates with Phase 2 duration.<sup>38</sup> This observation confirms our hypothesis that workload directly influences time constraints. In particular, a CO burdened with a lower workload experiences less time pressure and may devote more time to assisting suppliers with the formulation of their proposals.

However, we refrain from making causal claims due to the sample selection necessary for merging the contract and BAA data. This data combination markedly restricts the sample when introducing fixed effects to one-tenth of the contracts and two AFRL offices. Also, the lack of power of our instrument impedes instrumenting CO employment and running our IV.

**Tighter time constraints, poorer outcomes** Despite the evidence in our data that a higher workload leads to greater time constraints, linking time dedicated to a specific R&D contract and its outcomes remains an arduous dimension to investigate. Indeed, we cannot reasonably assert that reduced time allocation is sub-optimal given the idiosyncratic nature of each R&D contract. Accordingly, our approach for shedding light on how time pressure affects innovation outcomes (Claim 2) in our context is two-fold.

First, when a CO has less time available to award a contract (i.e., when time constraints are tighter), such a contract should be more heavily impacted by increased workload, as the CO is presumably forced to make decisions faster. This is most likely the case for contracts awarded close to the end of the fiscal year in the US federal procurement system as, by the end of the fiscal year, a contracting office needs to spend its procurement budget (i.e., all scheduled processes need to be started) to avoid losing allocated resources. To link faster decisions to worse outcomes, in Table (A9), we replicate our IV estimates by restricting the sample to awards signed in the last quarter of the fiscal year, i.e., July to September (see Panel a, column 1). Point estimates double in magnitude compared with the baseline, and stronger significant results suggest that when the CO has less time available to award a contract, an additional CO in the office improves R&D performance more strongly. In column 2, we repeat the analysis for the remainder of the year—that is, the first three quarters—and, consistent with our reasoning, the estimates turn insignificant.<sup>39</sup>

Second, based on the results-interpretation exercise discussed in Section 4.1, we rank the office-year pairs in terms of the residualized employment (i.e., inverse workload). We group the office-year

---

<sup>37</sup>We exploit different strategies. In a minority of cases, we are able to use the solicitation identifiers reported in the contract awards notice. For the vast majority of the matched cases, we relied on the textual similarity between the object of the BAA and the awarded contract description provided in FPDS and additional contract- and BAA-level information. BAA records are mainly available from two offices (i.e., FA8650 and FA9453). The merged rate for the contracts from these offices is 36%.

<sup>38</sup>Exploiting the text of the BAA, we also create an additional text-based metric that we use as an additional covariate in this analysis. It measures the number of words from the BAA texts that are not present in the standard English dictionary, which we interpret as gauging the inherent technicality and complexity of the procurement.

<sup>39</sup>The argument can be emphasized when focusing on time windows closer to the end of September. For instance, despite the weak instrument and little power, we find evidence of exceptionally strong results in a small sample of contracts awarded in the last week of the fiscal year.

units in the first tercile of residualized employment distribution and label them “overloaded”. In Table (A9) Panel a, we report the estimates performed on contracts assigned by the high-workload office (column 3) and others (column 4). Estimates hold only in the former sub-sample and turn negative and insignificant for buyers with a lighter workload. We argue that the estimated effect is non-linear in workload and applicable when workload levels are already relatively high.

**Tighter time constraints, poorer guidance** As discussed in Section 2.1, one of the main challenges for a CO in designing a proper acquisition strategy for the purchase of R&D services is to translate a rather abstract idea for each new procurement activity into a language that is clear to the prospective suppliers, even in situations in which a complete understanding of the work is not possible in advance (US Air Force 1967). An increased workload, coupled with tighter time constraints, can potentially hinder the CO’s ability to effectively transform intricate yet ambiguous technical objectives into precise instructions that can effectively steer the actions of potential contractors. Consequently, this could lead to a decline in the quality of guidance offered within solicitation documents and a reduction in opportunities for meaningful interactions between prospective vendors and the CO. The importance of such guidance for contract performance is likely to interact with suppliers’ characteristics. Less (or poorer) guidance, when required the most, may result in poorer performance by the supplier.

To substantiate this argument and provide suggestive evidence for Claim 3, in Panel b of Table (A9), we present the point estimates from two different sample splits. In particular, columns 1 and 2 restrict the sample to small and large contractors, respectively, while in columns 3 and 4 we limit the focus to inexperienced and experienced contractors, respectively (i.e., those with and without previous R&D awards in FPDS), according to our definitions in Section 3.2.<sup>40</sup> The workload effect is stronger when guidance is more necessary from the supplier’s perspective. This applies to contract recipients that are small or without previous experience in federal R&D procurement. By contrast, the results turn insignificant for large and experienced suppliers. Experienced sellers are likely less influenced by contract specifications and more likely to carry out the project irrespective of the quality of guidance from the CO.

## 6 Conclusions

As part of the broader transition from industrial- to knowledge-based economies, the activity of public administrations has evolved to encompass the increasingly important domain of technological governance. In public procurement, this structural change has been compelling governments to procure increasingly complex products and services. In the US, this trend has been associated with an increase in the number and value of procurement contracts. The increasing complexity of procurement has prompted the government to streamline acquisition rules and grant greater discretion to front-line officials (Carril 2021; Calvo et al. 2019; Giuffrida and Rovigatti 2022). However, human resources constitute a major bottleneck, as procurement agencies have faced difficulties recruiting and retaining talent. The result is an overburdened procurement workforce. Indeed,

<sup>40</sup>A zero effect on small and inexperienced dummies as outcomes in the 2SLS highlights how these contractor-level sample splits are orthogonal to both our IV and outcome.



staffing deficits represent a major capacity constraint for effective public spending in general, and investment in R&D in particular.

To quantify the implications of this friction, this paper sought to elucidate the impact of bureaucratic workload on R&D procurement outcomes. We combine several data sources to link tender, contract, patent, and procurement office records to the identity of COs. We focus on contracts awarded by the procurement offices of the AFRL to effectively count the officials actively involved in procuring R&D in a given fiscal year. We use this measure as an inverse indicator of an office’s workload, after controlling for its annual budget and purchases. To overcome a spurious correlation between contract outcomes and CO workload, we implement an IV strategy—combined with CO fixed effects—to identify how the latter impacts the former. The identification comes from unanticipated retirement shifts by COs, which we use as an instrument for workload. Our results are robust to several modifications and indicate that a large increase in patenting at the extensive margin occurs when the same officer is exposed to a declining workload. The results are also robust to outcome definition: A reduction in CO workload triggers a significant increase in the probability of a supplier being awarded a follow-on contract without competition. Decreased workloads lead to less stringent pre-award time constraints, as evidenced by longer proposal submission periods, which are set at the CO’s discretion. This results in better contract outcomes, particularly when time is already limited. Lower time pressure appears to improve the quality of CO guidance during the pre-award phase, disproportionately impacting smaller, less experienced contractors while not affecting the innovation output of larger, seasoned firms.

When CO personnel costs are taken into consideration, our findings suggest that the government could considerably augment the performance of R&D procurement if it dedicated part of its budget for R&D contracts to hiring additional COs. Future research could use our results as a starting point for undertaking further investigation into how the personal traits of procurement staff interact with our proposed causal mechanism. Topics of interest in this regard include the effect exerted by management practices and how COs with high levels of implicit motivation respond to excessive workload.

## References

- AAP (2007): “Report of the Acquisition Advisory Panel to the Office of Federal Procurement Policy and the United States Congress,” *Executive Office of the President*.
- ASCH, B., S. J. HAIDER, AND J. ZISSIMOPOULOS (2005): “Financial Incentives and Retirement: Evidence from Federal Civil Service Workers,” *Journal of Public Economics*, 89, 427–440.
- ASHRAF, N., O. BANDIERA, E. DAVENPORT, AND S. S. LEE (2020): “Losing Prosociality in the Quest for Talent? Sorting, Selection, and Productivity in the Delivery of Public Services,” *American Economic Review*, 110, 1355–94.
- AZOULAY, P., J. S. GRAFF ZIVIN, D. LI, AND B. N. SAMPAT (2019): “Public R&D investments and private-sector patenting: evidence from NIH funding rules,” *The Review of Economic Studies*, 86, 117–152.

- BAJARI, P. AND S. TADELIS (2001): “Incentives versus Transaction Costs: A Theory of Procurement Contracts,” *RAND Journal of Economics*, 32, 387–407.
- BALTRUNAITE, A., C. GIORGIANTONIO, S. MOCETTI, AND T. ORLANDO (2021): “Discretion and Supplier Selection in Public Procurement,” *The Journal of Law, Economics, and Organization*, 37, 134–166.
- BANDIERA, O., A. PRAT, AND T. VALLETTI (2009): “Active and Passive Waste in Government Spending: Evidence from a Policy Experiment,” *The American Economic Review*, 99, 1278–1308.
- BERRY JAEKER, J. A. AND A. L. TUCKER (2017): “Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity,” *Management Science*, 63, 1042–1062.
- BEST, M. C., J. HJORT, AND D. SZAKONYI (2023): “Individuals and Organizations as Sources of State Effectiveness,” *American Economic Review*, 113, 2121–67.
- BHATTACHARYA, V. (2021): “An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program,” *Econometrica*, 89, 2189–2224.
- BLACK, M., R. MOFFITT, AND J. T. WARNER (1990): “The dynamics of job separation: The case of federal employees,” *Journal of Applied Econometrics*, 5, 245–262.
- BLOOM, N., J. VAN REENEN, AND H. WILLIAMS (2019): “A Toolkit of Policies to Promote Innovation,” *Journal of Economic Perspectives*, 33, 163–84.
- BOLDRIN, M. AND D. K. LEVINE (2013): “The case against patents,” *Journal of Economic Perspectives*, 27, 3–22.
- BREITINGER, J., B. DIERKS, AND T. RAUSCH (2020): “World class patents in cutting-edge technologies: The innovation power of East Asia, North America, and Europe,” *mimeo*.
- BRUCE, J. R. AND J. M. DE FIGUEIREDO (2020): “Innovation in the U.S. Government,” Working Paper 27181, National Bureau of Economic Research.
- BRUCE, J. R., J. M. DE FIGUEIREDO, AND B. S. SILVERMAN (2019): “Public Contracting for Private Innovation: Government Capabilities, Decision Rights, and Performance Outcomes,” *Strategic Management Journal*, 40, 533–555.
- BRYAN, K. A. AND H. L. WILLIAMS (2021): “Innovation: Market Failures and Public Policies,” in *Handbook of Industrial Organization*, Elsevier, vol. 5, 281–388.
- BUCCIOL, A., R. CAMBONI, AND P. VALBONESI (2020): “Purchasing Medical Devices: The Role of Buyer Competence and Discretion,” *Journal of Health Economics*, 74.
- CALVO, E., R. CUI, AND J. C. SERPA (2019): “Oversight and efficiency in public projects: A regression discontinuity analysis,” *Management Science*, 65, 5651–5675.
- CALZOLARI, G. AND G. SPAGNOLO (2009): “Relational Contracts and Competitive Screening,” *CEPR Discussion Papers*.

- CARRIL, R. (2021): “Rules versus Discretion in Public Procurement,” Economics Working Papers 1765, Department of Economics and Business, Universitat Pompeu Fabra.
- CARRIL, R. AND M. DUGGAN (2020): “The Impact of Industry Consolidation on Government Procurement: Evidence from Department of Defense Contracting,” *Journal of Public Economics*, 184, 104141.
- CATTANEO, M. D., R. K. CRUMP, M. H. FARRELL, AND Y. FENG (2019): “On binscatter,” *arXiv preprint arXiv:1902.09608*.
- CHE, Y.-K., E. IOSSA, AND P. REY (2021): “Prizes versus contracts as incentives for innovation,” *The Review of Economic Studies*, 88, 2149–2178.
- CHIAPPINELLI, O., L. M. GIUFFRIDA, AND G. SPAGNOLO (2023): “Public procurement as an innovation policy: Where do we stand?” Discussion Paper 17873, CEPR Press.
- DE RASSENFOSSE, G., A. JAFFE, AND E. RAITERI (2019): “The procurement of innovation by the US government,” *PloS one*, 14, e0218927.
- DE RASSENFOSSE, G., G. PELLEGRINO, AND E. RAITERI (2020): “Do Patents Enable Disclosure? Evidence from the Invention Secrecy Act,” Working Papers 9, Chair of Innovation and IP Policy.
- DECAROLIS, F. (2014): “Awarding Price, Contract Performance and Bids Screening: Evidence from Procurement Auctions,” *American Economic Journal: Applied Economics*, 6, 108–132.
- (2018): “Comparing Public Procurement Auctions,” *International Economic Review*, 59, 391–419.
- DECAROLIS, F., G. DE RASSENFOSSE, L. GIUFFRIDA, E. IOSSA, V. MOLLISI, E. RAITERI, AND G. SPAGNOLO (2021): “Buyers’ Role in Innovation Procurement: Evidence from U.S. Military R&D Contracts,” *The Journal of Economics & Management Strategy*, Forthcoming.
- DECAROLIS, F., L. M. GIUFFRIDA, E. IOSSA, V. MOLLISI, AND G. SPAGNOLO (2020): “Bureaucratic Competence and Procurement Outcomes,” *The Journal of Law, Economics, and Organization*, 36, 537–597.
- DECHEZLEPRÊTRE, A., E. EINIÖ, R. MARTIN, K.-T. NGUYEN, AND J. VAN REENEN (Forthcoming): “Do Fiscal Incentives Increase Innovation? An RD Design for R&D,” *American Economic Journal: Economic Policy*.
- EDQUIST, C. (2015): “Innovation-related Public Procurement as a Demand-oriented Innovation Policy Instrument,” Papers in Innovation Studies 2015/28, Lund University, CIRCLE - Center for Innovation, Research and Competences in the Learning Economy, Lund.
- FELDMAN, D. C. (1994): “The decision to retire early: A review and conceptualization,” *Academy of management review*, 19, 285–311.
- FLEMING, L., H. GREENE, G. LI, M. MARX, AND D. YAO (2019): “Government-funded research increasingly fuels innovation,” *Science*, 364, 1139–1141.

- FREEMAN, M., N. SAVVA, AND S. SCHOLTES (2017): “Gatekeepers at work: An empirical analysis of a maternity unit,” *Management Science*, 63, 3147–3167.
- FRIEBEL, G., M. KOSFELD, AND G. THIELMANN (2019): “Trust the Police? Self-Selection of Motivated Agents into the German Police Force,” *American Economic Journal: Microeconomics*, 11, 59–78.
- FURMAN, J. L. AND S. STERN (2011): “Climbing atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research,” *American Economic Review*, 101, 1933–63.
- GAO (2017): “High-Risk series 2017,” Tech. rep., GAO-17-317, US Government Accountability Office.
- (2019): “DOD Increased Use of Human Capital Flexibilities but Could Improve Monitoring,” Tech. rep., GAO-19-509, US Government Accountability Office.
- GERARDINO, M. P., S. LITSCHIG, AND D. POMERANZ (2017): “Can Audits Backfire? Evidence from Public Procurement in Chile,” Working Paper 23978, National Bureau of Economic Research.
- GIUFFRIDA, L. M. AND G. ROVIGATTI (2022): “Supplier selection and contract enforcement: Evidence from performance bonding,” *Journal of Economics & Management Strategy*, 31, 980–1019.
- GUSTMAN, A. L. AND T. L. STEINMEIER (1993): “Pension portability and labor mobility: Evidence from the survey of income and program participation,” *Journal of Public Economics*, 50, 299–323.
- HANNA, R. AND S.-Y. WANG (2017): “Dishonesty and selection into public service: Evidence from India,” *American Economic Journal: Economic Policy*, 9, 262–90.
- HOOD, C. (1995): “The “new public management” in the 1980s: Variations on a theme,” *Accounting, organizations and society*, 20, 93–109.
- HOWELL, S. T. (2017): “Financing Innovation: Evidence from R&D Grants,” *American Economic Review*, 107, 1136–64.
- HOWELL, S. T., J. RATHJE, J. VAN REENEN, AND J. WONG (2021): “Opening up Military Innovation: Causal Effects of ‘Bottom-Up’ Reforms to U.S. Defense Research,” Working Paper 28700, National Bureau of Economic Research.
- JAFFE, A. B. (2000): “The U.S. Patent System in Transition: Policy Innovation and the Innovation Process,” *Research Policy*, 29, 531–557.
- KANG, K. AND R. A. MILLER (2021): “Winning by Default: Why is There So Little Competition in Government Procurement?” *The Review of Economic Studies*, 89, 1495–1556.
- KELMAN, S. (2021): “Bureaucracies as Innovative Organizations,” Working Paper 21-017, HKS Faculty Research.

- KIVIET, J. (2020): “Testing the impossible: Identifying exclusion restrictions,” *Journal of Econometrics*, 218, 294–316.
- LEWIS, G. B. AND D. PITTS (2018): “Deciding to Retire from the Federal Service,” *Review of Public Personnel Administration*, 38, 49–82.
- LIEBMAN, J. B. AND N. MAHONEY (2017): “Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement,” *American Economic Review*, 107, 3510–49.
- MAIT, J. N. (2005): “Making IT Happen: Transforming Military Information Technology,” Tech. rep., National Defense University Washington DC Center for Technology and National.
- MORETTI, E., C. STEINWENDER, AND J. VAN REENEN (2019): “The Intellectual Spoils of War? Defense R&D, Productivity and International Spillovers,” Working Paper 26483, National Bureau of Economic Research.
- NISKANEN, W. A. (1968): “The peculiar economics of bureaucracy,” *The American Economic Review*, 58, 293–305.
- PEÑA, V., S. V. HOWIESON, B. LAL, J. R. BEHRENS, B. L. ZUCKERMAN, M. V. MERRILL, AND J. L. ZHU (2017): *Early Stage Research and Technology at US Federal Government Agencies*, Institute for Defense Analyses.
- PETERS, B. G. (2018): *The politics of bureaucracy: An introduction to comparative public administration*, Routledge.
- RAITERI, E. (2018): “A Time to Nourish? Evaluating the Impact of Public Procurement on Technological Generality through Patent Data,” *Research Policy*, 47, 936–952.
- RAU, C. A. AND P. J. STAMBERSKY (2009): “Management and oversight of services acquisition within the United States Army,” Tech. rep., Naval Postgraduate School, Monterey, CA.
- RENDON, R. G., U. M. APTE, AND A. APTE (2012): “Services acquisition in the DoD: A comparison of management practices in the Army, Navy, and Air Force,” Tech. rep., Defence Acquisition University.
- RUTTAN, V. W. (2006): “Is war necessary for economic growth? Military procurement and technology development,” Tech. rep., University of Minnesota, Department of Applied Economics.
- SANTOLERI, P., A. MINA, A. DI MININ, AND I. MARTELLI (2022): “The Causal Effects of R&D Grants: Evidence from a Regression Discontinuity,” *The Review of Economics and Statistics*, 1–42.
- SHURTZ, I., A. EIZENBERG, A. ALKALAY, AND A. LAHAD (2022): “Physician workload and treatment choice: the case of primary care,” *The RAND Journal of Economics*, 53, 763–791.
- SPENKUCH, J. L., E. TESO, AND G. XU (2023): “Ideology and Performance in Public Organizations,” *Econometrica*, 91, 1171–1203.

US AIR FORCE (1967): “Air Force Research and Development Contracting Officers’ Handbook,” AFSC Pamphlet 70-2, Air Force System Command.

WARREN, P. L. (2014): “Contracting Officer Workload, Incomplete Contracting, and Contractual Terms,” *The RAND Journal of Economics*, 45, 395–421.

WILLIAMSON, O. E. (1975): “Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization,” *University of Illinois at Urbana-Champaign’s Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*.

XU, Y., T. F. TAN, AND S. NETESSINE (2022): “The Impact of Workload on Operational Risk: Evidence from a Commercial Bank,” *Management Science*, 68, 2668–2693.

## A Appendix: Additional Tables

Table A1: Cross-tabulation of Contracts and Patents per R&D Category and Stage

Category	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5	Stage 6	Total
AC1: Defense System (Aircraft)	99 ( 13 )	96 ( 51 )	91 ( 27 )	1 ( 0 )	1 ( 0 )	1 ( 0 )	289 (91 )
AC2: Defense System (Missile/Space Systems)	88 ( 48 )	38 ( 1 )	7 ( 1 )	0 ( 0 )	3 ( 6 )	3 ( 3 )	139 (56 )
AC5: Defense System (Weapons)	0 ( 0 )	1 ( 0 )	1 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	2 ( 0 )
AC6: Defense System (Electronics/Communication Equipment)	110 ( 72 )	257 ( 104 )	53 ( 11 )	14 ( 1 )	5 ( 0 )	6 ( 0 )	445 (188 )
AC9: Defense System (Miscellaneous Hard Goods)	2 ( 0 )	11 ( 1 )	5 ( 0 )	0 ( 0 )	1 ( 0 )	1 ( 0 )	20 (1 )
AD2: Defense Other (Services)	2 ( 0 )	2 ( 4 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	4 (4 )
AD6: Defense Other (Construction)	1 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	1 ( 0 )
AD9: Defense Other (Miscellaneous)	127 ( 6 )	573 ( 93 )	169 ( 59 )	5 ( 0 )	33 ( 2 )	0 ( 0 )	907 (160 )
AE3: Economic Growth (Manufacturing Technology)	0 ( 0 )	0 ( 0 )	2 ( 0 )	0 ( 0 )	38 ( 6 )	0 ( 0 )	40 (2 )
AJ4: General Science/Technology: Engineering	0 ( 0 )	0 ( 0 )	1 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	1 ( 0 )
AJ9: General Science/Technology (Other)	23 ( 2 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	23 (2 )
AN1: Medical (Biomedical)	0 ( 0 )	1 ( 1 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	1 ( 1 )
AR1: Space (Aeronautics/Space Technology)	3 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	0 ( 0 )	3 ( 0 )
AZ1: Other R&D	34 ( 8 )	49 ( 0 )	8 ( 0 )	3 ( 2 )	1 ( 0 )	0 ( 0 )	95 (10 )
Total	488 (149 )	1,028 ( 255 )	337 ( 98 )	23 ( 3 )	82 ( 14 )	11 ( 3 )	1,970 (519 )

*Notes:* Cross-tabulation of the total number of contracts and associated patents (in parenthesis) for each R&D category and stage in our dataset.

Table A2: Summary Statistics: Contract Level

	No Patents			With Patents		
	Mean	Median	St.Dev.	Mean	Median	St.Dev.
Award Amount (\$ ,000)	3432.95	959.83	14728.58	10703.81	3071.42	23041.69
Total Cost (\$ ,000)	4092.92	1042.87	17227.72	13314.69	4009.57	26729.48
Expected Duration (days)	967.09	823.00	632.40	1333.86	1194.00	785.99
Total Duration (days)	1075.47	1003.50	633.12	1495.86	1462.00	786.95
# Patents	.	.	.	2.97	1.00	4.97
Cost Plus (dummy)	0.94	.	0.23	0.96	.	0.20
Fully Completed (dummy)	0.81	.	0.39	0.79	.	0.41
Last Week FY (dummy)	0.09	.	0.29	0.09	.	0.28
Same State (dummy)	0.17	.	0.38	0.16	.	0.37
Specialist (dummy)	0.81	.	0.39	0.81	.	0.39
# Bids	6.70	2.00	15.89	9.48	2.00	25.24
N	1794			176		

Notes: The level of observation is the contract. The share of contracts associated with at least one patent is 8.9 percent.

Table A3: Summary Statistics: Office

	Mean	Median	S.D.
R&D Budget (\$ ,000)	536,270.00	394,194.22	468,614.23
# R&D Contracts	65.67	74.00	47.19
# COs	32.11	26.86	25.55
# GS-1102	292.10	111.00	321.25
# Top GS-1102	2.33	2.08	1.50
Non-retirement	53.83	22.0	51.40
Retirement Eligible	65.17	26.00	62.55
Retirement Actual	10.50	5.00	12.85
N	30		

Notes: The level of observation is the contracting office and fiscal year.



Table A4: Non-retirement Predictors

	Log-(Non-retirement)					
	(1)	(2)	(3)	(4)	(5)	(6)
(mean) University	-0.79 (0.53)	-0.80 (0.92)			-1.04 (0.44)	-0.29 (0.57)
(mean) Small	-0.18 (0.16)	0.50 (1.83)			-0.22 (0.25)	0.97 (0.53)
(mean) Inexperienced	0.029 (0.96)	-1.07 (3.22)			0.89 (0.71)	-1.85 (1.18)
(mean) Specialist	-0.092 (0.43)	-5.95 (0.75)			-0.67 (0.62)	-1.03 (0.68)
(mean) Same State	-0.33 (0.60)	-3.40 (1.75)			0.43 (0.69)	0.38 (0.51)
(mean) Last Week	0.56 (0.63)	0.043 (2.68)			1.02 (0.67)	-1.01 (1.03)
(mean) Log-Budget			-0.11 (0.24)	-0.13 (0.062)	-0.24 (0.34)	0.26 (0.23)
(mean) Log-Purchases			0.0085 (0.082)	0.19 (0.079)	-0.16 (0.12)	-0.086 (0.13)
(mean) # GS-1102			0.63 (0.21)	0.90 (0.042)	0.38 (0.36)	0.79 (0.12)
(mean) # Top GS-1102			-0.024 (0.020)	-0.022 (0.039)	-0.075 (0.044)	-0.081 (0.046)
Office FEs	Yes	No	Yes	No	Yes	No
Fiscal Year FE	No	Yes	No	Yes	No	Yes
R-Squared	0.98	0.78	0.98	0.98	0.99	0.99
N	30	30	30	30	30	30

*Notes:* The table presents two sets of possible predictors of the office-year non-retirement instrument. Columns (1) and (2) include contract and contractor characteristics demeaned at the office-year level. Columns (3) and (4) include office features. Columns (5) and (6) nest the set of covariates. OLS estimates include, alternatively, office and year fixed effects. Robust standard errors are in parentheses.

Table A5: First-stage Regressions

	Log-# COs				
	(1)	(2)	(3)	(4)	(5)
Log-(Non-Retirement)	0.152 (0.00753)	0.203 (0.0513)	0.206 (0.0512)	0.208 (0.0556)	0.210 (0.0566)
Weak Id.	408	16	16	14	15
Under Id.	270	23	24	23	24
CO FEs	No	Yes	Yes	Yes	Yes
Fiscal Year FE	No	Yes	Yes	Yes	Yes
Cost and Duration FEs	No	No	Yes	Yes	Yes
R&D Category FEs	No	No	No	Yes	Yes
R&D Stage FEs	No	No	No	Yes	Yes
Controls	No	No	No	No	Yes
N	1173	1173	1173	1173	1173

*Notes:* We report the Wald F statistic for weak identification (Kleibergen-Paap) and LM test statistic for under identification (Kleibergen-Paap). In addition to the reported fixed effects, controls for budget and purchases in the unit-year are included in all specifications. Controls include indicators for last week of fiscal year awards, seller and buyer located in the same state, assistance from a contract specialist. Robust standard errors are in parentheses.

Table A6: Reduced-form Regressions

	$\mathbf{1}(\# \text{ Patents} > 0)$				
	(1)	(2)	(3)	(4)	(5)
Log-(Non-Retirement)	0.024 (0.011)	0.22 (0.12)	0.23 (0.12)	0.27 (0.13)	0.28 (0.13)
CO FEs	No	Yes	Yes	Yes	Yes
Fiscal Year FE	No	Yes	Yes	Yes	Yes
Cost and Duration FEs	No	No	Yes	Yes	Yes
R&D Category FEs	No	No	No	Yes	Yes
R&D Stage FEs	No	No	No	Yes	Yes
Controls	No	No	No	No	Yes
N	1173	1173	1173	1173	1173

*Notes:* The coefficients report average marginal effects. In addition to the reported fixed effects, controls for budget and purchases in the unit-year are included in all specifications. Controls include indicators for last week of fiscal year, seller and buyer located in the same state, assistance from a contract specialist. Robust standard errors are in parentheses.

Table A7: Follow-on Contracts

<b>Panel a: R&amp;D vs. non-R&amp;D</b>						
1(# Follow-on Contracts > 0)						
	All		R&D		Non-R&D	
	(1)	(2)	(3)	(4)	(5)	(6)
Log-# COs	1.54 (0.62)	1.59 (0.59)	2.16 (0.50)	2.27 (0.45)	0.20 (0.73)	0.38 (0.71)
N	1,663	1,703	1,576	1,634	1,583	1,646

<b>Panel b: R&amp;D Stages</b>						
1(# Follow-on Contracts > 0)						
	R&D - St.1		R&D - St.2		R&D - St.3+	
	(1)	(2)	(3)	(4)	(5)	(6)
Log-# COs	0.57 (0.67)	0.43 (0.71)	0.64 (0.41)	0.73 (0.44)	1.58 (0.57)	2.07 (0.53)
N	1,010	1,056	1,125	1,255	1,509	1,568

*Notes:* Baseline model—column 5, Panel b, Table (1)—is reproduced with different specifications for follow-on (non-compete) contracts (awarded by the Department of Air Force) to the focal R&D contract as alternative binary outcomes. Even columns indicate at least one contract at issue in the year of completion; odd columns indicate at least one contract at issue in the year of completion or in the following year. Panel a: all contracts (1,2); R&D contracts only (3,4); non-R&D contracts only (5,6). Panel b: R&D contracts of stage 1 only (1,2); R&D contracts of stage 2 only (3,4); R&D contracts of stage 3 or higher only (5,6). Coefficients report average marginal effects.

Table A8: Bidding Period

	(Log-Phase 2)	(Log-Phase 2)
Log-# COs	0.90 (1.26)	2.28 (1.30)
N	145	145
BAA Controls	No	Yes

*Notes:* Baseline model—column 5, Table (1)—is reproduced in its non-instrumented version (Panel a) with the length (in log days) of the pre-award phase 2, as defined in Section 5.

Table A9: Mechanisms and sample splits

<b>Panel a: Time Constraints</b>				
	$\mathbf{1}(\# \text{ Patents} > 0)$			
	(1)	(2)	(3)	(4)
	(Q4)	(Q1-3)	(High Workl.)	(Low Workl.)
Log-# COs	2.44 (0.88)	1.26 (1.13)	0.38 (0.24)	-1.36 (1.70)
N	339	584	377	686

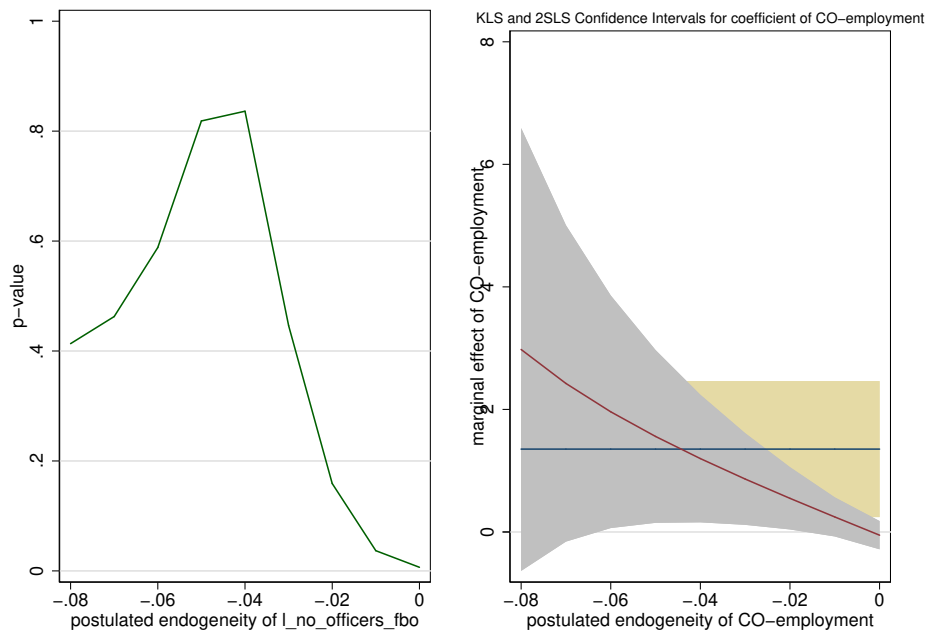
<b>Panel b: Size and Experience</b>				
	$\mathbf{1}(\# \text{ Patents} > 0)$			
	(1)	(2)	(3)	(4)
	(Small Firm)	(Large Firm)	(Inexp. Firm)	(Exp. Firm)
Log-# COs	2.54 (0.54)	-1.98 (1.80)	2.01 (1.01)	0.81 (0.56)
N	466	438	264	1,684

*Notes:* Baseline analysis is replicated: in the sub-sample of the last quarter of the fiscal year (Panel a, column 1) and of the first three quarters (Panel a, column 2), conditioning on overloaded offices (Panel a, column 3), not overloaded (Panel a, column 4); conditioning on small (Panel b, column 1), large (Panel b, column 2), inexperienced (Panel b, column 3), experienced (Panel b, column 4) winners. Coefficients report the average marginal effects of the IV probit regression. The coefficients are estimated via 2SLS in Panel b, columns 3 and 4 due to insufficient power. Robust standard errors are in parentheses.

## B Appendix: Further Discussion on the Exclusion Restriction

To further corroborate our IV strategy, we propose a state-of-the-art exercise hinging on the recent contribution from Kiviet (2020) to seek to actually *test* the exclusion restriction. Kiviet (2020) presents an approach by which, without exploiting any external instruments, general linear coefficient restrictions can be tested in a multiple regression model with an arbitrary number of endogenous regressors. The strategy requires a flexible assumption on the degree of endogeneity of all regressors. This approach allows for generating statistical evidence on the tenability of exclusion restrictions. When this yields an acceptable just-identifying or over-identifying set of instruments, it provides the essential underlying building block for a standard or a series of incremental Sargan–Hansen tests. We follow the Kiviet (2020) approach in our just-identified IV analysis to produce further insights into the tenability of our exclusion restriction hypothesis. The left graph of Figure (B1) shows different values of downward bias in Equation (3) (i.e., negative values of postulated endogeneity), the p-values of the single just-identifying exclusion restriction tests for CO employment. The instrument’s validity seems quite likely when it is close to -0.04, and it holds over most of the negative space. In the right-hand graph, Figure (B1) shows the 2SLS asymptotic 95 percent confidence interval (the yellow area) for  $\beta$ , which is invariant regarding the endogeneity and centered at the 2SLS estimate 1.14 (solid blue line). It also shows the KLS estimator (the solid red line), which varies with the postulated endogeneity, and the KLS asymptotic 95 percent confidence interval (the grey area). The graph also shows that the 95 percent 2SLS confidence interval, which is contingent on the validity of the instruments, conforms in width to a conservative KLS-based interval contingent on the supposition and exogeneity of the instrument. This evidence suggests that, with the likely size and direction of endogeneity, the 2SLS and KLS inference are similar, and the exclusion restriction underlying our IV strategy’s validity is satisfied.

Figure B1: Exclusion Restriction Test as in Kiviet (2020)



*Notes:* Values of single just-identifying exclusion restriction tests and inference of 2SLS and KLS based on non-orthogonality conditions.

## C Appendix: Robustness Analysis

### C.1 Main Robustness Checks

In this appendix, we summarize the findings from the main robustness checks. To simplify the exposition, we present the findings by categorizing them into three groups. First, to address potential concerns about the definition of the workload index, we execute the regression shown in Equation (3) with alternative endogenous employment measures sourced from the FedScope’s workforce data. Second, to assess the soundness of the identification strategy, we explore alternative definitions of the instrument by again using the FedScope dataset. Finally, we use additional variables from the FPDS to test for possible omitted variables that could bias our results. In essence, these additional results further exploit the richness of the dataset. While the overall qualitative results prove robust, these additional findings play an important role in strengthening the quality and depth of the analysis.

The combination of the FPDS and FedBizOpps allows us to identify the COs in charge of the procurement process and count the distinct officers active in a given contracting unit-year. We want to test the robustness of our results against alternative and less conservative specifications of CO employment. By relying on available information from FedScope, we are able to provide an alternative and less conservative count of CO employment. As stressed in Section 3.1, the GS-1102 count includes all COs and other contracting employees involved in the procurement process at different levels of the hierarchy and with different tasks. In Table (C1), we show how the estimates change relative to our baseline from Table 1—reported in column 1—when we change only the endogenous variable. The coefficients from column 2 suggest that replacing our baseline index of CO employment with the total count of GS-1102 has no statistical difference in terms of its effect on the outcome. To capture the effect heterogeneity that arises from GS-1102 having managerial responsibilities, we condition GS-1102 on being “top”, according to our definition. Although qualitatively the same, the magnitude of the effect is one-third of the baseline although more precisely estimated. Econometrically, this follows a stronger first-stage coefficient (i.e., 0.80 instead of 0.27 from Table A5) most likely due to a higher chance of a top contracting employee in the retirement-eligible population. One possible interpretation of a weaker second stage is that top contracting officials are not actively involved in contract administration and the labor supply shock associated with their non-retirement has less impact on outcomes.

Our instrument leverages the unexpected gaps between actual and expected retirement. In the baseline analysis, we construct this variable as the log difference between the two. As we rely on the same underlying variables, we want to test the qualitative stability of our results when we use alternative specifications of IV with a similar interpretation. In Table (C2), we benchmark the results in column 1, where we report the baseline. In column 2, we report the ratio of Retirement Eligible to Retirement Actual. In column 3, we use the logarithm of this ratio. The second-stage results are stable and statistically indistinguishable across all alternative linear or log-linear specifications of the non-retirement counts. Column 4 uses log counts of total retirements as in Warren (2014) as an alternative instrument for the workload. Again, the results are qualitatively and quantitatively stable. Finally, column 5 presents an over-identified IV model aiming at simultaneously addressing two possible concerns on our empirical strategy. On the one hand, our

non-retirement IV can only induce by construction positive shocks to COs from retirement delays. That means that our just-identified IV could only measure the LATE of decreasing workloads, not increasing workloads. Next, a second IV allows testing validity of over-identification restrictions, which would allow corroboration of our arguments on the exclusion restriction of our baseline IV. We employ deaths among the contracting staff as a second IV for 2SLS. We count the (log) number of #GS1102 deaths at  $t$ . Despite including all contracting stuff, deaths are a rare event in our sample: average and median are respectively 0.44 and 0. Only 11/34 office-year pairs present at least one dead, and 4 feature 2 death events. Despite such sparse figures, this second IV has a strong, negative first stage ( $\beta^{fs} = -.06$ ,  $F = 36.27$ ). The second stage estimates are presented in column 5 and are not statistically distinguishable from the baseline, which help generalize our baseline insights to the effects of workload increases in an ATE fashion. Test of over-identifying restrictions is not rejected ( $J = 0.72$ ), not indicating that the instruments may not be valid.

The results of Warren (2014) suggest that the decision to leave contracts less complete may also affect other procurement terms, in particular the extent to which a project is contested at the bidding phase and the pricing structure of the contract offered. Less complete contracts benefit less from the competition, so busier COs use less competitive mechanisms. Cost-plus contracts facilitate the management of contract renegotiations and are therefore preferred by COs in the current and foreseeable high workload. The author shows that an increased workload for COs due to workload spikes leads to fewer complete contracts and, consequently, higher use of noncompetitive and cost-plus agreements. Following these arguments, other dimensions of the design process observed via FPDS may be affected by workload. Although in the R&D realm contracts are highly incomplete by design, some variation at the intensive margin could still be captured by contract pricing (i.e., cost-plus vs. fixed price) and the choice of the officer to make the notice full and open to competition or to exclude some sources.

Another decision that the CO makes in the solicitation is the bidding process. In standard procurement processes, the bidding process usually boils down to a choice between a sealed low-bid auction and a negotiated proposal format. According to FAR, bidding procedures in R&D procurements can vary, and the path chosen depends heavily on the nature of the research features being procured. Whether pricing, competition, and tendering procedures also affect our outcome variable is again an open question for which there is neither empirical evidence nor theoretical modeling. However, we believe it is relevant to test the robustness of our results against the inclusion of these problematic controls that capture contract completeness decided by CO and could bias our baseline results. In column 2 of Table (C3), we add two binary variables to the baseline model specification—in column 1—indicating the cost-plus nature of the contract (as opposed to a fixed price) and open (vs. restricted) competition—whose statistics are displayed in Table (A2). The coefficient is insignificant, while the coefficient on our main variable is only marginally affected. As we are agnostic about the implications of different awarding procedures in the R&D context, we take a data-driven approach in column 3 and include a set of fixed effects for the different categories of procedures in the baseline model specification. Specifically, 81 percent is basic research (FAR § 6.102), 11 percent is negotiated proposals/quotes, whereas the remainder is split between sealed bids, single sourcing (FAR § 13.106), and multiple award fair opportunities (FAR § 16.505). Again, the main coefficient of interest is positive and significant,



with an indistinguishable magnitude compared with the baseline analysis for the same sample. Finally, column 4 includes both dummy variables and fixed effects with similar results. The results prove to be very robust to the inclusion of decision variables that are up to the discretion of the CO, in particular in the R&D contracting realm.

Table C1: Alternative Specifications of Endogenous Employment

	1(# Patents > 0)		
	(1)	(2)	(3)
Log-# COs	1.42 (0.61)		
Log-# GS-1102		1.02 (0.48)	
Log-# Top GS-1102			0.40 (0.090)
N	1173	1173	1173

*Notes:* Baseline results—column 5 of Table (1), reported in column 1—are replicated with alternative measures of contracting employment. # COs is replaced as endogenous variable by # GS-1102 and # Top GS-1102 in columns 2 and 3, respectively. Robust standard errors are in parentheses. \*\*p < .05, \*\*\*p < .01

Table C2: Alternative Instruments

	1(# Patents > 0)				
	(1)	(2)	(3)	(4)	(5)
Log-# COs	1.42 (0.61)	1.87 (0.52)	1.76 (0.48)	0.81 (0.35)	1.24 (0.51)
N	1173	1028	1028	1173	1173

*Notes:* Baseline results—column 5 of Table (1), reported in column 1—are replicated with alternative IVs: Retirement Eligible  $t - 1$  / Retirement Actual  $t$  in Column 2;  $\log(\text{Retirement Eligible } t - 1 / \text{Retirement Actual } t)$  in Column 3;  $\log(\text{Retirement Actual } t)$  in Column 4. Column 5 employs GS1102 deaths  $t$  as a second IV in an over-identified 2SLS model. Coefficients report average marginal effects. Robust standard errors are in parentheses.

Table C3: Inclusion of Endogenous Omitted Controls

	1(# Patents > 0)			
	(1)	(2)	(3)	(4)
Log-# COs	1.42 (0.61)	1.25 (0.56)	1.22 (0.58)	1.12 (0.51)
Cost Plus Pricing		0.031 (0.052)		0.029 (0.052)
Open Competition		0.047 (0.044)		0.051 (0.044)
Solicitation Proc. FEs	No	No	Yes	Yes
N	1173	1126	1172	1125

*Notes:* Baseline results—column 5 of Table (1), reported in column 1—are replicated with the inclusion of additional covariates: Cost Plus and Open Competition dummies in column 2; Solicitation Procedures fixed effects in column 3; Cost Plus and Open Competition dummies plus Solicitation Procedures fixed effects in column 4. Coefficients report average marginal effects. Robust standard errors are in parentheses.

**Additional Robustness Checks** We report additional robustness checks. For convenience, these results are subdivided into two groups depending on whether the robustness analysis involves 1) the specification of the standard errors; 2) the estimation method.

1. *Robustness to the definition of the standard errors.* Baseline results—column 6 of Table (1), reported in column 1—are replicated with different specifications of the standard errors: homoscedastic standard errors in column 2; clusterization at the supplier level in column 3; clusterization at the R&D category level in column 4; clusterization at the R&D stage level in column 5; clusterization at the level of the state of performance in column 6.
2. *Robustness to the estimation method.* The baseline model—column 5 of Table (1), reported in column 1—is replicated with different estimators using the sample from column 5 of Table (1). Columns 2 and 3 report results of a 2SLS estimation using the working sample and the full sample, respectively; column 4 retains perfect predictor variables in the maximization process of the IV probit. This option is typically not used and may introduce numerical instability. Normally, IV probit drops any endogenous or exogenous variables that perfectly predict success or failure in the dependent variable. The associated observations are also dropped. Results are robust to the employment of these alternative estimation methods.

Table C4: Robustness to the Standard Errors Definition

	1(# Patents > 0)					
	(1)	(2)	(3)	(4)	(5)	(6)
Log-# COs	1.42 (0.61)	1.42 (0.63)	1.42 (0.55)	1.42 (0.22)	1.42 (0.43)	1.44 (0.74)
Observations	1173	1173	1173	1173	1173	1170

*Notes:* Baseline results—column 5 of Table (1), reported in column 1—are replicated with different specifications of the standard errors: homoscedastic standard errors in column 2; clusterization at the supplier level in column 3; clusterization at the R&D category level in column 4; clusterization at the R&D stage level in column 5; clusterization at the level of the state of performance in column 6.

Table C5: Robustness to the Estimation Method

	1(# Patents > 0)			
	(1)	(2)	(3)	(4)
Log-# COs	1.42 (0.61)	1.84 (0.85)	1.43 (0.73)	1.42 (0.61)
N	1173	1173	1948	1173

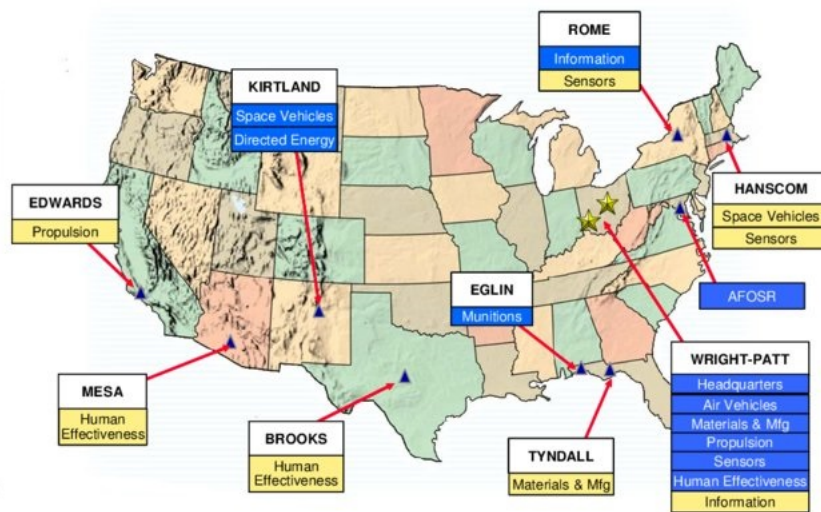
*Notes:* The baseline model—column 5 of Table (1), reported in column 1—is replicated with estimators other than the IV probit. Columns 2 and 3 report results of a 2SLS estimation using the working sample and the full sample, respectively; column 4 retains perfect predictor variables in the maximization process of the IV probit.

## D Appendix: Details on the Air Force Research Lab (AFRL)

The AFRL is a scientific research organization operated by the US Air Force Materiel Command dedicated to the discovery, development, and integration of air and space combat technologies, planning and executing the Air Force science and technology program, and providing war-fighting capabilities to the air, space, and cyberspace forces of the US. The laboratory is divided into eight technical directorates and the Air Force Office of Scientific Research, based on various research areas. The latter is primarily a funding body for external research, whereas the other directorates conduct research internally or on behalf of external entities. Figure (D1) shows a map of AFRL sites throughout the US territory.

Procurement of R&D activities at the AFRL is conducted through six different contracting offices in different branches of the agency. AFRL headquarters is located at Wright-Patterson Air Force Base in Ohio. Its primary functions are leadership, policy, and guidance. The AFRL's only contracting office in Ohio, i.e., FA8650, as coded in FPDS, is located there. The Space Vehicles Directorate is one of the branches of the AFRL. Its mission is to develop and implement space technologies for more effective and less costly war-fighting missions. The Directorate has two headquarters located at different Air Force Bases: Kirtland, New Mexico, and Hanscom, Massachusetts. Both Directorate headquarters conduct R&D procurements, which we track through FPDS with two separate contracting offices (FA9453 and FA8718, respectively). Rome Laboratory is the Air Force "superlab" for command, control, and communications R&D and is responsible for planning and executing the USAF science and technology program. The contracting office FA8750 is installed at Rome Laboratory. AFRL's only R&D procurements in Ohio, New Mexico, Massachusetts, and New York are performed by those offices. Tyndall and Eglin are two Air Force bases, both located in Florida. The AFRL's Florida R&D procurement is conducted by the contracting offices of the two bases (i.e., FA8651 and FA9200). Because these two contracting offices are located in the same state, we are unable to link FedScope information on separations and employment to specific AFRL contracting offices, so they are excluded from our sample. AFRL's four purchasing units provide a diverse science and technology portfolio, ranging from basic and applied research to engineering and operational systems development.

Figure D1: AFRL Sites



Notes: Operating locations of the AFRL in 2006. Source: Mait (2005).

## E Appendix: An example of Broad Agency Announcement (BAA)

Figure E1: Extract of a BAA (first page) from the AFRL

### **2-Step Broad Agency Announcement (BAA) Number: BAA-12-01-PKS Metamaterials for RF and Optical Applications**

#### **Overview Information**

**Federal Agency Name:** Air Force Research Laboratory, AFRL/RV Sensors Directorate

**Broad Agency Announcement Title:** Metamaterials for RF and Optical Applications

**Broad Agency Announcement Type:** This is the initial announcement.

**Broad Agency Announcement Number:** BAA 12-01-PKS

**Catalog of Federal Domestic Assistance (CFDA) Number(s):** 12.800\_AF

#### **THIS WILL BE A TWO-STEP SOLICITATION:**

**First Step: WHITE PAPER DUE DATE AND TIME:** The BAA is open and effective until 20 Feb 2016. White papers will be considered if received prior to 1400 EST on 20 Feb 2016. Submission of white papers will be regulated in accordance with FAR 15.208.

**Second Step: PROPOSAL DUE DATE AND TIME:** To be provided in response to the Requests for Proposals sent to offerors that submit White Papers considered to meet the needs of the Air Force.

**NOTE: White Paper/ Proposal receipt after the due date and time shall be governed by the provisions of FAR 52.215-1(c)(3).** It should be noted that this installation observes strict security procedures to enter the facility. These security procedures are NOT considered an interruption of normal Government processes, and proposals received after the above stated date and time as a result of security delays will be considered "late." Furthermore, note that if offerors utilize commercial carriers in the delivery of proposals, they may not honor time-of-day delivery guarantees on military installations. Early white paper submission is encouraged.

**Solicitation Request:** Air Force Research Laboratory, Sensors Directorate (AFRL/RVD), Wright Research Site is soliciting white papers on the research effort described below. White Papers should be addressed to the Contracting Point of Contact (POC) stated in Section VII of the Full Text Announcement. This is an unrestricted solicitation. Small businesses are encouraged to propose on all or any part of this solicitation. The NAICS Code for this acquisition is **541712**, and the small business size standard is 500 employees. White Papers/Proposals submitted shall be in accordance with this announcement. *There will be no other solicitation issued in regard to this requirement.* Offerors should be alert for any BAA amendments that may permit extensions to the white paper submission date.