Breaking Down the US Employment Multiplier Using MicroLevel Data*

Edoardo Briganti Holt Dwyer Ricardo Duque Gabriel

Victor Sellemi

December 2024

Abstract

We use restricted data from the Quarterly Census of Employment and Wages to link the universe of US establishments with the universe of contractors in the Federal Procurement Data System. Leveraging detailed institutional knowledge of federal acquisitions, we construct a new dataset of unanticipated contracts and examine their effects on employment growth. We find positive, significant, and persistent effects on firms with fewer than 150 employees. Using loan data from the US Federal Reserve (Y14-Q), we show that small firms expand their credit and experience lower interest rates after winning unanticipated contracts. At the regional level, we estimate a cost-per-job of $57,000 per year using unanticipated contracts—an order of magnitude lower than previous estimates based on all defense contracts. Lastly, we leverage the restricted census data to decompose the employment multiplier into a direct effect on contractors of 55% and an indirect effect on non-contractors of 45%.

*Acknowledgments and Disclaimers:

This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The research proposal was submitted to the BLS in September 2021. The first visit to the BLS was delayed until March 2024 because of the COVID-19 pandemic. The views expressed here are those of the authors do not necessarily reflect the views of the BLS, the US government, the Bank of Canada, the Federal Reserve Board or the Federal Reserve System. We are particularly thankful to an anonymous US government contracting officer for their extensive help with the institutional knowledge of federal contracting. We are also thankful for useful comments from Valerie Ramey, Johannes Wieland, Munseob Lee and other seminar participants at UCSD.

Emails: bred@bank-banque-canada.ca, bdwyer@ucsd.edu, ricardo.f.duquegabriel@frb.gov, vsellemi@ucsd.edu.

I. Introduction

Government spending plays an important role in macroeconomic stabilization, yet the mechanisms through which it operates are not fully understood. Government purchases, in particular, may influence employment by acting as a positive demand shock, either directly through contractors or indirectly via demand spillovers to other firms. Conversely, additional labor demand by contractors could crowd out employment at non-contractor firms. The relative contributions of these direct and indirect channels of government purchases are therefore, unclear. Moreover, direct evidence of the firm-level transmission mechanism operating through the extensive margin of labor remains scarce.

This paper leverages restricted access to firm-level data in the United States to demonstrate that a novel measure of unanticipated government contracts enables small firms to persistently increase their number of employees. Using restricted loan-level data, we show that this mechanism allows small firms to expand their credit and reduce the interest rates on their loans, supporting the existence of a *borrow-to-hire* mechanism. In the second part of the paper, we aggregate our new measure of unanticipated contracts at the regional level. We estimate a cost-per-job of approximately \$57,000 per year, an order of magnitude lower than existing estimates based on total defense contracts and consistent with the cost-per-job estimates for other types of government spending in the literature. Lastly, we find that, on average, 55% of the employment multipliers originate from the direct employment expansion of government contractors. We observe minimal evidence of crowding-out effects on impact. Despite the negative initial impact, non-contractors' contribution to the employment multiplier becomes positive, accounting for the remaining 45%. This indicates that demand shocks propagate through *multiplier effects* to firms not directly benefiting from procurement contracts. The observed 55-45% split suggests that direct and indirect effects are nearly equally important.

First, we provide empirical evidence of the employment channel of procurement contracts in the U.S. by utilizing restricted access to the complete universe of U.S. establishments in the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics (BLS). We link firms from QCEW with the universe of contractors in the Federal Procurement Data System (FPDS). We investigate the effects on employment of a new measure of contracts that is "*unanticipated*" by firms. However, after 2006, most federal procurement contracts stem from contract types whose effects firms might anticipate. To address this, we select only contracts that are (i) newly awarded, (ii) standalone and not part of any ongoing government agreement, (iii) competitively bid, and (iv) received at least two offers. Our findings suggest that unanticipated contracts have significant positive effects, particularly on small firms with fewer than 150 employees.

Using information on bank loans from the Y14-Q database of the Federal Reserve, we link the universe of firms receiving unanticipated contracts with available credit data since 2012. We find that winning a contract increases employment, enhances credit access, and reduces the average interest rate firms pay on their outstanding loans. This effect is evident only among small firms, suggesting the existence of credit constraints and the need to "*borrow-to-hire*." We refer to this transmission mechanism as the *credit channel of public procurement* as argued by Gabriel (2022) in the context of Portugal.

Policymakers focused on macroeconomic stabilization may be less interested in the effects of procurement on individual firms and more concerned with aggregate regional outcomes, such as regional employment multipliers and their implied *cost-per-job* per year. While positive firm-level estimates are valuable for providing direct evidence of underlying mechanisms (e.g., the credit channel), general equilibrium forces can potentially turn the multiplier negative. For instance, changes in factor prices (Barattieri, Cacciatore, and Traum (2023)) or the reallocation of workers from non-contractors to contractors could offset positive effects.

Widely accepted measures of fiscal shocks in the U.S. at the metropolitan statistical area (MSA) level (Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b)) yield significant and positive estimates of the cost-per-job per year, around $548,000. These estimates suggest positive net effects of contracts on employment but are notably higher than those associated with other types of government spending in the U.S., such as ARRA transfers (Chodorow-Reich (2019)). Why does this discrepancy exist? When aggregating our new measure of unanticipated contracts at the regional level, we find a much lower cost-per-job—$57,000 per year—closer to those for ARRA transfers. We propose that this difference is likely due to the more unexpected nature of these contracts compared to typical procurement contracts, which are often awarded via task and delivery orders under longer-term agreements known as indefinite delivery vehicles (IDVs). The effects of IDVs are more dispersed over time, making them harder to measure.

Even if estimates of employment multipliers are positive and significant, ruling out negative net crowding-out effects, it remains unclear how much of the employment channel arises from direct effects on contractors versus indirect multiplier effects. For policymakers aiming to maximize the employment impact of contracts, understanding the origin of this mechanism is crucial. Specifically, if direct effects on awardees drive most of the impact, policymakers should focus on identifying the types of contracts and firms that elicit the strongest responses. For instance: How do fixed-price contracts compare with cost-plus-fee contracts in terms of implied employment multipliers? Do longer contracts produce substantially higher multipliers? Conversely, if spillover effects dominate, targeting more responsive regions or sectors may be more effective. For example, do regions with higher shares of (i) hand-to-mouth consumers, (ii) input suppliers to prime contractors, or (iii) credit-constrained firms generate higher multipliers?

Answering these specific questions goes beyond the scope of this paper. Instead, we quantify how generally accepted measures of regional contract shocks (Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b)) are divided between direct and indirect effects. First, we construct regional time series for contractors' and non-contractors' employment. This is achieved by leveraging restricted access to the QCEW, allowing us to match the full set of contractors with all U.S. firms and decompose employment into contractor and non-contractor components. Second, using this newly constructed panel dataset, we find weak evidence of employment crowding-out from non-contractors in the initial period (impact multiplier), although the estimate is statistically insignificant. Furthermore, we estimate that direct effects on contractors account for 55% of the employment multiplier, while indirect effects on non-contractor firms account for the remaining 45%, on average. To the best of our knowledge, we are the first to provide a clear percentage breakdown of fiscal multipliers into direct and indirect effects.

Related Literature

This paper relates to the growing literature studying the effects of procurement spending on firms: Ferraz, Finan, and Szerman (2021) (Brazil), Gabriel (2022) (Portugal), Giovanni et al. (2023) (Spain), and Lee (2024) (South Korea). For the U.S. case, Hebous and Zimmermann (2020) examine the response of investment to contracts using a set of publicly traded firms, while Juarros (2022) estimates the effects of state-level shocks on a set of non-contractors from Orbis. We contribute to this literature by providing direct evidence of the positive effects of contracts on employment growth for contract winners in the U.S., leveraging (i) restricted data on the universe of firms from QCEW and (ii) a newly identified set of "*unanticipated*" contracts from FPDS.

The second part of this paper relates to the extensive literature studying the effects of procurement spending across regions and industries in the U.S.: Perotti (2007) and Nekarda and Ramey (2011) (industry-level effects); Nakamura and Steinsson (2014) and Dupor and Guerrero (2017) (state-level multipliers); and Demyanyk, Loutskina, and Murphy (2019), Auerbach, Gorodnichenko, and Murphy (2020b), Cox et al. (2024), Barattieri, Cacciatore, and Traum (2023), Juarros (2022), and Muratori, Juarros, and Valderrama (2023) (effects of contracts using FPDS). We are particularly related to the work of Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b), who also study employment multipliers of defense procurement spending at the MSA level during a similar time period.

We contribute to this literature by demonstrating that current estimates of cost-per-job per year implied by employment multipliers derived from defense contracts in the U.S. are an order of magnitude higher than those associated with other types of government spending. Using our novel measure of unanticipated contracts, we provide estimates that align with the broader fiscal policy literature.

This paper also contributes to the fiscal policy literature on cross-sectional employment multipliers and procurement spending in other countries: Serrato and Wingender (2016) (U.S. Census shocks), Adelino, Cunha, and Ferreira (2017) (U.S. municipalities), Corbi, Papaioannou, and Surico (2019) (Brazilian municipal transfers), Chodorow-Reich (2019) and Choi, Penciakova, and Saffie (2023) (ARRA transfers and political connections), Pinardon-Touati (2022) (crowding-out of credit in France), and Buchheim and Watzinger (2023) (German public investments).

Overall, we contribute to the fiscal policy literature by providing the first detailed breakdown of employment multipliers into direct and indirect effects. By disentangling the responses of firms that win contracts (contractors) from those that never win them (non-contractors), we identify a 55-45% split between direct and indirect effects, respectively.

The rest of the paper is organized as follows. In Section 2, we discuss the institutional background of federal procurement spending in the U.S. and the datasets used in the analysis. We then describe our empirical specification and present the firm-level results in Section 3. In Section 4, we link procurement contracts with balance sheet and loan data from the Federal Reserve and illustrate a credit channel transmission mechanism. Section 5 studies the aggregate effect of procurement contracts and provides a breakdown of the regional employment multiplier into contractor and non-contractor components. Finally, Section 6 concludes.

II. Institutional Background and Data

II.a The Federal Procurement Process

Federal procurement spending refers to any good or service purchased from private entities by federal agencies. Figure 1 shows the time series of federal procurement spending as measured by the National Income and Product Accounts, or NIPA (blue line), as shares of total government spending, G (left panel) and GDP (right panel) by fiscal years.

Over the last 20 years, federal procurement spending has accounted for 16% of total government spending (G) and about 3% of GDP, on average. Given its size and its direct effect on U.S. private companies, federal procurement represents a direct effective method to implement fiscal

Figure 1: Federal Procurement Spending Shares

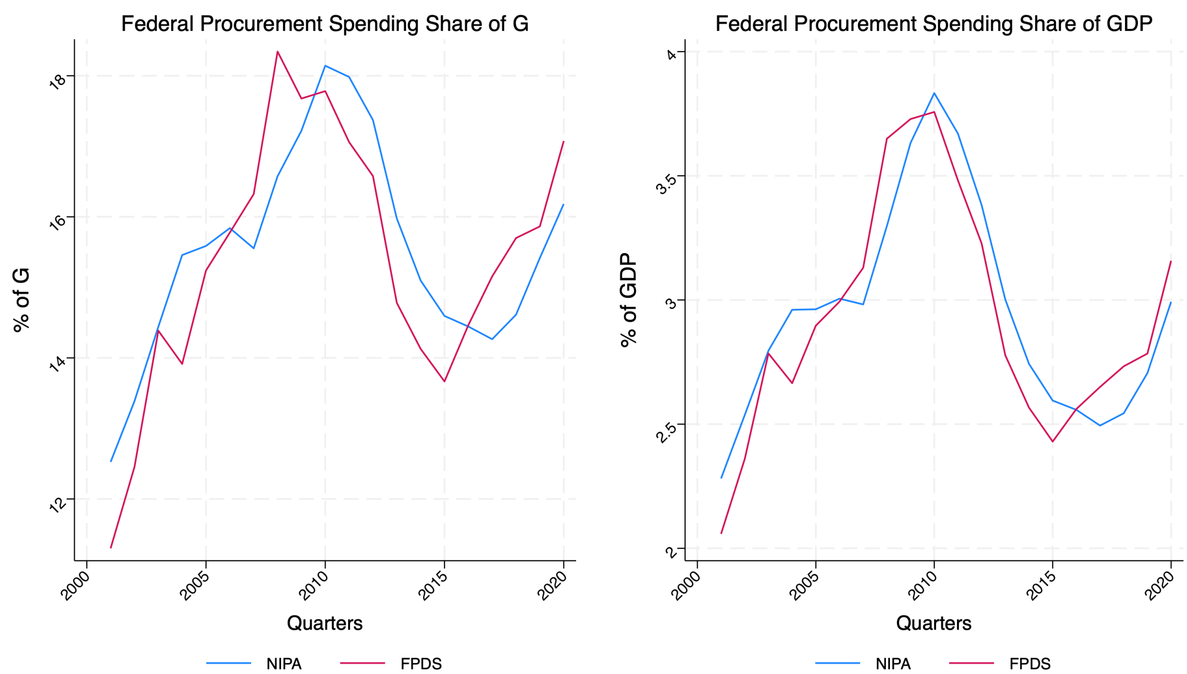

Notes: Federal procurement spending is calculated by adding up (i) NIPA federal government intermediates goods and service purchased and (ii) NIPA federal government gross investment in structure, equipment and software (see Cox et al. (2024) and Briganti and Sellemi (2023)). Federal spending is the summation of defense and non-defense spending.

policy.[1]

Starting from fiscal year 2001, the universe of federal procurement contracts is publicly available at USASpending.gov. The data is sourced from the Federal Procurement Data System (FPDS), the software used by federal contracting officers to input information of any federal action.[2] Data from FPDS is aggregated by fiscal year and plotted in red in Figure 1. Notice that the two data sources of federal procurement spending, (i) FPDS contracts and (ii) NIPA, match well, providing an incredibly detailed micro-origin of federal procurement spending, allowing researchers to study

[1]More than 90% of U.S. procurement spending originates from contracts awarded with a primary place of performance in the U.S. See Appendix B.5 for details.

[2]The same data used in Demyanyk, Loutskina, and Murphy (2019), Auerbach, Gorodnichenko, and Murphy (2020b), Hebous and Zimmermann (2020), Juarros (2022), Cox et al. (2024), Muratori, Juarros, and Valderrama (2023), Barattieri, Cacciatore, and Traum (2023) and Briganti, Brunet, and Sellemi (2025).

Figure 2: FPDS Contracts Breakdown

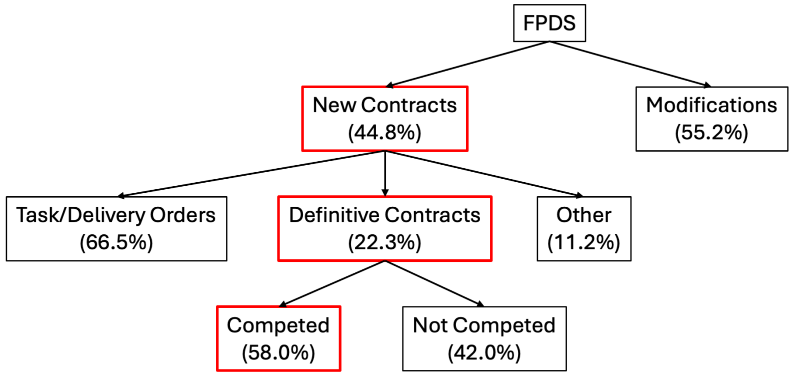

Notes: Data refers to averages over fiscal year shares. Fiscal years span from 2006 to 2019 to maintain consistency with the rest of the analysis in this paper. Shares refer to contract values, not the number of contracts (i.e., dollar-weighted shares).

the effects of procurement spending at highly disaggregated levels.

Overview of Federal Contracting. As noted in Auerbach, Gorodnichenko, and Murphy (2020b), behind each government contract lies a long history of transactions, with significant heterogeneity in the types of contracts awarded. While regional-level analyses can overlook individual contract details due to higher aggregation level, firm-level analyses of the effects of contracts cannot disregard this aspect. Figure 2 presents a breakdown of federal contracts by the most common categories.[3]

First, only 45% of all transactions in FPDS account for newly awarded contracts. The remainder refers to contract modifications, such as options, extra work, or administrative actions, all related to an existing contract.

Second, not all newly awarded contracts are necessarily "new." In fact, 66.4% of all new contracts is represented by task orders (for services) and delivery orders (for goods) issued under a pre-existing parent contract, called an indefinite delivery vehicle (IDV).[4] It is noteworthy that FPDS

[3] We are deeply thankful to a federal government contracting officer, who preferred to remain anonymous, for assisting us in sorting out the details of each contract type.

[4] IDVs are regulated by the Federal Acquisition Regulation (FAR) 16.5. Specifically, an IDV serves as a mechanism awarded to one or more vendors, streamlining the provision of supplies and services. This method is particularly advantageous for handling both expected and unforeseen needs, simplifying the procurement process by eliminating the need for a new solicitation for each task or delivery order and reducing the paperwork for these orders, among other

reports the competition details and bid counts from a parent IDV to its subsequent task/delivery orders, even if these orders were not competed. As a result, many contracts seemingly reported as "newly competed" in FPDS are neither newly awarded (since their parent IDV may have been awarded months or even years earlier) nor genuinely competed, as they represent potentially anticipated government demands granted by the originally competed IDV. Conversely, only 24.3% of new contracts represent standalone contracts, which are not part of any ongoing relationship with the government. These are technically referred to as "definitive contracts."

Third, contracts can be awarded either competitively or non-competitively. Non-competitive contracts are mainly for complex products, for which agencies often prefer to award "sole-sourced" contracts. For instance, products or services might be deemed available from a sole source if that source offers unique and innovative concepts or proposes a concept or service unavailable from other providers.[5] Figure 2 shows that 43.5% of newly awarded definitive contracts fall into this non-competed category.

The Competition Process. Contracts awarded competitively are subject to a public solicitation procedure with the goal of increasing competition.[6] In particular, since October 1, 2001, contract actions with an expected value of over $25,000 must be publicized in an online government platform, `sam.gov` (SAM).[7]

Using SAM data on the universe of federal contracts' notices from fiscal year 2006 to 2019, we construct the entire pre-award to award history of competed contracts: for example, from the oldest pre-solicitation date to the award notice date. We find that the median number of days from the oldest pre-solicitation to an award notice is 20 days, while for 75% of solicited contracts, the number of days is less than 52.[8] More details on contracts solicitation are available in Online Appendix A.1 and A.2.

benefits.

[5]See FAR 6.302-1-a.

[6]See FAR 5, *Publicizing Contract Actions*.

[7]Contract actions below the threshold might still be posted to increase visibility.

[8]We follow indications from Gonzalez-Lira, Carril, and Walker (2021) to work with solicitation data. We thank Andres Gonzales-Lira for directing us to the General Services Administration Technical Documentation for the FedBizOpps (FBO) website, whose information migrated to Contract Opportunities (SAM).

In response to contract solicitations, interested vendors submit an offer in the form of either a bid or a competitive proposal, depending on the nature of the acquisition procedure. Table 1 provides some summary statistics of the competition procedure of new definitive contracts.

Table 1: Solicitation Procedures of New Definitive Competed Contracts

Solicitation Procedure	*Fraction of Total Value*	*Number of Offers Received*			
		Q1	*Median*	*Q3*	*Mean*
One-step Sealed Bidding	18.90%	3.07	4.36	7.86	6.81
Two-step Sealed Bidding	5.90%	2.29	3.86	7.86	12.29
Competitive Proposal	67.40%	1.57	3.14	6.57	27.35
Other	7.80%	-	-	-	-

Notes: Values refer to fiscal year averages from 2006 to 2019. Fraction of dollar value is calculated relative to the total value of new, definitive, and competed contracts. `q1` and `q3` refer to the first and third quartile, respectively. Detailed information about the solicitation procedure can be found at this link.

Two thirds of the total value of new definitive competed contracts are awarded via competitive proposals, receiving a median number of offers equal to 3.14. One fourth of these contracts are awarded using auctions, either one- or two-step sealed bidding, with a median number of offers received of about four offers per contract.

Unanticipated Contracts. In light of these considerations, we select a specific subset of contracts that we believe are reasonably unanticipated. We build on the approach of Hebous and Zimmermann (2020), but we impose even more restrictive conditions. In particular, we focus on contracts that meet four conditions: (i) new, (ii) definitive, (iii) competed contracts, and (iv) contracts that received at least two offers.

Contracting officers have indicated that even a single-offer scenario, if open to full competition, is treated as competitive, since it potentially pressures the bidder to refine their proposal in anticipation of additional bids. However, they have also revealed that the number of bids is a good indicator of competitiveness, thereby ensuring the unanticipated nature of definitive contract awards. Overall, competition significantly diminishes the ability of firms to anticipate contract awards, especially when the bid count is high. Therefore, we further restrict our attention to contracts that have received at least two offers, consistent with the approach in Hebous and Zimmermann (2020).

It turns out that only 6% of total federal spending meets conditions (i) through (iv). We refer to this set of contracts as "*unanticipated contracts*." The median unanticipated contract is worth $114,900. The top categories of services purchased via unanticipated contracts are construction services and defense R&D, while the top categories of goods are food products and manufacturing goods related to defense hardware. The median duration of an unanticipated contract for a service is 283 days, while the median duration of an unanticipated contract for the purchase of goods is much shorter: 79 days.

Detailed descriptive statistics of unanticipated contracts, as well as additional information on their geographic distribution and sample coverage, are available in Online Appendix B.4.

II.b Data

Procurement Contracts. We use the Federal Procurement Data System (FPDS) to obtain the universe of all federal procurement contracts awarded from 2006 to 2019, at a daily frequency (including both defense and non-defense spending).[9] We focus on firms that have won at least one unanticipated contract during this period, identifying approximately 80,000 unique firms, identified using Dun & Bradstreet's data universal numbering system (DUNS). Large contractors, such as Lockheed, report different DUNS numbers for each specific subsidiary, resulting in DUNS being location-specific. For example, 97.5% of all defense contractors in the FPDS dataset are located in only one MSA, and 99.7% are located in no more than two MSAs.[10] For instance, Lockheed uses different DUNS for its Fort Worth-Dallas subsidiary, which appears as a separate firm.

We aggregate all contracts at the level of recipient firms (identified by DUNS number) by quarter, resulting in an unbalanced panel dataset. We break down contracts into two components. First, unanticipated contracts, which we treat as shocks, meeting conditions (i) through (iv); second, potentially anticipated contracts, which include all contracts not meeting all four conditions.

[9]Choice of years: we exclude 2020 to avoid the COVID-19 period; we start from 2006 because FPDS information appears to be more stable and fields are more populated.

[10]This statistic is constructed using all recipient DUNS numbers that received at least one contract from the Department of Defense from 2006 to 2019.

Firm-Level Outcomes. We leverage restricted data access to the Longitudinal Database of Establishments (LDBE), compiled by the Bureau of Labor Statistics (BLS). The data source is the Quarterly Census of Employment and Wages (QCEW), which collects data quarterly from Unemployment Insurance Tax agencies in all states. The data covers the universe of establishments in the 42 states that agreed to provide access for this project.[11]

This dataset provides comprehensive monthly employment and quarterly wages at the establishment level. It also includes information on the establishment's name, location (state, county, and town) and primary industry (six-digit NAICS). LDBE's employment count includes only filled jobs, whether full or part-time, temporary or permanent, by place of work. LDBE's wage measure is comprehensive and, therefore, wages can be considered total payments to employees.[12]

We identify firms in LDBE using the Employer Identification Number (EIN), which allows us to aggregate data across multiple establishments. In the 42 states analyzed, 96% of firms have only a single establishment within a state. However, firms may also have establishments in multiple states. To avoid cross-state spillovers, we focus on firms with a single establishment within a state. This simplification does not exclude large multi-establishment contractors from the sample, such as Lockheed Martin, which reports different EINs for its different establishments/subsidiaries.

Matching Firms with Contractors. We merge both the previous contracts panel (from FPDS) with the firm panel (from LDBE) using a string-matching algorithm.

We successfully match 13,000 firms that have received at least one unanticipated contract between 2006 and 2019.[13] Additionally, we eliminate firms that (i) have gaps in their time series, (ii) receive their only contract shock in their first four observed quarters, so we cannot control for four lags of employment, (iii) receive their only shock in their last two observed years, so we cannot estimate the full impulse response function for that firm, and (iv) have fewer than one employee on average. This leaves us with 5,317 firms.

[11] The states for which we do not have access are Florida, Kentucky, Massachusetts, Mississippi, New York, North Carolina, Rhode Island and Vermont.

[12] Wages include bonuses, stock options, severance pay, profit distributions, the cash value of meals and lodging, tips and other gratuities. Source: BLS website.

[13] We exclude 2020 from our analysis to avoid COVID-related contracts and employment dynamics.

We further divide our set of firms into two groups: (i) small firms, if the initial observed number of employees is fewer than 150, and (ii) large firms, if more than 150.[14] We identify 175 large federal contractors, while the remaining 5,142 firms are classified as small firms.

Controlling for sector and state fixed effects, we find that in-sample small firms have a statistically significant higher number of employees but do not pay different average wages. In contrast, in-sample large firms do not show any statistically significant difference from their out-of-sample counterparts. Detailed information about our sample of shocked firms and a discussion of external validity is available in Online Appendix C.

Finally, we observe around 10,800 shocks in our sample, meaning that each firm receives, on average, two unanticipated contracts over the observed time period. The average contract size is $700,000, while the third quartile of the contracts' size distribution is about $4,000,000.

III. Effects of Contracts on Firms

III.a Identification via Unanticipated Contracts

Most government contracts cannot be treated as quasi-random shocks. The majority of procurement spending takes place in the context of long-term agreements (e.g., IDVs) whose timing may be anticipated well in advance by the awardee. Similarly, many awardees are selectively chosen by contracting officers using non-competitive acquisition procedures. Furthermore, many contracts in the FPDS are merely modifications of existing agreements rather than new orders.

Thus, estimates that simply compare the growth of firms after receiving an award to the growth of other firms that did not receive an award would suffer from two forms of bias: selection bias, i.e., the firms winning awards may be positively selected and thus display higher counterfactual growth than comparison firms; and anticipation, i.e., when the awards are anticipated, they affect the firm prior to the award date, so the impulse response function does not fully capture the impact of the award.

[14] Shocked firms are classified by their employment in the period before the arrival of the shock.

We address each of these concerns through our empirical strategy. First, we study the effects of unanticipated contracts, introduced in Section 2. Unanticipated contracts isolate a subset of contracts that are (i) newly awarded (not modifications of existing agreements), (ii) standalone contracts, i.e., definitive contracts (not part of an ongoing series of purchases), (iii) competed, and (iv) have at least two bidders.[15] Conditions (iii) and (iv) are similar to those imposed by Hebous and Zimmermann (2020); conditions (i) and (ii) impose additional restrictions. As mentioned in Section 2, we refer to contracts satisfying conditions (i) through (iv) as "*unanticipated contracts*." Only 6% of federal procurement spending meets these four conditions.

Selection Bias. In the context of federal purchases, Nekarda and Ramey (2011) were the first to highlight that industry technological progress can endogenously drive medium-term changes in industry-level government purchases (Perotti (2007)), i.e., reverse-causality via selection bias. Indeed, government purchases driven by technological progress not only occur frequently, but they are specifically regulated by FAR: sole-source acquisition procedures (FAR 6.302-1-a). Our empirical strategy addresses this type of concern.

First, conditions (iii) and (iv) ensure that unanticipated contracts are not awarded because of a firm offering a new, innovative product, since these types of acquisitions fall in the non-competed category.

Second, we control for firm fixed effects in growth rates. Thus, the estimated impacts are not based on differences in average growth rates between firms but rather on whether a firm grows faster than its own average growth rate in the periods after winning an unanticipated contract. We also control for lags of the outcome variable and for industry-time and location-time fixed effects, so the results are not driven by other location- or industry-specific shocks that may be correlated with unanticipated contracts.

Third, we include only firms that are awarded an unanticipated contract at some point in our sample. This ensures that the comparison firms for a given shock are other firms that receive similar shocks at other times, which are arguably a more comparable sample than the full set of firms.

[15] We show in robustness tests that results are largely the same when imposing stricter limits on the number of bidders.

Of course, this comes at the cost of external validity; however, the aggregate value of contracts awarded to firms that win at least one unanticipated contract accounts for a large fraction of total federal procurement spending. This means that recipients of unanticipated contracts are not just a small special subset of contract recipients, but rather are representative firms within the universe of federal contractors.

Anticipation. Firms might anticipate the effects of contracts for two reasons: first, the contract is not a new one or is part of an existing ongoing relationship with the government; second, firms learn about a contract opportunity well ahead of the award notice and anticipate winning the contract. Hebous and Zimmermann (2020) also highlight the problem of contract anticipation by large public firms and show that future competed contract awards do not cause any effect on the current stock prices of the future awardee.

In the context of smaller private firms, we address concerns about anticipation by focusing solely on the effects of new standalone contracts that have been highly competed. Second, in Online Appendix D, we show that these contracts behave like one-time shocks: total government contracts jump on impact 1:1 with the size of the shock contract. Therefore, the timing of the shocks does not coincide with receiving other, potentially anticipated contracts. Moreover, there is no evidence of either anticipation or persistence, i.e., being shocked does not predict subsequent government spending at the firm to any meaningful extent.

Third, we verify that the median number of days between when firms learn about the existence of a contract opportunity (i.e., pre-solicitation) and the award notice is just 20 days. Therefore, any potential anticipation behavior is not measured at a quarterly frequency, which is what we use in our study. In the next section, we also conduct additional tests to directly rule out any potential anticipation behavior.

III.b Effect of Unanticipated Contracts on Firms' Growth

We use panel local projections à la Jordà (2005) to estimate the effect of \$1 of unanticipated contracts on employment, total wages and average wage. In particular, we estimate via OLS the following baseline equation for either small or large firms:

$$
\begin{aligned}
Y_{i,t+h} - Y_{i,t-1} = \beta^h \cdot \underbrace{\varepsilon^G_{i,t}}_{\text{Shock}_t} + \gamma^h_0 \cdot \tilde{G}_{i,t} + \ldots \\
\ldots + \underbrace{\sum_{j=1}^{4} \left\{ \rho^h_j \cdot \varepsilon_{i,t-j} + \gamma^h_j \cdot \tilde{G}_{i,t-j} + \phi^h_j \cdot (Y_{i,t-j} - Y_{i,t-1-j}) \right\}}_{\text{Lags (Control for serial correlation and pre-trends)}} + \ldots \\
\ldots + \underbrace{\alpha^h_i + \alpha^h_{s,t} + \alpha^h_{\ell,t}}_{\text{Fixed Effects}} + \varepsilon_{i,t+h} \quad h = 0, 1, \ldots, H, \qquad (1)
\end{aligned}
$$

where $Y_{i,t+h}$ denotes the h-period ahead outcome variable. $\varepsilon^G_{i,t}$ denotes the value of unanticipated contracts awarded to firm i in quarter t. $\tilde{G}_{i,t}$ indicates the dollar value of potentially anticipated contracts, that is, contracts which do not meet conditions (i) through (iv). The sum of unanticipated and potentially anticipated contracts corresponds to the total value of procurement contracts awarded in a given quarter to a firm and is denoted by

$$
G_{i,t} = \underbrace{\varepsilon^G_{i,t}}_{\text{Unanticipated}} + \underbrace{\tilde{G}_{i,t}}_{\text{Potentially Anticipated}} .
$$

To mitigate serial correlation, we control for four lags of the shock, four lags of potentially anticipated contracts and four lags of quarterly outcome's past changes. Nominal variables are deflated using the GDP price deflator.

Moreover, we include several fixed effects in our analysis. First, α^h_i represents a firm fixed effect designed to control for firm-selection bias and time-invariant firm characteristics; for instance, firms that are more productive may experience faster growth and win more contracts. This firm fixed effect isolates the variation in contracts awarded to a firm over time.[16]

[16] One limitation of the firm fixed effect is its inability to account for reverse causality, such as a firm receiving more

Second, $\alpha^h_{s,t}$ is a sector-time fixed effect intended to absorb any sectoral business-cycle effects. Essentially, if a sector is experiencing growth in a particular year due to breakthroughs, and contractors are winning more federal contracts as a result, the significance of β^h could erroneously attribute this growth to procurement effects rather than the underlying sectoral boom. Lastly, $\alpha^h_{\ell,t}$ represents a location-time fixed effect, capturing regional business-cycle effects within a state.

Baseline Results. We are interested in β^h, the effect of a dollar in extra federal procurement on firm-level outcome changes over any horizon h. We estimate Equation (1) via OLS for both small firms, i.e., initial level of employment is less than 150 employees, and large firms for each horizon h from 0 (impact) to 8 (two years). Firms' size is measured at the beginning of the sample (i.e., pre-treatment). The OLS estimates of β^h can be interpreted as impulse response functions (IRF) of the effect of an extra dollar of spending. Figure 3 shows the results for employment for both small (left panel) and large (right panel) firms.

The left panel shows the IRF for small firms. It displays a significant and positive effect. The maximum response to changes in employment from the shock occurs at horizon 2, where the response is 0.21, which means that after one year from the shock, \$1M of contract generates 0.21 more jobs, on average.[17] Following Chodorow-Reich (2019), we calculate the number of job-years by cumulating the impulse response function and then dividing by four, since our data are measured at quarterly frequency. We obtain a value of 0.322, which corresponds to a cost-per-job (with average duration of one year) of \$3,110,000.

Moreover, notice that the effects of an unanticipated contract are quite persistent and appear to survive even after five quarters from the shock. Since 75% of unanticipated contracts for services have a duration shorter than five quarters while 75% of unanticipated contracts for goods have a duration shorter than three quarters (see Online Appendix B.3), the effects of contracts appear to be very persistent and survive even after the termination of the contract. The persistence of the effects of procurement contracts is consistent with previous findings in the literature (Ferraz, Finan, and

contracts due to increased productivity. To address this, we examine a set of quasi-randomly assigned contracts.

[17] We divide contracts by 1,000,000.

Figure 3: Employment – Small vs Large Firms

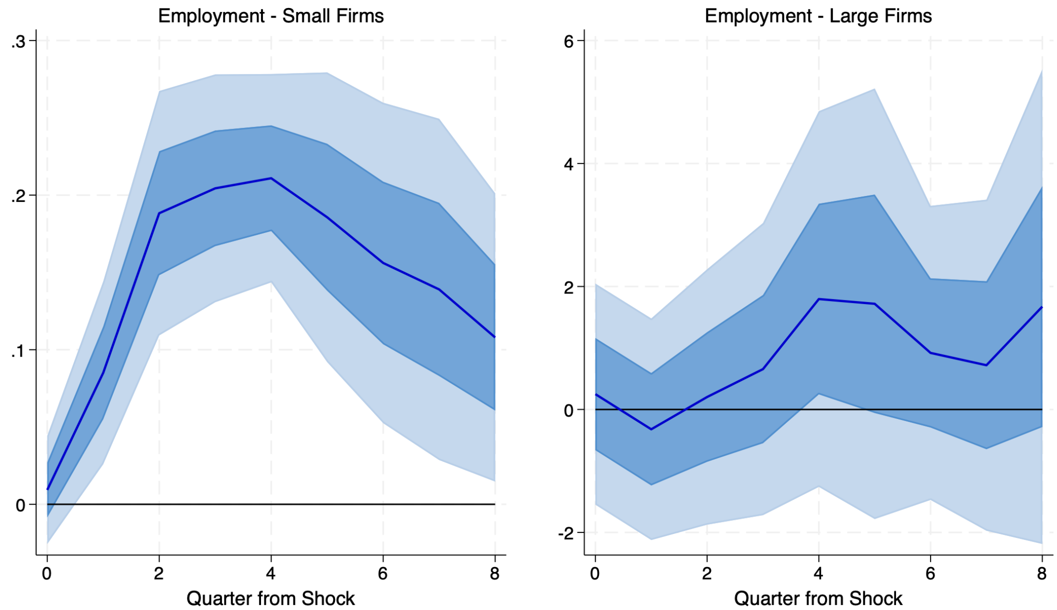

Notes: Firms are observed from 2006:1 to 2019:4, i.e. $T = 56$. Number of small firms is $N = 5,142$, while the number of large firms is $N = 175$. Standard errors are clustered at the state level. Small bands are 68% confidence. Large bands are 95% confidence.

Szerman (2021) for Brazil, Lee (2024) for South Korea, and Gabriel (2022) for Portugal).

In contrast, the response of large firms' employment (right panel) is positive but not significant. The implied cost-per-job (with average duration of one year) is $525,482. The lack of responsiveness of large firms to procurement contracts is consistent with the findings in the literature: Hebous and Zimmermann (2020) in the case of public U.S. firms and Gabriel (2022) using procurement data from Portugal. Additionally, Giovanni et al. (2023) find that firms with fewer fixed assets (a proxy for borrowing constraints), which are likely smaller, experience a larger crowding out of private sales using Spanish data. The literature generally agrees that the existence of credit constraints on small firms accounts for a significant part of the effects of procurement contracts.

In what follows, we implement several robustness checks.

III.c Robustness

Placebo. First, we carry out a placebo test to rule out the possibility that our specification is picking up some unknown source of spurious correlation. We do so by reshuffling the timing of the shocks within each firm. The new synthetic shock for firm i in quarter t is denoted by

$$\forall (i,t)\ \hat{\varepsilon}_{i,t} = \varepsilon_{i,\tau} \quad \text{with } \tau \in \{\text{2006:1,...,2019:4}\}.$$

We then re-estimate Equation (1), replacing the original shocks, $\varepsilon_{i,t}$, with the synthetic shocks, $\hat{\varepsilon}_{i,t}$, to carry out a placebo test. If our specification is capturing a spurious correlation instead of a causal effect of contracts on employment, we would expect to see positive and significant results even in response to synthetic shocks. On the contrary, the placebo test is passed only if the synthetic shocks have no significant effect on employment.

The right panel of Figure 4 shows the estimate of β^h using synthetic shocks and compares it to the original estimates displayed in the middle panel of the same figure.

Figure 4: Employment – Small Firms – Anticipation and Placebo Tests

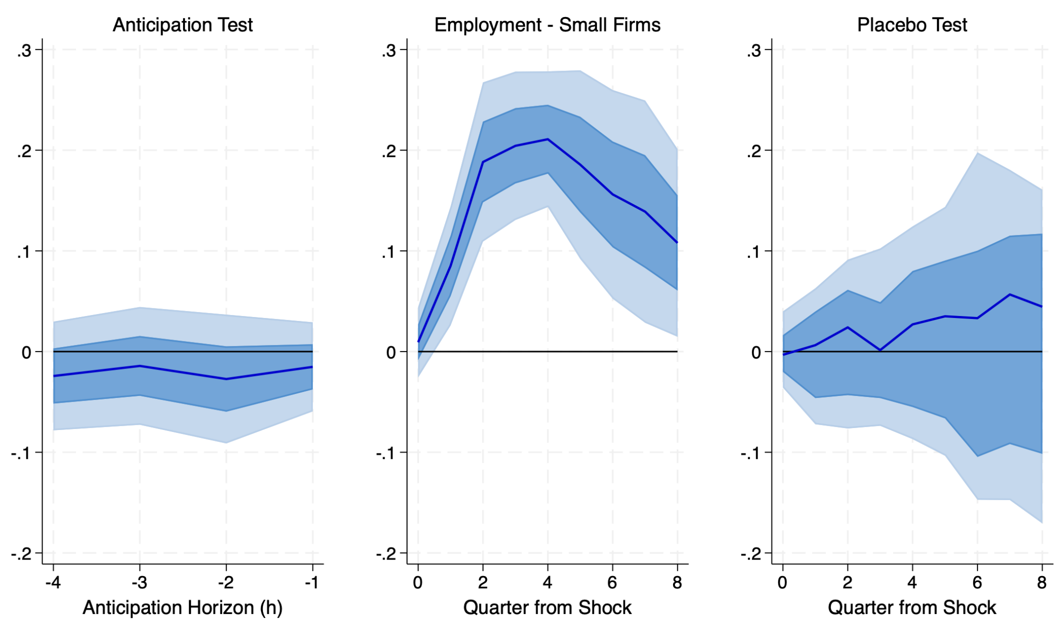

Notice that the synthetic shocks do not produce any significant effect on employment growth, suggesting that the original shocks are indeed capturing the causal effect of (unanticipated) contracts on firms' employment growth.

Anticipation Test. Second, we carry out an anticipation test to rule out the possibility that our shocks are anticipated. This test is motivated by two potential concerns. First, firms might be capable of predicting the win of a newly definitive competed contract that received several offers. Second, as mentioned in Section 2, contracts are solicited several days before the award notice, creating a potential anticipation horizon in those cases where the solicitation occurs more than a quarter before the award notice.

Therefore, we re-estimate Equation (1) for $h = 0$ while shifting forward the shocks $\varepsilon_{i,t}$ by either one, two, three or four periods. In particular, we are interested in the effect of future shocks $\varepsilon_{i,t+\tau}$, with $\tau = 1, \ldots, 4$, on current employment changes: $Y_{i,t} - Y_{i,t-1}$. If shocks were anticipated, we would expect to see a significant effect of future shocks on current employment growth, with a magnitude similar to that observed for current shocks (middle panel). We report the OLS estimates of the effect of $\varepsilon_{i,t+\tau}$ on current changes in employment for $\tau = 1, \ldots, 4$ in the left panel of Figure 4.

Notice that future shocks have neither a meaningful nor a significant effect on current employment changes at any point in the anticipation horizon. The result of this test rules out the possibility that contracts, which are considered unanticipated, are in fact anticipated by firms.

Analysis by Quartile of Small Firms. Third, we further subdivide the sample of small firms by analyzing each quartile of the size distribution of small firms separately. Firms in the first quartile have between 1 and 6 employees, firms in the second quartile have between 6 and 13 employees, and firms in the third quartile have between 13 and 28 employees. The fourth quartile is characterized by much greater dispersion in the number of employees: while the first three quartiles range from 1 to 28 employees, the last quartile ranges from 28 to 150, thus including much larger firms.

Therefore, we re-estimate Equation (1) for each quartile of the firms' size distribution separately to explore the robustness of the result across the sample. Figure 5 shows the IRFs of employment growth for each quartile.

Figure 5: Employment - Quartiles of Small Firms

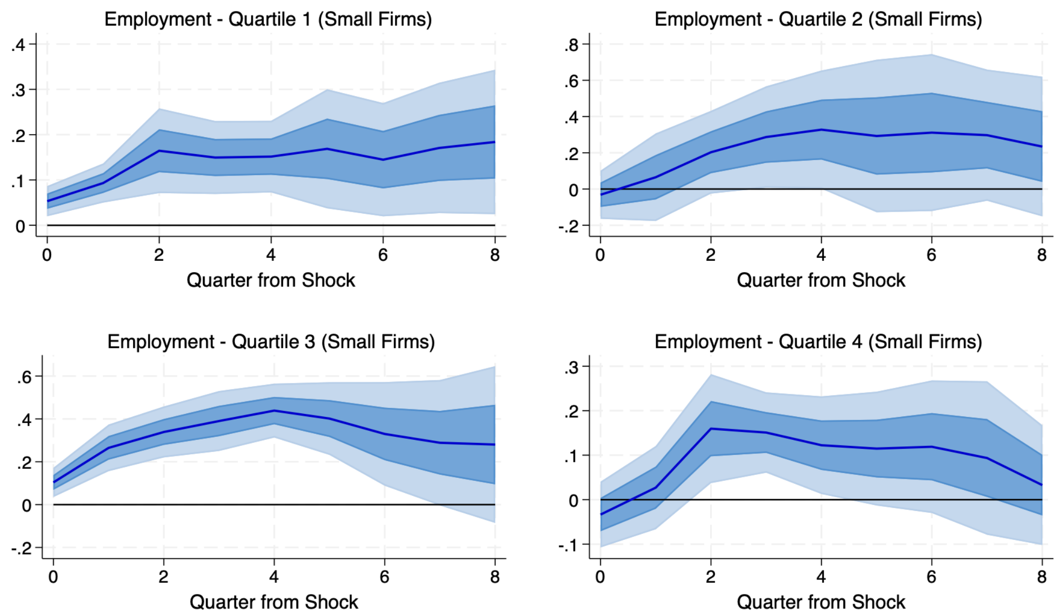

Note that the results appear to be robust across all four quartiles of the size distribution.

Lastly, we repeat both the anticipation test and the placebo test for each quartile of small firms. Both sanity checks are passed in all four quartiles of small firms. In the interest of brevity, the results are presented in the Online Appendix E.1.

Time-Varying Productivity Shocks. Fourth, we are concerned that firms might win contracts in response to temporary productivity shocks, which make them capable of outbidding their competitors and, consequently, outgrowing them (i.e., omitted variable bias). Note that our firm-fixed effects can remove the effects of systematic differences in productivity levels across firms, while our use of highly competed newly awarded definitive contracts rules out the possibility that contracts are awarded in response to the development of innovative products (i.e., sole sourcing). However, they are not capable of controlling for time-varying productivity shocks that make firms temporarily more productive.

Therefore, we re-estimate Equation (1) by augmenting the specification with four lags of wage-per-worker. According to Neoclassical theory, the marginal product of labor is equal to the real (product) wage. Consequently, changes in wage-per-worker should reflect changes in productivity levels. Thus, using lags of wage-per-worker enables us to control for time-varying productivity shocks.

The response of small firms' employment is robust to the inclusion of lags of wage-per-worker in the specification. In the interest of brevity, the results are presented in the Online Appendix E.2.

Response of Total and Average Wages. Fifth, we study the response of total wages and the average wage, or wage-per-worker, using the same specification as Equation (1).

We find no significant effect on wage-per-worker for either large or small firms. The same holds when we look at quartiles of small firms. According to Neoclassical theory, an increase in hours worked should decrease, *ceteris paribus*, the marginal productivity of labor, which, in turn, is proportional to the wage-per-worker. For example, Nekarda and Ramey (2011) find that industries receiving relatively more government spending pay lower wages. The fact that wage-per-worker does not fall in response to government spending shocks at the firm level is consistent with two potential explanations: (i) increasing productivity levels via learning-by-doing (e.g., Ilzetzki (2023)), and/or (ii) downward nominal wage rigidity (e.g., Born et al. (2023)). Discriminating between the two goes beyond the scope of this paper.

In light of the null response of wage-per-worker (i.e., $d\bar{w}_{i,t} = 0$) and the positive response of employment (i.e., $dE_{i,t} > 0$), we would expect to observe a positive response in total wages:

$$W_{i,t} := \bar{w}_{i,t} \cdot E_{i,t} \implies dW_{i,t} = \bar{w}_{i,t} \cdot \underbrace{dE_{i,t}}_{>0} + \underbrace{d\bar{w}_{i,t}}_{=0} \cdot E_{i,t} > 0.$$

In fact, the results indicate a significant and positive response of total wages for small firms, while responses for large firms are insignificant. In the interest of brevity, the results for both average and total wages are presented in Online Appendix E.3.

Goods vs Service. Muratori, Juarros, and Valderrama (2023) show that regional multipliers in the U.S. are driven by spending in services rather than goods. If large firms mainly produce goods and small firms specialize in services, we would then observe insignificant effects for large firms and significant effects for small firms. In this section, we show that our firm-level results are not driven by a service vs goods channel.

We follow Muratori, Juarros, and Valderrama (2023) and break down procurement spending into either service or goods spending. In our context of firm-level contracts, we break down (unanticipated) contracts into contracts for goods, i.e., $\varepsilon^g_{i,t}$, and (unanticipated) contracts for service, i.e., $\varepsilon^s_{i,t}$. Notice that these two categories are mutually exclusive:

$$\forall(i,t)\ \varepsilon_{i,t} = \varepsilon^g_{i,t} + \varepsilon^s_{i,t}.$$

In our dataset, firms that receive at least one unanticipated contract specialize in either goods suppliers or service providers. In fact, the distribution of the average share of government contracts for service is a highly bi-modal distribution with peaks at both zero, i.e., pure goods suppliers, or one, i.e., pure service providers. Moreover, we find that there are also many more service providers in the sample than goods suppliers because a larger fraction of procurement spending after 2006 goes for services rather than goods. In turn, this implies higher standard errors for the goods multipliers than those for service (see Online Appendix B.2).

We re-estimate Equation (1) by breaking down unanticipated contracts, i.e., $\varepsilon_{i,t}$, into a goods component and a services component, thus estimating the effects of the two separately. All the rest is identical to the baseline settings, where the equation is estimated via OLS separately for small and large firms. Results are shown in Figure 6.

There appears to be no difference between service and goods employment multipliers for large firms, with neither being statistically significant (right panel of Figure 6). In contrast, both the

Figure 6: Employment - Small vs Large Firms - Goods vs Service

Notes: Firms are observed from 2006:1 to 2019:4, i.e. $T = 56$. Number of small firms is $N = 5,142$, while the number of large firms is $N = 175$. Standard errors are clustered at the state level. Small bands are 68% confidence. Large bands are 95% confidence.

service and goods multipliers are positive and significant for small firms (left panel). This suggests that the type of product purchased, whether service or goods, does not explain the small versus large differences discussed in the baseline results, i.e., Figure 3.

IV. The Credit Channel of Public Procurement

Demyanyk, Loutskina, and Murphy (2019) use U.S. data to show that households' credit conditions matter for the size of regional fiscal multipliers. We now provide evidence on the credit channel of public procurement, which states that firms use procurement contracts to increase cash flow–based lending and increase access to bank credit. This credit channel is particularly relevant for smaller firms that borrow not only to invest but also to hire (Gabriel (2022)). Yet, there is no empirical evidence linking procurement contracts to bank credit in the United States. Hence, we tackle this with the innovative approach of linking procurement-level data with the Federal Reserve's supervisory data on bank credit.

Data source. Besides the procurement contract data presented before, our credit data source is Schedule H.1 of the FR Y-14Q that the Federal Reserve collects as part of the Comprehensive Capital Analysis and Review (CCAR) process. The latter dataset is quarterly, starts in 2013, and contains information on the loan portfolios of the largest banks in the United States which, in 2019, the last year of our sample, originated about 81% of all U.S. commercial and industrial lending and account for 86% of assets in the banking sector (Brown, Gustafson, and Ivanov (2021), Beyhaghi et al. (2024)). Banks are required to report all corporate loans and leases inclusive of all term loans and lines of credit with a committed balance greater than or equal to $1 million, which accounts for over 97% of these banks' corporate exposures (Beyhaghi (2022)). The dataset also contains detailed firm-level balance sheet information on borrowing firms updated annually, which we use to classify firms according to their size and categorize them into small and large firms.[18]

Data cleaning and merging. We merge the FPDS and Y-14Q databases relying on a name-matching exercise using the same method explained in Section 2.2. For this exercise, we only keep exact matches. Our final sample includes 770 firms with unique Taxpayer Identification Numbers (TIN) for which we have information on procurement, credit, and balance sheets. We rely on name matching again because there is no available crosswalk between DUNS and TIN. Regarding credit information, we rely on loan-level data that contains information such as banks' private assessments of loan risk, loan amount, interest rate, maturity, collateral, and guarantees. We then compile loan-level data into quarterly aggregates by adding all loan amounts to different banks in a given quarter and by computing the weighted average interest rate in the same quarter. As a last step and given the absence of employment data in the Y-14Q, we use the total assets of the matched firms to categorize them as small, if they are in the first 7 deciles of the size distribution, or large, if they are in the upper 3 deciles of the size distribution. This allows us to have a firm-level quarterly

[18] Detailed data instructions are available on the Federal Reserve's website at https://www.federalreserve.gov/apps/reportingforms/Report/Index/FR_Y-14Q.

dataset with credit information and procurement shocks.

Methodology and Results. We estimate Equation (1) via OLS for both small firms and large firms, for each horizon h from 0 (impact) to 4 (one year). We are again interested in β^h, which can be interpreted as the effect of a million dollars in extra federal procurement on firm-level outcome changes over any horizon h. Figure 7 shows the results for both utilized credit and average interest rates for both small (right panel) and large (left panel) firms.

Figure 7: Credit and Interest Rates – Small vs Large Firms

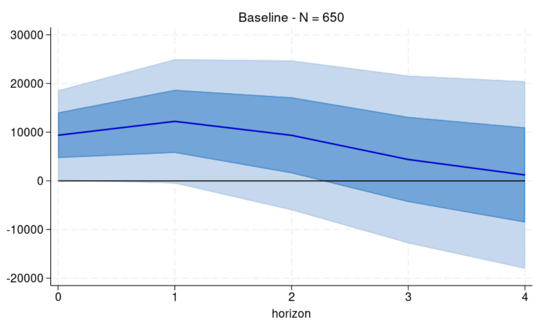

(a) Credit – Small Firms

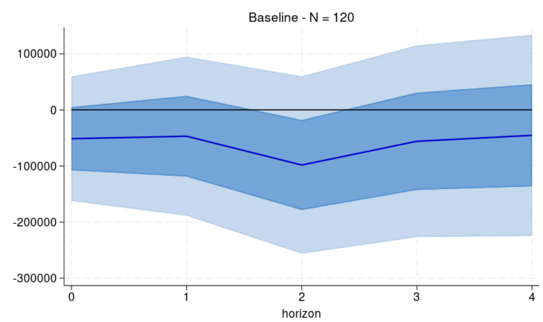

(b) Credit – Large Firms

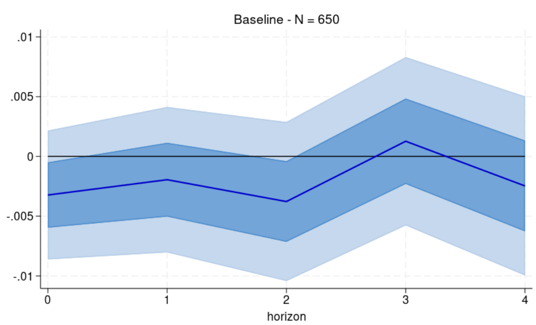

(c) Interest Rate – Small Firms

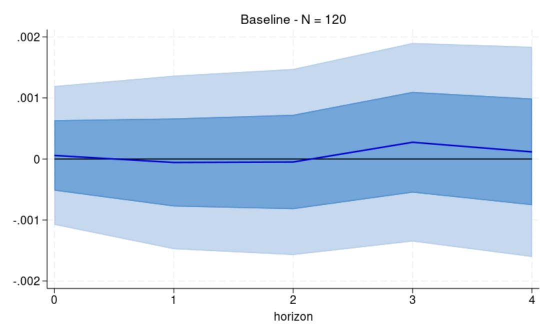

(d) Interest Rate – Large Firms

Notes: Firms are observed from 2013:1 to 2019:4, i.e., $T = 28$. Number of small firms is $N = 650$, while the number of large firms is $N = 120$. Standard errors are clustered at the state level. Small bands are 68% confidence. Large bands are 95% confidence.

Figure 7 shows that for small firms an unexpected contract amounting to USD 1M leads to a USD 10,000 increase in total credit available to the firm and around a 0.3 p.p. decline on average interest rates. We can thus conclude that small firms tend to increase borrowing following the announcement of an unexpected procurement contract. This pattern supports the notion that small firms are typically more financially constrained than large firms that have easier access to financial markets and thus view procurement contracts as opportunities to improve their access to credit.

The lower interest rates these firms experience upon securing procurement contracts likely reflect the improved risk profile perceived by lenders, due to the steady cash flows associated with government contracts and the associated positive reputation effect. First, procurement contracts may help resolve information asymmetries by serving as a government endorsement, signaling to banks that the firm is a reliable borrower and prompting broader credit access at competitive rates—similar to the certification effects observed in other government certification programs such as the one analyzed by Bonfim, Custódio, and Raposo (2023). Second, the procurement contracts themselves may act as collateral, strengthening a firm's credit profile and incentivizing lenders to offer more favorable credit terms (Gabriel (2022)). These mechanisms, whether through signaling or collateral enhancement, highlight the unique ways procurement awards reduce financing costs, particularly for smaller, constrained firms.

In opposition, the results also indicate that large firms show minimal or non-significant responses in both credit uptake and employment when they secure contracts. This discrepancy may stem from the fact that large firms often have more diversified sources of capital and less dependence on a single source, like government contracts, for their operational or growth financing. Consequently, large firms are not as financially constrained, which limits the incremental impact of procurement-induced credit.

Taken together, these findings reinforce the hypothesis that the credit channel is particularly important for small firms, which then see the most statistically significant effect on employment. Given the increased credit access (and lower interest rates) upon receiving these unexpected contracts, small firms can alleviate financing constraints, enabling them to hire in response to demand expansions from the government. The lack of similar credit responses in large firms ties to their insignificant employment response as well. Without a meaningful financial accelerator effect through borrowing, large firms likely do not expand their workforce significantly in response to procurement contracts.

Notice that standard Neoclassical theory would predict the opposite: when $G \uparrow$, labor supply also increases (i.e., negative income effect); therefore, the marginal product of capital increases, too, putting upward pressure on interest rates. This credit channel could in part explain why, historically, fiscal shocks in the U.S. have not increased interest rates (Jørgensen and Ravn (2022), Murphy and Walsh (2022)). In demand-driven output models, fiscal expansions increase aggregate income, which not only boosts the demand for credit but also raises its supply. The resulting balance between credit demand and supply can lead to lower interest rates rather than the increases predicted by Neoclassical theory (Murphy and Walsh, 2022). In tandem, some empirical studies have shown that government spending can be associated with lower interest rates at the firm level (Gabriel (2022)), the local level (Auerbach, Gorodnichenko, and Murphy (2020a)), and even the national level (Ramey (2011), Corsetti, Meier, and Müller (2012), and D'Alessandro, Fella, and Melosi (2019)).

Notwithstanding, the observed reduction in interest rates for winning firms should not be interpreted as evidence that government spending broadly lowers equilibrium interest rates. Similar to the findings of Hebous and Zimmermann (2020) and Gabriel (2022), the link between firm-level procurement shocks and aggregate interest rates is not straightforward. In this context, government procurement shocks do not necessarily indicate an overall increase in government spending or budget deficits but may instead reflect a reallocation of public funds.

V. Breaking Down the Regional Employment Multiplier

In the previous sections, we provided direct evidence that firms hire more employees when they win additional contracts, highlighting a transmission mechanism of fiscal policy operating through the extensive margin of employment. Moreover, using loan-level data from the Federal Reserve, we also demonstrated that when (small) firms win more contracts, they receive more credit and experience lower interest rates, suggesting the existence of a financial accelerator mechanism (*borrow-to-hire*).

Nonetheless, the aggregate effects of unanticipated contracts on employment remain unclear.

Specifically, if contractors "stole" employees from competitors, the net effect on employment would be zero. Conversely, if contractors hired from the unemployed or from new labor force entrants, we would observe positive employment multipliers.

In this section, we therefore study the effects of unanticipated contracts at the regional level (MSA) and examine how much of the regional employment multiplier originates from (i) the direct response of contractors to procurement spending shocks and (ii) the response of firms not directly affected by procurement.

V.a The Aggregate Effects of (Unanticipated) Contracts

Following Nakamura and Steinsson (2014), Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b), we estimate regional fiscal multipliers in the U.S. by employing the following equation:

$$\frac{E_{\ell,t+h} - E_{\ell,t-1}}{E_{\ell,t-1}} = \beta_h \cdot \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda_t + \alpha_\ell + \varepsilon_{\ell,t}, \tag{2}$$

where $E_{\ell,t}$ represents MSA-level employment in region ℓ and year t, and $Y_{\ell,t-1}$ is annual regional personal income. The use of personal income as a normalizing weight for changes in defense spending aligns with Muratori, Juarros, and Valderrama (2023) and is similar to the use of earnings in Auerbach, Gorodnichenko, and Murphy (2020b). The terms α_t and λ_ℓ are time and location fixed effects, respectively. The government spending measure $G_{\ell,t}$ represents defense contracts from FPDS, aggregated by MSA-year. Unlike Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b), who spread Department of Defense contracts over their duration, we use contracts as recorded in FPDS to avoid missing any anticipatory effect.[19] The estimand of interest, β_h, measures the percentage increase in regional employment in response

[19] Even at fiscal year frequency, defense contracts lead defense spending as measured from either NIPA or FPDS when contracts are spread over their duration. The inherent delay in NIPA spending, as discussed in Brunet (2023) and Briganti and Sellemi (2023), stems from NIPA accounting practices. To avoid underestimating the multiplier due to delays in the defense spending measure, we prefer aggregate changes in military contracts over defense spending for constructing the instrument. Notably, Demyanyk, Loutskina, and Murphy (2019) report higher spending estimates on several outcomes when using contracts instead of spending (see interaction terms in Table 2 and Table 7).

to a 1% increase in defense spending relative to personal income.

Mintz (1992) notes that regional defense procurement spending may be endogenous due to political factors. To address this, Nakamura and Steinsson (2014) suggest a Bartik (1991)-type instrument, exploiting exogenous time variation in national defense spending. In our context, the Bartik-type instrument is defined as

$$Z_{\ell,t+h} := \frac{1}{Y_{\ell,t-1}} \cdot \underbrace{(G_{t+h} - G_{t-1})}_{\text{Shift}} \cdot \underbrace{s_\ell}_{\text{Share}},$$

where $G_{t+h} - G_{t-1}$ represents the aggregate change in lumpy defense contracts in FPDS (*shift*), and s_ℓ denotes the long-run exposure of regions to defense spending (*share*), calculated as $1/T \cdot \sum_{\tau=1}^{T} G_{\ell,\tau}/Y_{\ell,\tau}$, with Y being personal income.

In this framework, either the shift or the share must be uncorrelated with the outcome variable (Borusyak, Hull, and Jaravel (2022)). Although the shares—regional exposure to defense spending—might correlate with unobservables influencing long-term outcomes, the validity of the econometric approach relies on the exogeneity of the shift. This reliance aligns with the fiscal policy literature, which has long used national changes in defense spending—equivalent to the shift in this framework—as an instrument for government spending (Barro (1981), Ramey and Shapiro (1998)). Actually, in a regional setup, the required identification assumption is even weaker, as national military build-ups are less likely to correlate with regional business-cycle conditions (Nakamura and Steinsson (2014)).

The left panel of Figure 8 shows the 2SLS estimates of β_h, representing employment multipliers estimated using instrumented defense contracts. The results closely align with the employment multipliers estimated in Auerbach, Gorodnichenko, and Murphy (2020b) (Table 3, Panel B, left column).

We construct the implied number of jobs-year following Chodorow-Reich (2019) and Muratori,

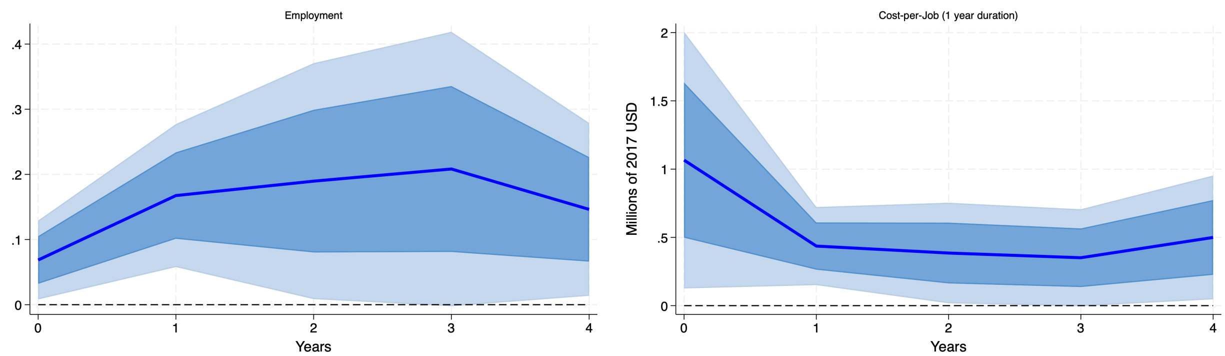

Employment Multipliers | Cost-per-Job (per year)

Figure 8: Effect of Contracts on MSA Employment

Notes: Sample spans 2006–2019 ($T = 14$) with $N = 329$ MSAs. Standard errors are clustered at the MSA level. Narrow bands indicate 68% confidence levels; wide bands indicate 90% confidence levels. Cost-per-job standard errors are computed using the Δ-method.

Juarros, and Valderrama (2023):

$$\text{jobs-year}_h := \beta_h \cdot \frac{1}{N \cdot T} \sum_{\ell=1}^{N} \sum_{t=2006+1}^{2019} \frac{\$1,000,000}{Y_{\ell,t-1}} \cdot E_{\ell,t-1},$$

where β_h represents the estimated employment multipliers reported in the left panel of Figure 8. In the right panel of the same figure, we plot the implied cost-per-job (with an average duration of one year) by taking the reciprocal of the number of jobs-year. The average cost-per-job is \$547,771.

To provide context, Chodorow-Reich (2019) review the effects of ARRA transfers and report a cost-per-job range (with an average duration of one year) between \$76,000 and \$393,000. Other estimates in the literature include Serrato and Wingender (2016), who find a cost-per-job of \$30,000 per year using Census data from U.S. regions; Adelino, Cunha, and Ferreira (2017), who report \$20,000 per year based on exogenous variation in local spending in low-employment U.S. municipalities; Corbi, Papaioannou, and Surico (2019), who estimate \$8,000 per year using Brazilian municipal transfers; and Buchheim and Watzinger (2023), who find \$24,000 per year based on German public investment in improving the energy efficiency of school buildings.

Given these benchmarks, it is worth investigating why the cost-per-job implied by defense procurement spending in U.S. regions is so high. To this end, we disaggregate regional defense spending, $G_{\ell,t}$, by aggregating at the MSA-year level the two components used in the firm-level analysis from Section 3: (i) unanticipated contracts ($\varepsilon_{\ell,t}$) and (ii) potentially anticipated contracts ($\tilde{G}_{\ell,t}$). Specifically, we augment our specification in Equation (2) by decomposing $G_{\ell,t}$:

$$\frac{E_{\ell,t+h} - E_{\ell,t-1}}{E_{\ell,t-1}} = \gamma_h \cdot \frac{\varepsilon_{\ell,t+h} - \varepsilon_{\ell,t-1}}{Y_{\ell,t-1}} + \delta_h \cdot \frac{\tilde{G}_{\ell,t+h} - \tilde{G}_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda_t + \alpha_\ell + \varepsilon_{\ell,t}. \tag{3}$$

When aggregating firm-level unanticipated contracts at the MSA-year level to study regional employment, we face a potential endogeneity threat: for political reasons, more (unanticipated) contracts might be awarded in regions with poor economic performance (i.e., low growth). This issue persists even if firms within the region cannot predict contract winners due to the competitive allocation system. Such concerns have driven the literature to adopt Bartik-type instruments to study regional government purchases, as used in the previous section (Nakamura and Steinsson (2014), Demyanyk, Loutskina, and Murphy (2019), and Auerbach, Gorodnichenko, and Murphy (2020b)). Accordingly, we construct two instruments:

$$\begin{aligned} Z^{\varepsilon}_{\ell,t+h} &:= \frac{\varepsilon_{t+h} - \varepsilon_{t-1}}{Y_{\ell,t-1}} \cdot s^{\varepsilon}_{\ell} \\ Z^{\tilde{G}}_{\ell,t+h} &:= \frac{\tilde{G}_{t+h} - \tilde{G}_{t-1}}{Y_{\ell,t-1}} \cdot s^{\tilde{G}}_{\ell}, \end{aligned}$$

where ε_t represents the yearly aggregated value of unanticipated contracts, $\tilde{G}_t$ denotes the annual value of potentially anticipated contracts, s^{ε}_{ℓ} is the long-run share of regional unanticipated contracts (i.e., $1/T \sum_t \varepsilon_{\ell,t}/\varepsilon_t$), and $s^{\tilde{G}}_{\ell}$ is the analogous share for potentially anticipated contracts. A similar two-instrument approach was recently proposed for analyzing $G_{\ell,t}$ by disaggregating goods and services spending (Muratori, Juarros, and Valderrama, 2023).

We estimate Equation (3) using $Z^{\varepsilon}_{\ell,t+h}$ and $Z^{\tilde{G}}_{\ell,t+h}$ as instruments for the two types of spending. The 2SLS estimates of γ_h (unanticipated contracts) are shown in red in the top panel of Figure 9, while the estimates of δ_h (potentially anticipated contracts) are shown in green.

The employment multipliers estimated using unanticipated contracts (red line in the top panel of

Employment Multipliers by Contract Type

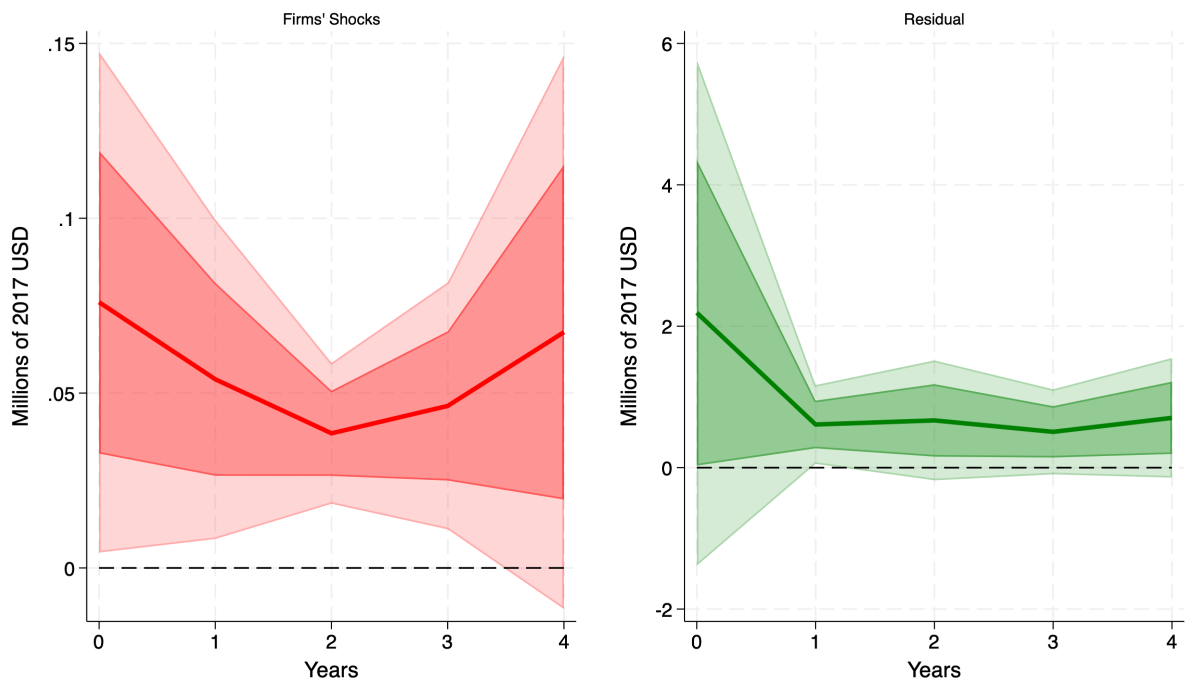

Cost-per-Job (per year) by Contract Type

Figure 9: Effect of Contracts on MSA Employment by Contract Type

Notes: Sample spans 2006–2019 ($T = 14$), with $N = 329$ MSAs. Standard errors are clustered at the MSA level. Narrow bands indicate 68% confidence levels; wide bands indicate 90% confidence levels. Cost-per-job standard errors are computed using the Δ-method.

Figure 9) are an order of magnitude larger than those estimated for the residual components. Since most defense procurement spending originates from the residual components (i.e., potentially anticipated contracts), the estimated multipliers closely resemble those estimated using $G_{\ell,t}$ (Figure 8), which, in turn, align with the estimates reported in Auerbach, Gorodnichenko, and Murphy (2020b). Similarly, the cost-per-job estimates for potentially anticipated contracts (green line, bottom-right panel of Figure 9) are nearly identical to those for all defense procurement contracts. Conversely, the cost-per-job estimates for unanticipated contracts are much lower, averaging \$56,444 per year. This value is consistent with literature estimates for other types of government spending, suggesting that the low cost-per-job estimates for defense procurement spending arise from its measurement.

As discussed in Section 2, most defense procurement spending in the U.S. consists of modifications or delivery/task orders awarded under prior indefinite delivery vehicles. These contracts, often awarded years earlier, may have minimal or diluted effects over extended periods. In contrast, unanticipated contracts capture the lumpy, unpredictable component of procurement spending, whose effects are more likely to materialize shortly after their award.

Therefore, contrary to previous estimates, defense procurement spending appears capable of stimulating employment to a degree comparable to other components of government spending when properly identified. Moreover, the positive and significant employment multipliers suggest that the net regional effects of unanticipated contracts are positive, with the new employees likely drawn from new labor force entrants or the unemployed (assuming minimal migration and cross-MSA commuting effects).

V.b Direct vs Indirect Regional Employment Channels

In the previous section, the positive regional employment multiplier indicates that when firms win more contracts and hire more employees, they likely do so from either the unemployed or new labor force entrants. However, part of the employment response might not only be due to firms directly affected by procurement contracts, i.e., contractors. In fact, firms not directly involved with contracting might well be affected by regional shocks to contracts (Juarros (2022)) and decide to increase employment due to several reasons: (i) input-output connections (Auerbach, Gorodnichenko, and Murphy (2020b)), (ii) subcontracting or (iii) more general "*multiplier-effects*", such as increased local consumption (see Demyanyk, Loutskina, and Murphy (2019)).

Generally speaking, *how much of the employment multiplier originates from direct effects of contracts on contractors and how much from indirect effects on non-contractors?*

To answer this question, we leverage again our restricted data access to the universe of establishments in the U.S. (QCEW). Specifically, the BLS regional employment public data is constructed by aggregating the micro (restricted) data from QCEW, to which we obtained access:

$$\underbrace{E_{\ell,t}}_{\text{Public}} = \sum_i \underbrace{E_{i,\ell,t}}_{\text{Restricted}},$$

where i denotes an establishment operating in period t in region ℓ, identified by its physical location address.

Thus, for each region, we break down employment into two components:

$$\begin{aligned} E_{\ell,t} &= \sum_{i \in \mathcal{C}} E^{\text{Contractors}}_{i,\ell,t} + \sum_{i \notin \mathcal{C}} E^{\text{Non-Contractors}}_{i,\ell,t} \\ &= E^{\text{Contractors}}_{\ell,t} + E^{\text{Non-Contractors}}_{\ell,t}. \end{aligned}$$

Essentially, we identify the set of establishments that win at least one contract over the sample period, $\mathcal{C}$, and aggregate employment for these establishments into the component associated with defense contractors. The residual employment component represents firms that are never directly involved with defense contracting. This breakdown of employment is implemented by matching the universe of defense contractors from FPDS with the universe of establishments from the restricted QCEW. The matching is restricted to establishments located in all of our 42 signatory states and is carried out using the same string-matching algorithm employed in Section 3. The total aggregated amount of defense spending resulting from aggregating contracts awarded to our set of matched defense contractors accounts for more than 90% of total defense spending for each state included in the analysis, on average.

Employment Breakdown. Having constructed MSA-annual time series of defense contractors' and non-contractors' employment, which add up exactly to the publicly available QCEW data, we can break down the left-hand side of Equation (2) into two components:

$$\frac{E_{\ell,t+h}-E_{\ell,t-1}}{E_{\ell,t-1}}=\frac{E^{\text{Contractors}}_{\ell,t+h}-E^{\text{Contractors}}_{\ell,t-1}}{E^{\text{Contractors}}_{\ell,t-1}}\cdot\underbrace{\frac{E^{\text{Contractors}}_{\ell,t-1}}{E_{\ell,t-1}}}_{:=s_{\ell,t-1}}+\ldots$$
$$+\frac{E^{\text{Non-Contractors}}_{\ell,t+h}-E^{\text{Non-Contractors}}_{\ell,t-1}}{E^{\text{Non-Contractors}}_{\ell,t-1}}\cdot\left(1-\underbrace{\frac{E^{\text{Contractors}}_{\ell,t-1}}{E_{\ell,t-1}}}_{:=s_{\ell,t-1}}\right).$$

Essentially, the percent change in total regional employment is a weighted average of the percent changes in contractors' and non-contractors' employment, with the weight $s_{\ell,t-1}$ representing the pre-shock fraction of contractors' employment in a region.

Therefore, we break down the 2SLS estimate of β_h from equation (2) in two components by estimating the following two equations:

$$\frac{E^{\text{Contractors}}_{\ell,t+h}-E^{\text{Contractors}}_{\ell,t-1}}{E^{\text{Contractors}}_{\ell,t-1}}\cdot s_{\ell,t-1}=\beta^{\text{c}}_h\cdot\frac{G_{\ell,t+h}-G_{\ell,t-1}}{Y_{\ell,t-1}}+\lambda^{\text{c}}_t+\alpha^{\text{c}}_\ell+\varepsilon^{\text{c}}_{\ell,t} \tag{4}$$
$$\frac{E^{\text{Non-Contractors}}_{\ell,t+h}-E^{\text{Non-Contractors}}_{\ell,t-1}}{E^{\text{Non-Contractors}}_{\ell,t-1}}\cdot(1-s_{\ell,t-1})=\beta^{\text{nc}}_h\cdot\frac{G_{\ell,t+h}-G_{\ell,t-1}}{Y_{\ell,t-1}}+\lambda^{\text{nc}}_t+\alpha^{\text{nc}}_\ell+\varepsilon^{\text{nc}}_{\ell,t}. \tag{5}$$

We estimate Equation (4) and (5) using the same Bartik instruments used before. In practice, we are breaking down the employment multiplier, estimated using the widely accepted MSA-level government spending shocks of Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b), into a contractor and non-contractor component. In fact, thanks to the linearity of the 2SLS estimator and the fact that the right hand side of Equations (2), (4) and (5) are the same, the 2SLS estimates of β^c_h and β^{nc}_h from (4) and (5) respectively add up to the 2SLS estimate of β_h:

$$\hat{\beta}_h^{\text{2SLS}} = \underbrace{\hat{\beta}_h^{c,\text{2SLS}}}_{\text{Direct}} + \underbrace{\hat{\beta}_h^{nc,\text{2SLS}}}_{\text{Indirect}}.$$

In particular, the 2SLS estimate of β_h^{c} represents the **direct** effect of defense spending on defense contractors, while the estimate of β_h^{nc} represents the **indirect** effect of defense spending.

The top-left panel of Figure 10 shows the estimates of β_h^c, the component of the employment multiplier that is due to the direct response of contractors (Equation (4)). The top-right panel of Figure 10 shows the estimates of β_h^{nc}, the component that stems from non-contractors (Equation (5)).

Both the responses of the (weighted) employment components to regional defense spending shocks are positive and significant with the only exclusion of non-contractors employment on impact, whose estimate is negative but insignificant, suggesting only potentially very mild employment crowding-out effects.[20]

The bottom panel of Figure 10 shows the contractors' share of the employment multiplier in blue (i.e., $\hat{\beta}_h^c / \left(\hat{\beta}_h^c + \hat{\beta}_h^{nc}\right)$. The red dash line is the average contractors' share of the employment multiplier calculated from horizon 1 through 4: 55% of the regional multiplier originates from the response of contractors, while the remaining 45% is driven by indirect effects. This novel result suggests that "*multiplier*" effects are almost as important as the direct effects in the propagation of fiscal policy shocks.

VI. Conclusion

In this paper, we use institutional knowledge of the federal procurement process in the United States to identify a special set of "*unanticipated contracts*" that act as firm-level demand shocks. First, by leveraging restricted access to the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics, we link the universe of U.S. establishments with the universe of fed-

[20]Notice that the sum of the two components does not add up exactly to the value estimated in Figure 8 but is very close to it. The reason for this discrepancy arises from the fact that the sample we used at the BLS only has 271 MSAs, since we only had access to 42 signatory states. However, the employment multipliers estimated using this smaller sample were almost identical to the ones reported here using all MSAs, serving as an extra robustness check.

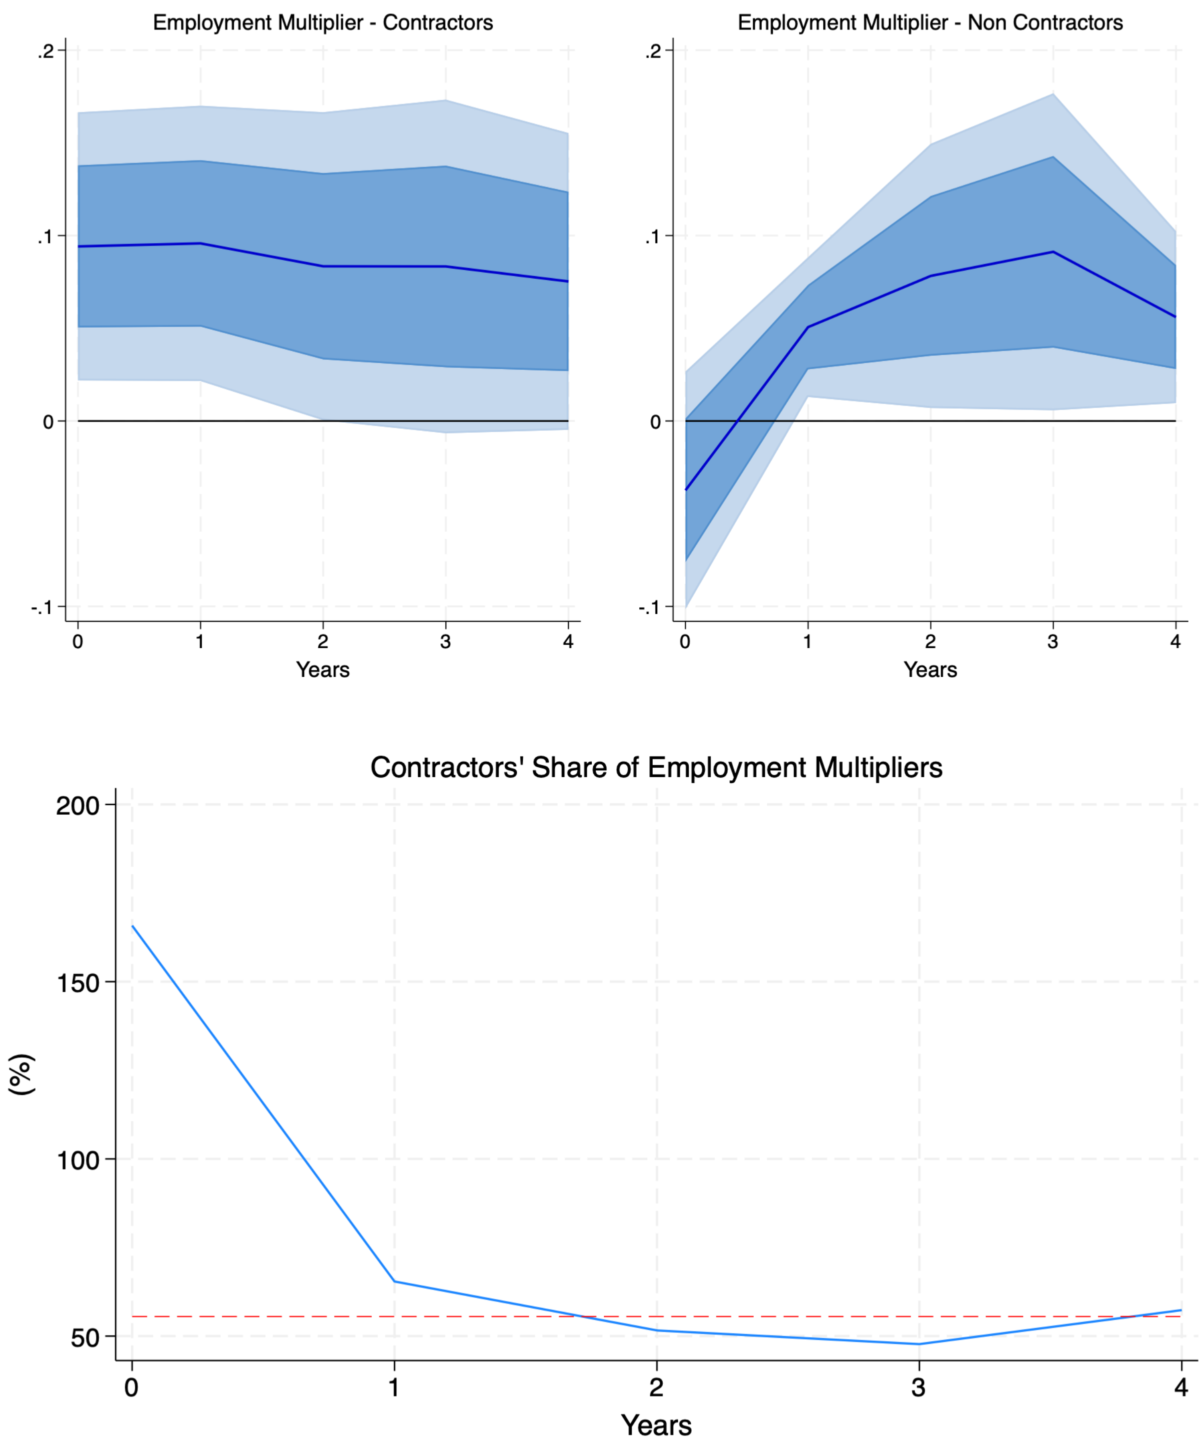

Figure 10: Employment Multipliers: Contractors vs Non-Contractors

Notes: Sample spans 2006–2019 ($T = 14$), with $N = 271$ MSAs. Standard errors are clustered at the MSA level. Narrow bands indicate 68% confidence levels; wide bands indicate 90% confidence levels.

eral contractors from the Federal Procurement Data System (FPDS). We find that when firms win (unanticipated) contracts, they increase their employment, with this increase persisting beyond the average duration of contracts. This finding provides direct evidence of an important transmission mechanism of fiscal shocks: the extensive margin of employment. Second, we link contractors from FPDS with firm loan bank-level data from the Federal Reserve (Y14Q) to show that unanticipated contracts also lead small firms to increase their credit and lower their borrowing costs. We refer to this credit channel of procurement contracts as the "*borrow-to-hire*" mechanism.

In the second part of the paper, we aggregate unanticipated contracts to the MSA level to study their regional effects. We find that unanticipated contracts lead to positive estimates of the employment multiplier, which are an order of magnitude higher than current estimates. The implied cost-per-job is $57,000 per year, aligning the effects of defense procurement spending on employment with other estimates in the literature based on either different types of government spending (e.g., ARRA transfers) or data from other countries (e.g., Brazil and Germany).

Lastly, we construct new regional time series for employment among defense contractors and non-defense contractors, again leveraging restricted QCEW access. Using these data, we decompose the employment multiplier, as estimated using the widely accepted MSA-level shock of Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020b), into (i) direct and (ii) indirect ("*multiplier*") effects. We find that non-contractors experience a contraction in employment only on impact (i.e., crowding-out), but these estimates are not statistically significant. From year 1 through year 4 after the fiscal shock, 55% of the employment multiplier originates from contractors' responses, while the remaining 45% stems from non-contractors.

These results have important implications for policymakers. First, governments can stimulate employment persistently through government acquisitions, with minimal crowding-out effects. Second, access to credit can help small firms grow out of financial constraints. Third, optimal fiscal policy should equally consider both direct and indirect effects when designing procurement policies. For instance, factors such as firm size (small vs large), contract type (cost-plus-fee vs fixed price), contract duration (short vs long) and product type (services vs goods) are all critical elements that influence the size of the direct component of the employment multiplier. Simultaneously, regional characteristics, such as the share of financially constrained households, sectoral biases, the production network location of purchased products, and regional economic slack, are key determinants of the indirect component of the employment multiplier.

Bibliography

Adelino, **Manuel**, **Cunha**, **Igor**, and **Ferreira**, **Miguel A.** (Sept. 2017). “The Economic Effects of Public Financing: Evidence from Municipal Bond Ratings Recalibration”. *The Review of Financial Studies* 30.9, pp. 3223–3268. ISSN 0893-9454, 1465-7368.

Auerbach, **Alan**, **Gorodnichenko**, **Yuriy**, and **Murphy**, **Daniel** (May 2020a). “Effects of Fiscal Policy on Credit Markets”. *AEA Papers and Proceedings* 110, pp. 119–124. ISSN 2574-0768, 2574-0776.

— (Mar. 2020b). “Local Fiscal Multipliers and Fiscal Spillovers in the USA”. *IMF Economic Review* 68.1, pp. 195–229. ISSN 2041-4161, 2041-417X.

Barattieri, **Alessandro**, **Cacciatore**, **Matteo**, and **Traum**, **Nora** (Sept. 2023). “Estimating the Effects of Government Spending Through the Production Network”. *NBER Working Paper* 31680.

Barro, **Robert** (Dec. 1981). “Output Effects of Government Purchases”. *Journal of Political Economy* 89.6, pp. 1086–1121.

Bartik, **Timothy J.** (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, Mich: W.E. Upjohn Institute for Employment Research. ISBN 978-0-88099-114-8 978-0-88099-113-1.

Beyhaghi, **Mehdi** (Mar. 2022). “Third-Party Credit Guarantees and the Cost of Debt: Evidence from Corporate Loans”. *Review of Finance* 26.2, pp. 287–317. ISSN 1572-3097, 1573-692X.

Beyhaghi, **Mehdi** et al. (2024). “Monetary Policy and Corporate Investment: The Equity Financing Channel”. *SSRN Electronic Journal*. ISSN 1556-5068.

Bonfim, **Diana**, **Custódio**, **Cláudia**, and **Raposo**, **Clara** (Mar. 2023). “Supporting Small Firms Through Recessions and Recoveries”. *Journal of Financial Economics* 147.3, pp. 658–688. ISSN 0304405X.

Born, **Benjamin** et al. (Sept. 2023). “Mr. Keynes Meets the Classics: Government Spending and the Real Exchange Rate”. *Journal of Political Economy*, p. 727707. ISSN 0022-3808, 1537-534X.

Borusyak, **Kirill**, **Hull**, **Peter**, and **Jaravel**, **Xavier** (Jan. 2022). “Quasi-Experimental Shift-Share Research Designs”. *The Review of Economic Studies* 89.1. Ed. by **Dirk Krueger**, pp. 181–213. ISSN 0034-6527, 1467-937X.

Briganti, **Edoardo**, **Brunet**, **Gillian**, and **Sellemi**, **Victor** (2025). “When Does Government Spending Matter? It’s All in the Measurement”. *Working Paper*.

Briganti, **Edoardo** and **Sellemi**, **Victor** (Mar. 2023). “Why Does GDP Move Before Government Spending? It’s all in the Measurement”. *UCSD Manuscript*.

Brown, **James R.**, **Gustafson**, **Matthew T.**, and **Ivanov**, **Ivan T.** (Aug. 2021). “Weathering Cash Flow Shocks”. *The Journal of Finance* 76.4, pp. 1731–1772. ISSN 0022-1082, 1540-6261.

Brunet, **Gillian** (Jan. 2023). “When Does Government Spending Matter? Evidence from a New Measure of Spending”. *Working Paper*.

Buchheim, **Lukas** and **Watzinger**, **Martin** (Feb. 2023). “The Employment Effects of Countercyclical Public Investments”. *American Economic Journal: Economic Policy* 15.1, pp. 154–173. ISSN 1945-7731, 1945-774X.

Chodorow-Reich, **Gabriel** (May 2019). “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?” *American Economic Journal: Economic Policy* 11.2, pp. 1–34. ISSN 1945-7731, 1945-774X.

Choi, **Joonkyu**, **Penciakova**, **Veronika**, and **Saffie**, **Felipe** (July 2023). “Political Connections, Allocation of Stimulus Spending, and the Jobs Multiplier”. *Working Paper*.

Corbi, **Raphael**, **Papaioannou**, **Elias**, and **Surico**, **Paolo** (Oct. 2019). “Regional Transfer Multipliers”. *The Review of Economic Studies* 86.5, pp. 1901–1934. ISSN 0034-6527, 1467-937X.

Corsetti, **Giancarlo**, **Meier**, **André**, and **Müller**, **Gemot J** (Nov. 2012). “Fiscal Stimulus with Spending Reversal”. *The Review of Economics and Statistics* 94.4, pp. 878–895.

Cox, **Lydia** et al. (Oct. 2024). “Big *G*”. *Journal of Political Economy* 132.10, pp. 3260–3297. ISSN 0022-3808, 1537-534X.

D'Alessandro, **Antonello**, **Fella**, **Giulio**, and **Melosi**, **Leonardo** (Aug. 2019). "Fiscal Stimulus with Learning-by-Doing". *International Economic Review* 60.3, pp. 1413–1432. ISSN 0020-6598, 1468-2354.

Demyanyk, **Yuliya**, **Loutskina**, **Elena**, and **Murphy**, **Daniel** (Oct. 2019). "Fiscal Stimulus and Consumer Debt". *The Review of Economics and Statistics* 101.4, pp. 728–741. ISSN 0034-6535, 1530-9142.

Dupor, **Bill** and **Guerrero**, **Rodrigo** (Dec. 2017). "Local and Aggregate Fiscal Policy Multipliers". *Journal of Monetary Economics* 92, pp. 16–30. ISSN 03043932.

Ferraz, **Claudio**, **Finan**, **Frederico**, and **Szerman**, **Dimitri** (2021). "Procuring Firm Growth: The Effects of Government Purchases on Firm Dynamics". *Working Paper*.

Gabriel, **Ricardo Duque** (2022). "The Credit Channel of Public Procurement". *SSRN Electronic Journal*. ISSN 1556-5068.

Giovanni, **Julian di** et al. (June 2023). "Buy Big or Buy Small? Procurement Policies, Firms' Financing, and the Macroeconomy". *Working Paper*.

Gonzalez-Lira, **Andres**, **Carril**, **Rodrigo**, and **Walker**, **Michael S** (Jan. 2021). "Competition under Incomplete Contracts and the Design of Procurement Policies". *Working Paper*, p. 104.

Hebous, **Shafik** and **Zimmermann**, **Tom** (Sept. 2020). "Can Government Demand Stimulate Private Investment? Evidence from U.S. Federal Procurement". *Journal of Monetary Economics*, S0304393220301100. ISSN 03043932.

Ilzetzki, **Ethan** (Jan. 2023). "Learning by Necessity: Government Demand, Capacity Constraints, and Productivity Growth". *Working Paper*.

Jordà, **Òscar** (Feb. 2005). "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review* 95.1, pp. 161–182. ISSN 0002-8282.

Jørgensen, **Peter L.** and **Ravn**, **Søren H.** (Jan. 2022). "The Inflation Response to Government Spending Shocks: A Fiscal Price Puzzle?" *European Economic Review* 141, p. 103982. ISSN 00142921.

Juarros, **Pedro** (Nov. 2022). "Fiscal Stimulus, Credit Frictions and the Amplification Effects of Small Firms". *Working Paper*.

Lee, **Munseob** (2024). “Government Purchases and Firm Growth”. *American Economic Journal: Applied Economics*.

Mintz, **Alex** (1992). *The Political Economy of Military Spending in the United States*. Routledge. ISBN 978-0-415-07595-4.

Muratori, **Umberto**, **Juarros**, **Pedro**, and **Valderrama**, **Daniel** (Mar. 2023). “Heterogeneous Spending, Heterogeneous Multipliers”. *IMF Working Papers* 2023.052, p. 1. ISSN 1018-5941.

Murphy, **Daniel** and **Walsh**, **Kieran James** (May 2022). “Government Spending and Interest Rates”. *Journal of International Money and Finance* 123, p. 102598. ISSN 02615606.

Nakamura, **Emi** and **Steinsson**, **Jón** (Mar. 2014). “Fiscal Stimulus in a Monetary Union: Evidence from US Regions”. *American Economic Review* 104.3, pp. 753–792. ISSN 0002-8282.

Nekarda, **Christopher** and **Ramey**, **Valerie** (Jan. 2011). “Industry Evidence on the Effects of Government Spending”. *American Economic Journal: Macroeconomics* 3.1, pp. 36–59. ISSN 1945-7707, 1945-7715.

Perotti, **Roberto** (Jan. 2007). “In Search of the Transmission Mechanism of Fiscal Policy [with Comments and Discussion]”. *NBER Macroeconomics Annual* 22, pp. 169–249. ISSN 0889-3365, 1537-2642.

Pinardon-Touati, **Noemie** (Feb. 2022). “The Crowding Out Effect of Local Government Debt: Micro- and Macro-Estimates”.

Ramey, **Valerie** (Feb. 2011). “Identifying Government Spending Shocks: It’s All in the Timing”. *The Quarterly Journal of Economics* 126.1, pp. 1–50. ISSN 0033-5533, 1531-4650.

Ramey, **Valerie** and **Shapiro**, **Matthew** (1998). “Costly Capital Reallocation and the Effects of Government Spending”. *Carnegie-Rochester Conference on Public Policy* 48.1998, pp. 145–194.

Serrato, **Juan Carlos Suárez** and **Wingender**, **Philippe** (July 2016). “Estimating Local Fiscal Multipliers”. *NBER Working Paper*, w22425.