Buyers' Workload and R&D Procurement Outcomes: Evidence from the US Air Force Research Lab*

Leonardo M. Giuffrida and Emilio Raiteri

June 14, 2022

Abstract

Is excessive workload a bottleneck for public agencies? Examining R&D procurements by the US government, we link contract, patent, and office records to the officer responsible to estimate how workload affects contract execution. Unanticipated retirement shifts among officers are used to instrument workload. When an officer's workload declines, we find a large increase in patenting—keeping procurement budget and the number of purchases fixed, an additional officer leads to a 2.5 percentage point increase in the probability that a contract generates patents, representing 28% of the sample variation. We provide suggestive evidence that overworked officers cannot devote sufficient time to key contract specifications, resulting in poorer guidance to R&D suppliers.

JEL codes: D23; H57; O30.
Keywords: Workload, Public Procurement, Contracting Officer, R&D, Patents.

*Leonardo M. Giuffrida, ZEW Mannheim, MaCCI, and CESifo, email: leonardo.giuffrida@zew.de; Emilio Raiteri, School of Innovation Sciences, Eindhoven University of Technology, email: e.raiteri@tue.nl. For their helpful comments and suggestions, the authors would like to thank Adriano De Leverano, Francesco Decarolis, Janne Tukiainen, colleagues at their institutions as well as participants at departmental seminars (ZEW Mannheim, University of Turku), workshops (EPPE 2020 at KU Leuven, MaCCI IO Day 2020 at UniMannheim) and conferences (EARIE 2021 Meeting in Bergen, EPIP 2021 Meeting in Madrid) where earlier drafts of this study were presented. Giuffrida is thankful to the Leibniz Association for its support through the project "Market Design by Public Authorities". The usual disclaimers apply.

1 Introduction

Governments use their extensive purchasing power to promote an array of policy objectives. In particular, recent studies stress how technical and scientific progress depends on public spending (Fleming et al. 2019; Moretti et al. 2019). Procurement of R&D from businesses and higher education institutes is one of the main channels through which governments fund innovative activities. In the US, government 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 grants, cooperative agreements, 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 (FAR 35.003), policymakers and scholars increasingly view this type of public spending as a crucial innovation policy tool. Anecdotal and historical evidence suggests that US federal procurement has been critical to the development of some of the most influential technologies invented in the 20th and 21st centuries, including computers, the Internet, GPS systems, and robotics (Ruttan 2006). Recent empirical contributions find support for the existence of such a link (Raiteri 2018).

Given the relevance of R&D procurement contracts for innovation, understanding and quantifying the factors determining their success or failure has policy implications. Innovation scholars emphasize the role of public buyers in clearly identifying their needs and specifying functional requirements for the procured activity or good (Edquist 2015). Empirical evidence on standard procurement confirms that buyers may have a considerable impact on contract performance (Decarolis et al. 2020; Best et al. 2017; Bucciol et al. 2020; Baltrunaite et al. 2021). The demand-side relevance for publicly funded R&D is highlighted by Bruce et al. (2019), who emphasize that federally funded R&D grants in the US have poorer innovation output than cooperative agreements precisely because the buyer has greater discretion over the latter.[1] Decarolis et al. (2021) provide the first empirical quantification of the role of public buyers for the success of R&D procurement. In particular, the authors exploit the variation of employees' death events across agencies and time to estimate the probability that contracts deliver patented inventions. However, the data and methodology used do not allow the authors to provide conclusive evidence on which contracting employment features—and their interplay with suppliers'—are mostly affected by the sudden loss of specialized human capital and matter for the procured R&D performance. Thus, the buyer-innovator relationship remains a black box, unpacking which can provide a better sense of how bureaucracies can promote (or suffocate) innovation (Kelman 2021). The main objective of this paper is to shed light on this relationship. To do so, we closely examine the official responsible for the procurement process—commonly referred to as the contracting officer (CO).

The CO is the leading agent in the acquisition process (Rendon et al. 2012), having the authority to enter into and terminate contracts and make related determinations and findings on behalf of the US Federal government (FAR 1.602). Specifically, they identify the requirements that the procured good or service needs to satisfy, conduct primary and secondary market research to assess the market availability of solutions satisfying the agency's needs, issue the formal request for proposals, and, eventually, select the best source. They are hence responsible for the

[1] Grants or cooperative agreements should be used when the primary purpose of the transaction is to stimulate or support R&D for another public purpose (FAR 35.003)

elaboration of the acquisition strategy and the drafting of requests for information, requests for proposals—establishing the procedure for contract award—contract terms, and the level of competition. Depending on the agency, CO might be also in charge of contract administration but (in particular at larger agencies) other contracting personnel often take the leading role in the contract management phase.[2] The CO has wide discretionary power in performing these tasks, and the FAR clearly states that CO "should be allowed wide latitude to exercise business judgment" (FAR 1.602-2). Arguably, CO's decisions deeply affect the contract awarding phase and the potential frictions the CO might encounter in this process may have a substantial impact on contract performance. This paper focuses on workload, as a possible source of capacity constraints for contracting.

Warren (2014) argues that, in the context of standard procurement, as the workload increases, a CO may have less time to spend on each task, leading to suboptimal contract specifications and poorer performance. This paper, for the first time, devotes special attention to the relevance of CO's workload for the performance of the R&D contracts they award. The role of the CO in the acquisition of R&D work is particularly complex as the CO needs to allow contractors the freedom to exercise innovation and creativity (FAR 35.005). The acquisition strategy needs to be individually tailored to attain the desired objectives of the R&D work, objectives for which the work or methods oftentimes cannot be precisely described in advance (FAR 35.002). One of the key problems for the CO in this setting is describing the work requirement in a way that is clear to prospective contractors, even in situations in which the CO does not have a complete understanding of the work in advance. The CO is responsible for translating a rather abstract idea into a language that is contractually clear to the prospective contractors (US Air Force 1967). To succeed in such a complex task, the CO generally interacts closely with the internal scientific staff and must develop an ad-hoc approach for each new procurement activity (FAR 35.005). Clearly, a higher workload for the CO may be particularly disruptive in this context as it may harm their capacity to allocate enough time to work through the complexity of a new R&D contract. Particularly problematic would be the case where an increase in R&D budgets intended to produce additional technical knowledge is slowed down by capacity constraints in contracting units.

Despite the relevance of the issue, no prior work has analyzed the impact of a CO's workload on the performance of R&D contracts they are responsible for. We do so by combining several different data sources for the first time. First, we use the Federal Procurement Data System (FPDS), which contains information on every contract awarded by US federal agencies. Second, we use the 3PFL Database of Federally Funded Patents (3PFL)—collected by de Rassenfosse et al. (2019) and employed in Decarolis et al. (2021)—to trace patented inventions associated with a federal contract and thus measure R&D contract performance in line with the literature (Bruce et al. 2019). Third, we collect extensive information about the tender stage of the awards from the Federal Business Opportunities (FedBizOpps) website, including information about the identity of the CO. The latter 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

[2] Generally, a contracting officer's representative (COR) assists the CO and is in charge of the technical monitoring or administration of a contract (FAR 1.604).

Research Laboratory (AFRL) and given its high degree of specialization in R&D procurement, we focus on contracts awarded by the AFRL contracting offices. We count the number of COs actively involved in procuring R&D projects in a given year. Following Warren (2014) and controlling for the number of R&D awards and budget for the office, we use this measure as an inverse indicator of the workload. After this data selection and merging process, we end up with a sample of 1,970 R&D contracts representing the universe of AFRL's R&D procurement processes, including COs' records for 2005 through 2012.

Identifying the effect of the buyer's workload on R&D contract performance presents multiple empirical challenges. First, the quality of the CO assigned to a particular contract is correlated with both the outcome and the workload. In case of a sudden and unpredictable jump in workload, the office manager (also referred to as the program manager) could still assign the most cumbersome projects that require extra effort and expertise to the highest quality COs. 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. Thanks to the wealth of information made available by FedBizOpps, we solve this issue by including CO fixed effects. A second concern involves the average complexity of the office's yearly procurement activity, which is unobservable by the econometrician but likely anticipated by the program manager when planning budget and employment. This omitted information also correlates both with the patentability odds of a project and workload. To handle this and additional potential sources of bias, we use an instrumental variable (IV) method that exploits exogenous changes in contracting employment based on unexpected retirement-postponing decisions. For this purpose, we use a fourth data source (i.e., FedScope) which contains detailed characteristics of the public 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. Therefore, we consider the difference between the number of contracting employees eligible for retirement and actual retirees as a good potential indicator of an unanticipated workload shock. 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 staff levels in an office. We show that our IV is strong and does not correlate with any observable characteristics of the work unit.

Our IV estimation strategy augmented with CO fixed effects allows us to estimate a causal effect of *individual* workload on R&D award outcomes, which is more than one order of magnitude larger than estimated by the corresponding endogenous model. Our results stress that the same CO exposed to increased workload determines a decrease in the average probability of awards 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 for an R&D contract to generate patents. The effect corresponds to 28 percent of the average patentability in the sample. 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, patenting would be 50 percent less likely (i.e., the number of patents per contract would be 13 percent lower).

Our findings are robust to the definition of contract performance. Based on the logic of relational contracts (Calzolari and Spagnolo 2009)—which are well-suited for R&D projects, given the difficulty of stipulating quality in the contract definition—we expect firms that fulfill the scope of the contract to be rewarded by the same buyer in the aftermath with further non-competed awards (Che et al. 2021). We propose follow-on contracts—based on the information on the universe of the Department of Air Force procurements—as an alternative metric for contract success. Consistently with our focal analysis, we find that a reduction in the CO workload triggers a significant increase in the probability of the contractor being awarded a follow-on R&D contract without competition.

Our results are also robust to several modifications in the empirical specification, emphasizing the variation of workload as a driver of post-award procurement outcomes—in this vein, advancing the related literature[3]—and have relevant policy implications. First, they confirm the evidence provided in the literature on the importance of the buyer and show that the government should pay particular attention to the workload of contracting employees during the tender stage and ensure adequate staffing. Second, 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 data do not allow us to observe nor proxy the actual effort put in by the CO for an individual contract. Despite this, the heterogeneity of our sample of contracts, buyers, and suppliers allows us to provide indirect evidence of the importance of the CO in drafting solicitations and contractual agreements in a clear work statement. In particular, we show that the CO's workload 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. Less time available for a CO for the award results in poorer guidance embedded in the contract. In fact, we show that suppliers that suffer under high workloads are those which require the guidance the most—that is, small and inexperienced firms—whereas we find that workload has no impact on the innovative output of large and experienced firms. Consistently, workload matters the most in cases where the role of procurement personnel is more central to the R&D process. Such relevance reaches its maximum for contracts awarded for the intermediate stage of development—i.e., applied research—where contracts for the procurement of R&D work are relatively definable. Although we are aware these results do not provide ultimate evidence on the effect channels, they strongly suggest that overworked COs cannot devote enough time to tender and contract specification, and this matters especially when guidance is necessary.

To the best of our knowledge, this work is the first to empirically investigate the relationship between the workload of the CO and the performance of R&D contracts awarded by the same bureaucrat. Given the high complexity of the tasks, the discretionary power of the CO, and the relevance of the pre-award phase in R&D procurement, our empirical analysis validates the

[3]This work relates to the growing economic analysis of the effects of different designs and institutions on procurement outcomes. Examples include awarding design (Decarolis 2014, 2018), wasteful year-end spending (Liebman and Mahoney 2017), external audits (Gerardino et al. 2017), industry consolidation (Carril and Duggan 2020), performance-based insurance schemes (Giuffrida and Rovigatti 2018), and the impact of centralized purchase agreements (Bandiera et al. 2009).

hypothesis that workload is particularly disruptive in this setting. Consistent with our findings, Warren (2014) shows that when procuring supplies and services, contracting offices that experience workload spikes are more likely to choose contract terms that ultimately result in poorer contract performance. Specifically, contracts awarded by busier offices undergo more modifications to the original contract and obligate more dollars. On the other hand, we show that the decreasing supervision of the process induced by a higher workload removes surplus from the firm in the form of less patenting. From the viewpoint of the supplier then, our findings differ from Warren (2014)'s in the sense that the buyer's workload does not induce neutral (i.e., limited competition) or positive—depending on the contract structure—(i.e., cost overruns) effects; on the contrary, we show that higher buyer's workload induces a less innovative R&D activity for the seller.

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 and the robustness analysis. In Section 5, we discuss potential mechanisms underlying our findings. Section 6 draws conclusions.

2 Institutional and Empirical Setting

2.1 Trends in the US contracting staff's workload

In the previous section, we briefly discussed the role of COs in the public procurement process and argued that they are particularly sensitive to an increase in workload, especially in the complex realm of R&D contracting, due to the multitude of tasks involved. But problems arising from overworked contracting personnel would be a second-order problem if spikes in agency workloads were sporadic and temporary events. However, over the past twenty years, 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-á-vis the 1990s. The panel reported a significant mismatch between the demands on the acquisition workforce and the personnel available to meet them and 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.[4] 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 well short of its workforce growth goal, particularly in areas such as contracting and audit (GAO 2017). In 2017, the GAO highlighted persistent problems with the acquisition

[4] Source: `www.usaspending.gov`.

workforce in its High-Risk list and emphasized that DoD agencies were still facing challenges maintaining sufficient staffing levels and overseeing their acquisition workforce.[5] In 2020, the Federal Acquisition Institute claims a shortage of qualified contracting professionals in its professional field brochure.[6]

The increase in the federal contracting budget and institutional concerns are not the only indications that contracting personnel are struggling with problems stemming from excessive workloads. For example, studies based on survey data confirm that federal procurement personnel cite understaffing as one of the primary problems within their work unit (Rau and Stambersky 2009). Specifically, Rendon et al. (2012) show that in two of the DoD's largest subagencies, the Department of the Army and 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.

2.2 Research design and prima facie evidence

As discussed above, peaks in workload strongly influence the discretionary decisions of COs, and previous research has consistently found that, in a non-R&D context, the performance of contracts worsens.The primary objective of this paper is to assess whether an increase in workload within contracting offices procuring R&D negatively affects contract performance.[7] 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 project. Second, we need to consider potential factors that might challenge the causal interpretation of a negative relationship between our measures of workload and project performance.

To address the measurement issues, we rely largely 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 a 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 a proxy for the workload. As in Warren (2014), we control for the number of purchases; in addition, we use the total amount of obligations (*budget*, from now on) of the contracting agency during a fiscal year. The main difference is that our focus is exclusively on R&D contracting. Therefore, our workload measure counts only COs responsible for R&D procurement, and in the same vein, we control only for the purchases and budget for 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; these and other details about the office and the definition of the CO are presented in Section 3.1.

[5] 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`.

[6] Source: `https://www.fai.gov/sites/default/files/1102-Career-Field-Brochure.pdf`.

[7] 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.

Regarding the measurement of R&D contract performance, we follow the recent contribution of 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.[8] 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 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 to be 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.[9] 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. As the figure shows, for 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 essentially 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 since then. Interestingly, the number of patents associated with contracts awarded by the offices (dotted line) seems to show a negative correlation with workload: it decreases significantly between 2005 and 2008 and reverses 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 the procurement

[8] Fixed-price and cost-plus contracts are two different types of contracts commonly used in procurement. In fixed-price contracts, the buyer offers the seller a predetermined price to complete the project. In a cost-plus contract, no price is set, but the contractor is reimbursed for the cost plus a markup. Cost-plus contracts are generally preferred for the procurement of R&D. Approximately 80 percent of R&D contracts over $1 million awarded yearly by federal agencies are cost-plus contracts (source: `usaspending.gov`).

[9] 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

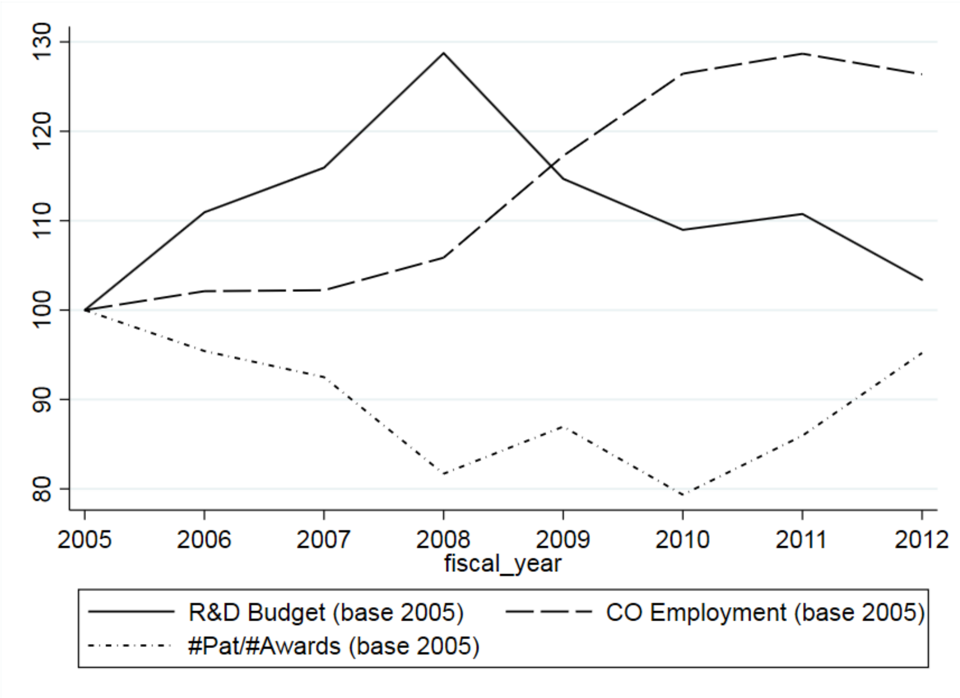

Notes: Average annual R&D obligated budget (solid line), contracting force (dashed line), and number of patents connected to contracts (dotted line). The source is the FPDS database and all agencies awarding at least one R&D contract in the sample are considered.

of 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. Unfortunately, even with a small sample of contracting offices, conducting such an experiment is not feasible. Nonetheless, 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 sufficient information to estimate the following linear probability model:

$$Patent_{i,k,t} = c + \beta \mathrm{Log\,\#\,CO_{k,t}} + \gamma \mathrm{X_i} + \delta Z_{k,t} + \zeta_k + \zeta_t + \epsilon_{i,k,t}, \tag{1}$$

where the variable $Patent_{i,k,t}$ reports whether R&D contract i awarded by the office k in the fiscal year t led at least to a patented invention and the $Log\,\#\,CO_{k,t}$ reports the *CO employment* in log-terms. Let X_i 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 ζ_k and ζ_t include a set of office- and fiscal-year fixed effects. The former control for the unit 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, the use of 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).[10] Thus, after taking into account 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 $t-1$ 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, as discussed by Warren (2014), 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 technology procurement, especially in the procurement of R&D. The size of a contract could depend heavily on the fixed costs of the R&D project to be performed, but 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 explain 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 cumbersome tasks will require more clerk time and effort, 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, $Log\,\#\,CO_{k,t}$, 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 specific 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 cumbersome research project is more likely to lead to inconclusive and thus unpatentable results. As omitted project complexity is negatively correlated with the variable $Patent_{i,k,t}$ and positively correlated with our variable of interest $Log\,\#\,CO_{k,t}$, we would expect the estimate of our coefficient of interest β in Equation (1) to be downward biased.

[10]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 begin preparing 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 in order to operate. The Department of the Air Force, which is the focus of our empirical analysis, sets 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`.

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 the workload and adjust the CO employment to such needs, the manager would still be able to make some adjustments to minimize the impact of an increase in 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 effort and expertise to the higher quality (or more experienced) COs who have 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 particular contract. In particular, we use data obtained 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 reported sparingly and incompletely. 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

$$Patent_{i,k,t,o} = c + \beta \mathrm{Log}\,\#\,\mathrm{CO_{k,t}} + \gamma \mathrm{X_i} + \delta Z_{k,t} + \zeta_k + \zeta_t + \zeta_o + \epsilon_{i,k,t,o}, \tag{2}$$

where ζ_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 cannot be predicted by the program manager. Warren (2014) considers a similar problem and approach and proposes the number of retiring contracting employees as an IV. The main idea behind this choice is that, as discussed by Asch et al. (2005), in the US civil service, retirements are mainly driven by the rules that govern pension obligations. More specifically, the retirement decision is strongly influenced by the attainment of the threshold years of service that qualify an employee for immediate pension benefits. However, at time $t-1$, the head of a given contracting office knows how many of the 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 is likely to 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 (OPM) 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.

At t, actual retirement is realized. Actual retirement 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 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.[11]

3 Data

3.1 Data sources and description of variables

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

FPDS To retrieve contract-specific information, we rely on the 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 (2018), Decarolis et al. (2020). Since 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.[12] 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."[13]

[11]In Section 4.3, we show that our results hold when the actual retirement counts among the contracting staff are used as an IV as in Warren (2014).

[12]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 projects are typically very high, and we can confidently state that we observe the universe of these projects over the period under analysis.

[13]The variable consists of two alphabetic and two numeric digits. The first digit is always the letter "A" to identify R&D; the second digit is alphabetic "A" through "Z" to identify the major category of R&D; the third digit 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 digits. The fourth digit 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 directed 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 design concept and/or prototype.

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 it is awarded in the last week of the fiscal year (i.e., the last seven days of September).[14]

3PFL de Rassenfosse et al. (2019) exploit the FAR to trace patented inventions directly related to federal contracts. The 3PFL database covers USPTO patents granted between 2005 and 2018. 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 total 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 choose to 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. So 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.[15] Nevertheless, as discussed in de Rassenfosse et al. (2020) and Decarolis et al. (2021), 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.

FedBizOpps The federal acquisition process begins when an agency determines its requirements and how to procure them. If the agency's 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 post a solicitation on the FedBizOpps platform.[16] 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.

[14] 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.

[15] The Invention Secrecy Act of 1951 governs this process.

[16] 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.

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 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 3.3, this variable is the main explanatory variable for the project innovation outcome under study. In addition, we define *Specialist* as a dummy variable indicating whether the CO is assisted by a contract specialist for a particular R&D process.[17]

FedScope 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.[18] 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, were 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 all second-order or 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)—to check the robustness of our results to the definition of the endogenous employment variable in Section 4.3.

Importantly, the FedScope dataset allows us to build our IV as outlined in Section 2.2. The retirement eligibility of federal civilian employees is determined by age and the number of years of creditable service. People that 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).[19] To identify the number of retirement-eligible

[17] 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 and administer the contract once awarded. A specialist is not always necessary—they appear in 81 percent of contracts in our sample, see Table (2)—, while a CO is always required for a purchasing process to be initiated.

[18] 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.

[19] For details, see: `https://www.opm.gov/retirement-services/fers-information/eligibility/`.

Figure 2: AFRL Sites

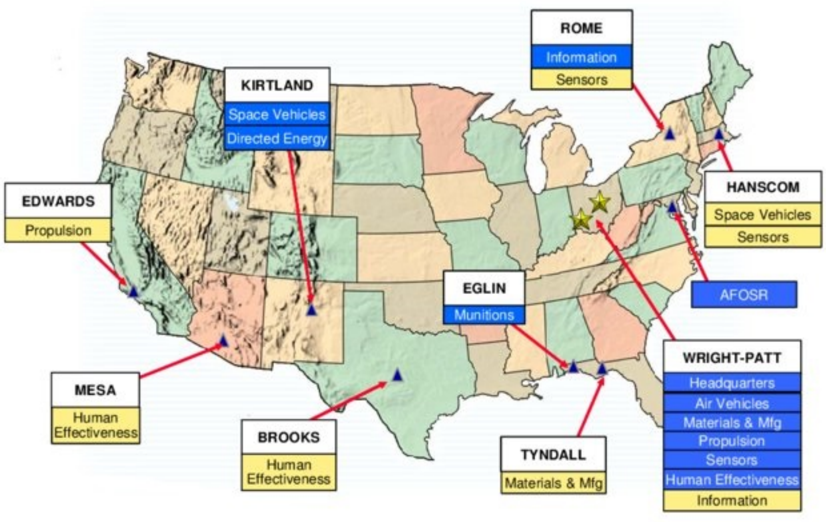

Notes: Operating locations of the AFRL back in 2006. Source: Mait (2005).

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.[20] Instead, *Retirement Actual* indicates the 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 Air Force Research Lab

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 the geographic information to link FedScope to the other data sources.

For the sake of clarity, 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 subunit of the Department of the Air Force. 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 warfighting 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

[20]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 available at: `https://www.fedscope.opm.gov/employment.asp`.

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 (2) 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. According to the FPDS data, procurement of R&D activities at the AFRL equals to $2.22 billion per year in 2005-2012.[21] 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 warfighting 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 are 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.[22] AFRL's four purchasing units provide a diverse science and technology portfolio, ranging from basic and applied research to engineering and operational systems development. 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 contracting office is located.

3.3 Sample selection and descriptive analysis

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

[21] To benchmark this budget, we report that the SBIR program obligated a similar average yearly amount during the same period. See `https://www.sbir.gov/awards/annual-reports`. Moreover, such spending amount represents one-third of the total NSF's budget for FY 2009. See `https://www.nsf.gov/news/speeches/bement/09/alb090514_budget.jsp`.

[22] For a full description of the sample selection, see Section 3.3.

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, but we have relied on an external dictionary that maps the variable "Contracting Office Agency ID" in FPDS to the variable "AGYSUB" of Fedscope.[23] Following the discussion above, we limit our focus to the combined project-level information associated with AFRL awarding agencies FA9453, FA8650, FA8718, FA8750, which represent 88 and 83 percent, respectively, of the spending and contract counts in the AFRL raw sample.

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.[24] 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).

Table (1) shows the details of the R&D activities included in the 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.[25]

[23]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.

[24]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, these contracts are intended to assist certain small businesses in conducting innovative activities aimed at their eventual commercialization, not their patentability (Howell 2017; 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 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.

[25]See Bajari and Tadelis (2001) for a detailed study of the trade-off between time and cost to completion induced

Table 1: 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 2: 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 Competed (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 contract associated with at least one patent is 8.9 percent.

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 (2). 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. (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 (3) shows the characteristics of the contracting units. Each office spends an average of \$0.54 billion per fiscal year on 65 different R&D projects. In the median office, the GS-1102 are 111, 27 of which are COs and 2 have managerial responsibilities.[26] We refer to the latter category as *Top GS-1102*.[27] 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

by contract pricing format.

[26] 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.

[27] 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.

percent of the Retirement Eligible). This results in high non-retiree counts, i.e., 21 per office.[28]

Table 3: 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.

Finally, we define firm-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 firm 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

4.1 Baseline results

We begin the presentation of our results with Table (4) row (1), which displays the estimates corresponding to the binary version of Equation (2), that is,

$$Pr\left(Patent_{i,k,t,o} = 1 | \mathrm{Log\#CO_{k,t}}, \mathrm{X_i}, Z_{k,t}, \zeta_t\ , \zeta_o\right) = \Phi\left(c + \beta \mathrm{Log\#CO_{k,t}} + \gamma \mathrm{X_i} + \delta Z_{k,t} + \zeta_t + \zeta_o\right), \qquad (3)$$

where Φ is the cumulative standard normal distribution function. Thus, we estimate the binomial response to our variable of interest β via a probit model.[29] 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 each additional CO colleague as an actual reduction in the office's total procurement workload. Column

[28]Since 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 (3) differs from the mean difference between Retirement Eligible and Retirement Actual.

[29]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 accommodate well both methodologies.

2 contains fixed effects for the CO and the fiscal year. The former are key to our identification, as discussed in Section 2.2; also, they bring in 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 3.3 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 are 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 category 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 be simultaneously correlated with outcome and treatment. Also, we enrich this model specification by adding *Small*, *Inexperienced*, and *University* as firm-level controls covariates accounting for the different frictions and incentives to channel innovation through patent filing. This is our favored specification.[30] To facilitate interpretation of the estimates, we report as coefficients the average marginal effects with robust standard errors. In Appendix C, we discuss 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 2.2. To address these concerns, we implement an IV strategy based on non-retirement as the instrument.

For this approach to be valid, we are to satisfy two conditions. First, the instrument must cause the variation in the treatment variable; in other words, the variation in non-retirement rates over time must have some causal power in explaining the variation in CO employment. Second, the instrument must not affect the outcome variable directly but only indirectly through the treatment variable. For this exclusion restriction to be satisfied, we need to show that the instrument affects

[30] Please note that the working sample amounts to 1,173 observations instead of the 1,970 presented in Table (2). This sample reduction is mostly due to the CO fixed effects predicting success or failure perfectly. For the sake of 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 the use of a 2SLS using the full sample of contracts.

Figure 3: Visual Representation of the First Stage

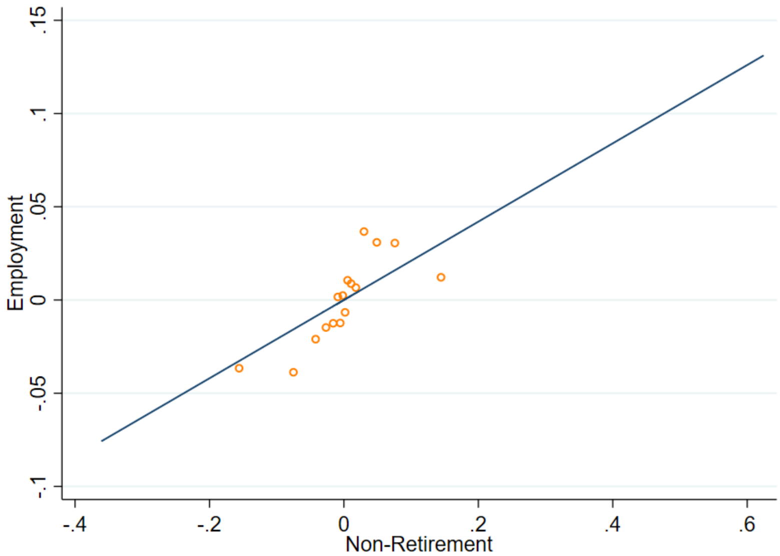

Notes: Graphical representation of the relationships between # COs and Non-retirement. The variables are residualized, including as controls those from column 5 of Table (B2). Each graph is a binned scatterplot. This means that each point represents the mean statistic of the residualized IV and the residualized endogenous variable inside each bin. The selected number of bins is 122, and it is optimal in minimizing the (asymptotic) integrated mean-squared error following Cattaneo et al. (2019).

the outcome variable only through the treatment (i.e., the non-retirement must relate to the patents because it correlates with CO employment), conditional on other possible confounding effects (i.e., the instrument must be independent of potential outcomes).

As for the strength of the instrument, Table (B2) in Appendix B 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 points 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.[31] Standard statistical tests of the performance of our instrument—reported at the bottom of Table (B2)—reject weak and under identification, and advocate a strong first-stage relation. Figure (3) 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 (4) as controls and grouping them in binner scatterplots as in Cattaneo et al. (2019). More specifically, each dot represents the residualized IV's mean statistic and the residualized endogenous variable within each bin. This graphical evidence further stresses the existence of a positive effect of our IV on project innovativeness.

In terms of the exclusion restriction of our instrument, we believe that the sudden non-separation refers to an existing experienced employee who has already reached full productivity

[31] We emphasize again that the instrument is constructed using FedScope data and the GS-1102 classification, which do 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.

and not to a newly hired officer who has not yet reached full productivity. As a result, we expect this exogenous labor surplus as good for the R&D office outcomes, but only through a variation in workload. This is confirmed by the reduced-form relationship between patent and the instrument, as the coefficients on the instrument tend to enter with a positive and significant effect on our outcome variable (see Table B3 in Appendix B). More specifically about the exclusion restriction, we 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 retirement now versus later is idiosyncratic and depends on personal circumstances, in addition to 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 the workplace characteristics affect the individual decision to postpone retirement. This would create a potential reverse causality problem, in particular when large changes in workload somehow induce people to stay at work even though eligible for retirement. To test this, we collapse the data at the office-year level and run an auxiliary regression analysis—presented in Table (B1) in Appendix B—to detect possible observable determinants of our instruments and provide evidence for our exogeneity argument. Based on the way we construct the instrument, we find that sudden postponement of retirement is 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 that we include through our data and that appear in Table (B1): 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 A.

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 (4) shows that the estimated effects of decreasing workload in the patentability of R&D contracts are not significant. This result would suggest that additional CO colleagues in the office do not affect the innovativeness of their purchases. Offices with more COs are not more or less likely to generate patents given the same budget and purchases. However, the IV probit results shown in Table (4) row (2) suggest that the structural probit results are misleading when exogenous changes in the number of COs are taken into account. The set of controls is identical in all columns and is the same as those of Table (4). According to the baseline IV probit estimates from column 5, in our sample an additional CO in the office—corresponding to a 3 percent increase in average CO employment—during the award year 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.[32] Compared with the probit estimates, the magnitude of the IV probit estimates is more than one order of magnitude larger than and exceeds the 95 percent confidence interval of the structural relationship. Moreover, the Wald chi-squared test of exogeneity of the IV is rejected at 95% level, highlighting the endogeneity in the structural model, the inappropriateness of the regular probit regression, and the need for the IV Probit.[33]

To provide a more transparent economic 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. To do so, we first collapse our information at the office-year level. Then, we regress employment in contracting offices 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 on a yearly basis). This is an economically large effect. For instance, Breitinger et al. (2020) recently estimated 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 huge 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[34]—to the size of the median R&D contract awarded by a DoD research

[32] From Table (2), 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 111 percentage points increase in the probability of patent equal to 1 (or equivalently, an increase in the expected probability of $0.089 + 1.11 = 1.199$). Such an implausible number results from the huge increase in the underlying predictor. From the perspective of the CO employment variable itself, # COs, being natural-log-transformed, is multiplied by approximately 2.718, representing an approximately 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 111 percentage points by one log-units 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 to 32.11*(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 $1.11/55.16 \approx 2.5$ percentage points.

[33] 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 firm-level controls. In an auxiliary exercise, we include fixed effects for the supplier and remove firm-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 firm characteristics interplay with our results in a battery of sample splits.

[34] `https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/2020/general-schedule/`

Table 4: Baseline Results

	$\mathbb{1}$(# Patents > 0)				
	(1)	(2)	(3)	(4)	(5)
(1) Log-# COs (Probit)	0.16	0.049	0.085	0.087	0.092
	(0.048)	(0.11)	(0.11)	(0.12)	(0.12)
(2) Log-# COs (IV Probit)	0.14	1.14	1.15	1.37	1.42
	(0.071)	(0.64)	(0.60)	(0.61)	(0.61)
Wald χ^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 last week of 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.

laboratories—such as the Air Force Research Lab—, which is approximately \$1 million over 3 years. Hiring COs results 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 Further results

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 proxy for innovation are well known and 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. A follow-on contract is defined 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 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, that 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 ground 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 identify a link between an initial R&D contract and a potential follow-on contract. The main advantage of using patents as 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 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 one if the company involved in the procurement of our focal R&D contract i wins a sole source, non-competed contract with the Department of Air Force in the year of completion of the focal contract i, and zero 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 disentangle between follow-on contracts awarded for the performance of R&D work and contracts awarded for the performance of other types of services or products supply. Panel a of 5 reports the results of the estimates obtained when using our measure(s) of follow-on contract as outcome variable(s) and adopting the same estimation strategy as in the focal analysis (colum 5 of Table 4). Column 1 and column 2 show that a reduction in the workload of the CO result in a significant increase in the probability of the focal contractor to win a non-competed 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 an non-R&D contracts. The results appear to be entirely driven by non-competed 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 lead 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 contract exclusively and constructing three additional follow-on variables that accounts for three separate R&D stages: basic, applied, and developmental R&D. Clearly, if our interpretation is correct and non-competed, 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-competed contracts awarded for more advanced R&D stages that builds on the one carried out in previous phases. Panel b in table 5 reports the results of this exercise and confirm our hypothesis. The positive effect of a reduction in the CO's workload on the likelihood of winning a subsequent non-competed 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.

4.3 Robustness analysis

We conclude this section with a brief summary of the findings from the main robustness checks. To simplify the presentation, 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. We refer the reader to Appendix B for figures and tables plus additional robustness checks.

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 (4)—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 B2) 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. Finally, 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.

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 competed 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, we believe that 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.

Table 5: Follow-on Contracts

Panel a: R&D vs. non-R&D

	$\mathbf{1}$(# Follow-on Contracts > 0)					
	(1)	(2)	(3)	(4)	(5)	(6)
Log-# COs	1.54	1.59	2.16	2.27	0.20	0.38
	(0.62)	(0.59)	(0.50)	(0.45)	(0.73)	(0.71)
N	1663	1703	1576	1634	1583	1646

Panel b: R&D Stages

	$\mathbf{1}$(# Follow-on Contracts > 0)					
	(1)	(2)	(3)	(4)	(5)	(6)
Log-# COs	0.57	0.43	0.64	0.73	1.58	2.07
	(0.67)	(0.71)	(0.41)	(0.44)	(0.57)	(0.53)
N	1010	1056	1125	1255	1509	1568

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

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 (2). The coefficient is insignificant, while the coefficient on our main variable is only marginally affected. As we do not know the difference between solicitation 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 proposal/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.

5 Mechanisms

The results of the analyses presented in the previous section show that the workload of the officer in charge of the award of an R&D procurement process matters for the contract performance as measured by patents, but also follow-on contracts. In this section, we discuss possible mechanisms underlying these findings. In particular, we show that a key dimension affected by workload is the reduced time a CO is forced to dedicate to the contract and the resulting poorer guidance to the prospective R&D supplier.

Although alternative channels are potentially in question, our results suggest that the impact of the buyer's workload on contract performance generally passes through tender and contract specifications. Indeed, we stress that the non-retirement shocks occur in the year of contract award that, together with the multi-year nature of R&D contracts—see Section 3.3—and the multi-month duration of the award phase, implies that the effect we estimate chiefly involves the pre-award period. This interpretation is in line with the main findings of Decarolis et al. (2021), who show that especially the disruptions happening in a contracting office in the pre-award phase affect R&D contract performance, whereas disruptions occurring during the execution and management phases of a contract only have a minor impact on the probability for a contract to be associated with a patented invention. The relatively lower importance of the ex-post monitoring and contract management phase is likely linked to the intrinsic difficulty for CO to monitor the advancement of research projects vis-á-vis more standard procurement and also explained by the fact that, as discussed in the introduction, the CO is not necessarily the contracting professional who directly supervise the execution stage of R&D.[35] All in all, this evidence speaks in favor of the relevance

[35] 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

of the design and award stage relative to ex-post contract monitoring for demand-side drivers of R&D contract performance. Yet, this leave us with the following question: Why are contracts awarded by CO affected by heavier workload in the pre-award phase performing more poorly down the road?

Workload induces the tightening of time constraints Our main argument is that a marginal increase in workload intensifies time pressure for COs providing them with less time than necessary to be devoted to contract drafting; in turn, this constrained time allocation makes guidance to contractors in tender and contract drafting poorer.

The time dedicated to the contract is an arduous dimension to investigate. The actual time the CO dedicates to a project is not recorded nor reported in any way in official data, and, even if it were, we could not claim its appropriateness for a specific R&D process. Our two-fold approach to show the relevance of time pressure in our context is the following. First, contracts for which the CO has less time available to award should be more heavily impacted by an increase in workload as, arguably, the CO is 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. In Table 6 Panel a column 1, we restrict the sample to awards signed in the last quarter of the fiscal year, i.e., July to September. 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 does have a substantial positive effect on R&D performance.[36] Second, based on the same results-interpretation exercise presented in Section 4.1, we rank the office-year pairs in terms of residualized employment (i.e., inverse workload). We group the office-year units in the first tercile of residualized employment distribution and label them as "high-workload". In Table 6 panel a, we report the estimates performed on contracts assigned by the high-workload office (column 2) and others (column 3). Estimates hold only in the former sample split and turn negative and insignificant for buyers with lighter workload. We argue that the estimated effect is non-linear in the underlying workload and applicable when workload levels are relatively high in our sample of offices and years.

Contract Drafting and Guidance As discussed in the introduction, 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 contractually 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). A higher workload that tightens time constraints may harm the CO's capacity to translate tender-specific but uncertain technical objectives into contractually binding behavior for prospective contractors. As a result, the quality of the indications provided in request for proposals and contractual agreement that are supposed to guide the

https://media.defense.gov/2007/Sep/28/2001712655/-1/-1/1/07-130.pdf.

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

sellers decreases. The importance of such guidance for contract performance is likely to interact with suppliers' characteristics. Less (or poorer) guidance, when it is required the most, may result in the poorer performance of the supplier. To advocate this argument, in Panel b of Table 6, we exclude potential channels of supplier characteristics and present a battery of sample splits to corroborate the relevance of this channel. In particular, column 1 (2) restricts the sample on small (large) firms, while in column 3 (4) we limit the focus on inexperienced (experienced) firms. The workload effect is stronger when the guidance is more necessary. This applies to small firms and winners with no previous experience in federal R&D procurement. On the contrary, results turn insignificant for large and experienced suppliers. Experienced sellers are likely less influenced by contract features and more likely to carry out the project irrespective of buyer's frictions.

Table 6: Mechanisms and sample splits

Panel a: Time Constraints

	$\mathbf{1}$(# Patents > 0)		
	(1)	(2)	(3)
Log-# COs	2.44	0.38	-1.36
	(0.89)	(0.24)	(1.70)
N	339	377	686

Panel b: Size and Experience

	$\mathbf{1}$(# Patents > 0)			
	(1)	(2)	(3)	(4)
Log-# COs	2.54	-1.98	2.01	0.81
	(0.54)	(1.80)	(1.01)	(0.56)
N	466	438	264	1684

Panel c: R&D Stages

	$\mathbf{1}$(# Patents > 0)		
	(1)	(2)	(3)
Log-# COs	0.06	2.47	1.85
	(3.05)	(0.89)	(2.29)
N	484	1014	450

Notes: Baseline analysis is replicated: in the subsample of last quarter of fiscal year (Panel a, column 1), conditioning on highly workloaded offices (Panel a, column 2), not highly workloaded (Panel a, column 3); conditioning on small (Panel b, column a), large (Panel b, column 2), inexperienced (Panel b, column 3), experienced (Panel b, column 4) winners; and in the subample of basic research (i.e., stage 1), applied research (i.e., stage 2), development (i.e., stage 3-6) contracts in columns 1, 2, and 3 of Panel c, respectively. R&D stages are omitted in specifications 1 and 2 of Panel c due to collinearity. Coefficients report average marginal effects of the the IV probit regression. The coefficients are estimated via 2SLS in Panel b, column 3 and 4, and Panel c, columns 1-3 due to insufficient power. Robust standard errors are in parenthesis.

Further evidence on the time/guidance channel Although we can neither observe nor proxy the actual effort that the CO puts in for specific procurement activities and, therefore, we cannot determine if the increases in workload lead to a reduction in the effort provided for a single contract, the heterogeneity of the R&D activities in our sample allows us to provide indirect evidence of the importance of the COs in drafting solicitations and contractual agreements nonetheless. In this fashion, the contracts in our sample are awarded for the performance of projects at different stages of the R&D process. Table (1) reports that 29 percent of the contract was awarded for the performance of basic research, 49 percent for the performance of applied research, and the remaining 22 percent for development research.[37] The importance of the CO in the tender specification is likely to change with the stage of the R&D work procured. In the procurement of basic research whose direct applicability is still uncertain, agencies often rely on a Broad Agency Announcement (BAA).[38] BAAs are broad in their subject matter and may be used by agencies to fulfill their requirements for scientific study and experimentation directed toward advancing the state-of-the-art or increasing knowledge or understanding rather than focusing on a specific system or hardware solution. Generally, a BAA describes the agency's research interest, either for an individual program requirement or for broadly defined areas of interest, but does not include a clear and complete work statement concerning the area of exploration and the end objectives of a contract. In this context, the importance of the CO in drafting a request for proposal and a subsequent contract that would provide the right incentive to the prospective contractor appears to be rather limited. However, contracts for the procurement of R&D work that happens at more advanced stages of development are relatively definable. In such cases, it is rather straightforward for the CO to draft a request for proposal and the resulting contract in a clear and complete fashion.

We expect that the relevance of the CO's ability in translating an abstract idea in a contractually clear work statement reaches its maximum for contracts awarded for more advanced stages of development. This applies in particular to the applied research (i.e., intermediate stage) where projects for the procurement of R&D work are enough definable for the CO to have an impact but not too definable—like in the development phases—to make their contribution mostly formal. Hence, if the negative effect of increases in workload on contract performance is connected to a reduction in the effort of the CO in the key moments of solicitation and contract drafting, we should expect our results to be stronger for contracts awarded for applied research than for the other stages. To test this hypothesis, we split our sample into three groups based on the three R&D stages and ran the same model as in the focal analysis. Columns 1-3 in Table (6) panel c reports the results of this split-sample analysis. As the table shows, the overall negative baseline effect seems to be largely driven by contracts awarded for applied research (column 2), with a insignificant effect for basic research (column 1) and development stages (column 3). Although we are aware this result does not provide conclusive evidence—and point estimates across these sample splits are not statistically different—it vigorously hints at the proposed idea that overworked COs are indeed unable to devote enough time to tender and contract specification, resulting in a reduction in the performance of the contracts they award when the contribution of the CO is most required.

[37]Due to the high comparability in their nature, we group the subsequent development stages (i.e., stages 4 to 6) together with stage 3 within the "broad development stage" umbrella.

[38]See FAR 35.016.

6 Conclusions

As modern economies have evolved from industrial- to knowledge-based, the activity of public administrations has also changed over time, and the governance of technology has become a crucial part of their activity. In public procurement, these structural changes are compelling governments to procure increasingly complex products and services. In the US, this increasing complexity is going hand in hand with an increasing number and value of procurements. In response, the government has focused on streamlining acquisition rules and giving more discretion to front-line officials (Carril 2021; Calvo et al. 2019; Giuffrida and Rovigatti 2018). However, human resources are lagging behind due to the inability to retain and recruit talent at all levels of the administration. The result is an overburdened contracting workforce, which generates a major capacity constraint for effective public spending in general, and investment in R&D in particular.

To provide a quantification of this bottleneck, this paper sheds light on the microeconomic implications for innovative activities of the workload of US federal buyers of R&D. Our research is in the spirit of Furman and Stern (2011)'s as we are interested in understanding how institutions shape new knowledge and what frictions they encounter, and complements Warren (2014)'s findings on how a overworked contracting personnel affects standard procurements' performance outcomes. We combine several data sources to link tender, contract, patent, and office records to the identity of the CO. We focus on contracts awarded by the contracting offices that compose 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 the office's workload, after controlling for the budget and purchases. To overcome endogeneity in the contract outcomes and buyer's workload relationship, we implement an IV strategy—combined with CO fixed effects—to identify the effect of the latter on the former. The identification comes from unanticipated retirement shifts among COs, which we use as an instrument for the workload. Our results are robust to several modifications and stress 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 the outcome definition: A reduction in the CO workload triggers a significant increase in the probability of the contractor to be awarded a follow-on contract without competition. The results appear to be entirely driven by contracts awarded for the performance of R&D work, in particular for the development stages. We show that a key dimension affected by workload is the reduced time a CO is forced to dedicate to the award process and the resulting poorer guidance to R&D supplier embedded in the contract. Consistently, workload matters the most in cases where the role of procurement personnel is more central to the R&D process. This applies to applied research, where translating the still abstract results from the basic research phase into detailed contract requirements, a typical task of the CO, is as complex as it is necessary for successful project performance.

Our findings acknowledge that the functioning of public bureaucracies is a key driver of government effectiveness (Weber 1921) and, in turn, the ability of the state to govern effectively is a crucial determinant of economic prosperity (Acemoglu 2005). Combined with the contracting personnel costs, we suggest that the government would gain considerably further innovation power if it dedicated part of the budget for R&D contracts to hire additional COs. Future research could use our results as a starting point to further investigate how excessive workload constrains bureaucratic effort and how the work environment and characteristics of contracting personnel interact with our proposed channel. Of particular interest is the mitigating role of management practices and the response of bureaucrats with high levels of implicit motivation 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.*

Acemoglu, D. (2005): "Politics and economics in weak and strong states," *Journal of Monetary Economics*, 52, 1199–1226.

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.

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.

Best, M. C., J. Hjort, and D. Szakonyi (2017): "Individuals and Organizations as Sources of State Effectiveness," NBER Working Paper 23350, National Bureau of Economic Research, Cambridge, MA.

Bhattacharya, V. (2021): "An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program," *Econometrica*, 89, 2189–2224.

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," .

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.

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.

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.

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.

Fleming, L., H. Greene, G. Li, M. Marx, and D. Yao (2019): "Government-funded research increasingly fuels innovation," *Science*, 364, 1139–1141.

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 (2018): "Can the Private Sector Ensure the Public Interest? Evidence from Federal Procurement," ZEW Discussion Paper No. 18-045, ZEW - Centre for European Economic Research, Mannheim.

Howell, S. T. (2017): "Financing Innovation: Evidence from R&D Grants," *American Economic Review*, 107, 1136–1164.

Kang, K. and R. A. Miller (2021): "Winning by Default: Why is There So Little Competition in Government Procurement?" *The Review of Economic Studies*, forthcoming.

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.

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.

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.

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.

WEBER, M. (1921): *Economy and Society: An Outline of Interpretive Sociology (1921; German: Wirtschaft und Gesellschaft. Grundriß der verstehenden Soziologie)*, vol. 1, University of California Press (1978).

A Appendix: Discussion on the Exclusion Restriction

To corroborate further 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 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 (A1) 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 graph, Figure (A1) shows the 2SLS asymptotic 95 percent confidence interval (the yellow area) for β, 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 A1: Exclusion Restriction Test as in Kiviet (2020)

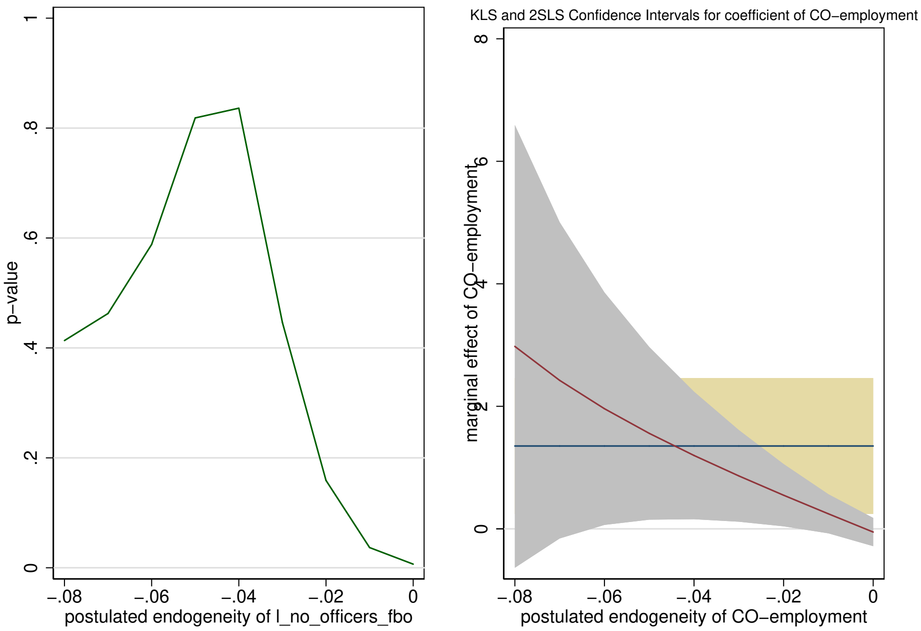

Notes: Values of single just-identifying exclusion restriction tests and inference of 2SLS and KLS based on non-orthogonality conditions.

B Appendix: Additional Tables

Table B1: Non-retirement Predictors

	Log-(Non-retirement)					
	(1)	(2)	(3)	(4)	(5)	(6)
(mean) University	-0.79	-0.80			-1.04	-0.29
	(0.53)	(0.92)			(0.44)	(0.57)
(mean) Small	-0.18	0.50			-0.22	0.97
	(0.16)	(1.83)			(0.25)	(0.53)
(mean) Inexperienced	0.029	-1.07			0.89	-1.85
	(0.96)	(3.22)			(0.71)	(1.18)
(mean) Specialist	-0.092	-5.95			-0.67	-1.03
	(0.43)	(0.75)			(0.62)	(0.68)
(mean) Same State	-0.33	-3.40			0.43	0.38
	(0.60)	(1.75)			(0.69)	(0.51)
(mean) Last Week	0.56	0.043			1.02	-1.01
	(0.63)	(2.68)			(0.67)	(1.03)
(mean) Log-Budget			-0.11	-0.13	-0.24	0.26
			(0.24)	(0.062)	(0.34)	(0.23)
(mean) Log-Purchases			0.0085	0.19	-0.16	-0.086
			(0.082)	(0.079)	(0.12)	(0.13)
(mean) # GS-1102			0.63	0.90	0.38	0.79
			(0.21)	(0.042)	(0.36)	(0.12)
(mean) # Top GS-1102			-0.024	-0.022	-0.075	-0.081
			(0.020)	(0.039)	(0.044)	(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 B2: First Stage Regressions

	Log-# COs				
	(1)	(2)	(3)	(4)	(5)
Log-(Non-Retirement)	0.152	0.203	0.206	0.208	0.210
	(0.00753)	(0.0513)	(0.0512)	(0.0556)	(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 B3: Reduced-form Regressions

	$\mathbf{1}(\#\text{ Patents} > 0)$				
	(1)	(2)	(3)	(4)	(5)
Log-(Non-Retirement)	0.024	0.22	0.23	0.27	0.28
	(0.011)	(0.12)	(0.12)	(0.13)	(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.

C Appendix: Robustness Analysis Tables and Additional Discussion

Robustness Checks Tables We report tables from Section 4.3.

Table C1: Alternative Specifications of Endogenous Employment

	$\mathbf{1}$(# 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 (4), 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 in parenthesis. **p < .05, ***p < .01

Table C2: Alternative Instruments

	$\mathbf{1}$(# Patents > 0)			
	(1)	(2)	(3)	(4)
Log-# COs	1.42	1.87	1.76	0.81
	(0.61)	(0.52)	(0.48)	(0.35)
N	1173	1028	1028	1173

Notes: Baseline results—column 5 of Table (4), reported in column 1—are replicated with alternative IVs: Retirement Eligible $t-1$ / Retirement Actual t in Column 2; log(Retirement Eligible $t-1$ / Retirement Actual t) in Column 3; Log(Retirement Actual t) in Column 4. Coefficients report average marginal effects. Robust standard errors are in parenthesis. **p < .05, ***p < .01

Table C3: Inclusion of Endogenous Omitted Controls

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

Notes: Baseline results—column 5 of Table (4), 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 in parenthesis.

Additional Robustness Checks We report additional robustness checks. For convenience of the latter, 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 (4), 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 (4), reported in column 1—is replicated with different estimators using the sample from column 5 of Table (4). 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

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

Notes: Baseline results—column 5 of Table (4), 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

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

Notes: The baseline model—column 5 of Table (4), 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.