Breaking Down the U.S. Employment Multiplier Using Micro-Level Data

Edoardo Briganti*
Bank of Canada

Holt Dwyer†
UC San Diego

Ricardo Duque Gabriel‡
Federal Reserve Board

Victor Sellemi§
UC San Diego

This Version: April 2026

First Version: November 2024

Abstract

Using newly matched U.S. defense contract and restricted administrative employment data, we show that the regional employment effects of defense procurement are costly, concentrated, and slow to diffuse. Employment gains are initially driven by large existing contractors and come at a high cost of approximately $290,000 per job-year, well above benchmark estimates in the fiscal policy literature. Non-contractors are crowded out on impact, but positive spillovers emerge gradually and account for more than half of regional employment gains by the third year. A newly identified set of unanticipated contracts reveals persistent employment gains at recipient establishments, but these account for only about 18% of contractor job creation within one year, with the remainder arising through indirect channels within contractor networks.

*Email: ebriganti@bankofcanada.ca Webpage: edoardobriganti.com

†Email: bdwyer@ucsd.edu Webpage: holtdwyer.com

‡Email: ricardo.f.duquegabriel@frb.gov Webpage: ricardogabriel.com

§Email: vsellemi@ucsd.edu Webpage: victorsellemi.com

This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views expressed here are those of the authors and do not necessarily reflect the views of the BLS, the U.S. government, the Bank of Canada, the Federal Reserve Board, or the Federal Reserve System. We are grateful to an anonymous U.S. government contracting officer for extensive assistance with institutional knowledge of federal contracting, and to Jessica Helfand of the BLS for assistance in interfacing with the LDBE data. We also thank Miguel Bandeira, Yvan Becard, Gabriel Chodorow-Reich (discussant), Joonkyu Choi, Leland Crane, Bill Dupor, Manuel Garcia Santana, Yuriy Gorodnichenko, Jim Hamilton, Munseob Lee, Karel Mertens, Umberto Muratori, Daniel Murphy, James Poterba, Valerie Ramey, Felipe Saffie, Paolo Surico, Kieran James Walsh, Johannes Wieland, Sarah Zubairy, and other seminar participants at the 2025 NBER Conference: Fiscal Dynamics of State and Local Governments, UCSD, Bank of Canada, Banco Central do Brasil, Federal Reserve Board, IAAE 2025, EEA 2025, PEJ 2025, and SGE 2025 for helpful comments. Keywords: Fiscal Policy, Employment, Public Procurement, Military Spending. JEL: E62, H57, J21.

I. Introduction

Government procurement is one of the largest components of government spending. Defense procurement, which accounts for approximately two-thirds of total federal procurement in the United States, has been of particular interest in macroeconomics, where variation in defense spending has served as a key source of identification for estimating the effects of fiscal policy (Ramey (2011), Nakamura and Steinsson (2014), Auerbach, Gorodnichenko, and Murphy (2020), and Cox et al. (2024)). The hundreds of billions of dollars of defense spending annually channeled to the private sector flow disproportionately to specific firms and regions, potentially shaping supply chains and regional labor markets in the process. Yet, despite its scale and policy importance, the channels through which defense procurement generates local employment are still not fully understood.

Understanding these mechanisms matters both for evaluating the macroeconomic consequences of defense buildups and for anticipating their incidence across firms and local labor markets. Rising geopolitical tensions and rearmament efforts have renewed policy attention to defense procurement and its broader economic effects. As President Biden put it during a 2022 visit to a missile factory:

> "[b]eing the arsenal of democracy also means good-paying jobs for American workers in Alabama and the States all across America where defense equipment is manufactured and assembled." (Joseph R. Biden, 2022)

This paper breaks down the employment gains generated by procurement spending to shed light on the channels through which job creation takes place. We document four facts. First, defense procurement raises employment, but generates relatively few jobs per dollar spent: the implied cost per job-year lies above the benchmark range commonly cited in the recent fiscal policy literature. Second, the employment response is highly concentrated among large firms, which account for the bulk of both contract awards and aggregate job creation. Third, employment gains are initially concentrated among contractors and diffuse only gradually to non-contractors, with broader local spillovers emerging over time. Fourth, only a small share of contractor employment gains reflects direct hiring at recipient establishments; most job creation within the contractor network appears to operate through indirect channels, including subcontracting relationships, supply-chain linkages, and reallocation across establishments within firms.

Taken together, our results show that defense procurement raises employment, but in a way that is costly, concentrated among large firms, slow to diffuse beyond the contractor network, and largely indirect. By breaking down the employment response across firm size, contractor status, and direct contract recipient status, the paper traces where procurement-induced jobs arise and how they propagate through local labor markets.

We establish these results by combining the universe of federal contracts from the Federal Procurement Data System (FPDS) with several sources of employment data. At the regional level, we construct annual defense procurement for U.S. metropolitan statistical areas (MSAs) and merge these data with public employment counts from the Business Dynamics Statistics (BDS) and Local Area Unemployment Statistics (LAUS) datasets. We then complement this regional analysis with restricted administrative microdata from Longitudinal Database of Business Establishments (LDBE) at the BLS. Merging the FPDS and LDBE allows us to distinguish contractors from non-contractors and to trace the direct employment response of recipient establishments to plausibly unanticipated contracts. Our main findings can be summarized as follows.

First, defense procurement in the 21st century raises regional employment at a high cost per job. Our baseline regional estimates follow the standard cross-sectional defense-spending design of Auerbach, Gorodnichenko, and Murphy (2020), using Jordà (2005) local projections to study the effects of spending at different horizons, and an instrument for regional procurement that interacts national changes in defense spending between 2001 and 2019 with long-run regional exposure to defense contracts. We find that defense procurement generates modest but persistent employment gains. A shock equal to 1 percent of regional wages and salaries raises regional employment by roughly 0.1 percent over a three-year horizon.

The implied cost per job-year is substantially higher than the $25,000–$125,000 benchmark range emphasized in the recent fiscal multiplier literature (Chodorow-Reich, 2019). This comparison should be interpreted with caution, since the studies reviewed by Chodorow-Reich (2019) largely focus on transfers, which differ from defense procurement in instrument, sample period, and sectoral composition. That said, recent post-2000 studies of defense procurement also find similarly high cost-per-job estimates. While we discuss several potential explanations for this high cost per job, the most plausible is the growing concentration of defense procurement in sectors employing highly paid workers (Bartal and Becard, 2024).

Second, regional defense spending shocks primarily affect large firms, which receive more than 70% of total contracts and account for over 80% of the estimated employment gains. We break down the regional multiplier estimate by firm size and show that the employment response is overwhelmingly concentrated in large firms. Although firm-level studies often find that small businesses respond more strongly to a given procurement shock, our results show that, at the regional level, these firms play only a limited role.

Third, the employment effects of defense procurement are entirely among contractors within the first year and diffuse only gradually to non-contractors. Using matched LDBE data, we next break down the regional multiplier estimates into contractor and non-contractor responses. Employment gains are initially concentrated among contractors, while non-contractors experience mild crowding out on impact. Broader spillovers emerge only gradually and account for roughly half of total employment gains in the medium run. This delayed diffusion suggests that defense procurement is ill-suited to generate rapid, broad-based employment stimulus, but can support employment growth across the local economy over time.

Finally, we find that procurement contracts generate persistent but small employment gains at recipient establishments, accounting for only about 18% of the total contractor employment response to defense shocks. We isolate the direct effect of contract receipt at recipient establishments by exploiting institutional features of the federal procurement process to identify a set of plausibly unanticipated contracts. We find statistically significant and persistent employment gains at winning establishments, but these direct effects are quantitatively modest relative to the overall contractor response. Recipient establishments account for approximately 18% of contractor-side job creation. Most employment gains within the contractor network therefore appear to arise through indirect channels: internal reallocation across establishments, subcontracting, and supply-chain linkages.

In sum, our findings show that defense procurement raises employment, but in a way that is costly, concentrated, and slow to diffuse. The paper therefore contributes to the study of fiscal transmission not only by quantifying the employment effects of modern defense procurement, but also by providing a nested breakdown of the margins along which those effects arise from the regional economy, to the contractor network, to the individual establishment. To our knowledge, no existing study of government procurement traces the employment response across all three levels of aggregation. In doing so, the paper helps discipline claims about the job-creation potential of defense spending and provides evidence on the extent to which defense procurement can support broader economic objectives.

Related Literature and Contribution. This paper contributes to the literature on the subnational employment effects of government purchases in the United States, which has been studied at the industry level (Perotti, 2007; Nekarda and Ramey, 2011; Acemoglu, Akcigit, and Kerr, 2016; Komarek, Butts, and Wagner, 2022; Barattieri, Cacciatore, and Traum, 2025), the state level (Chodorow-Reich et al., 2012; Nakamura and Steinsson, 2014; Dupor and Guerrero, 2017), and at the regional or substate level (Demyanyk, Loutskina, and Murphy, 2019; Auerbach, Gorodnichenko, and Murphy, 2020; Muratori, Juarros, and Valderrama, 2023; Auerbach, Gorodnichenko, and Murphy, 2024). Overall, this literature finds that government spending raises employment and hours worked and lowers unemployment.[1] The main debate has therefore centered less on the sign of employment effects than on their magnitude and determinants. We complement this literature by moving beyond the regional multiplier itself to trace the distribution of employment gains across large and small firms, between contractors and non-contractors over time, and between direct and indirect effects within contractor networks.

Our paper also relates to the literature on the effects of procurement contracts at the firm level. Evidence from Austria (Gugler, Weichselbaumer, and Zulehner, 2020), Brazil (Ferraz, Finan, and Szerman, 2021), Portugal (Gabriel, 2024), South Korea (Lee, 2024), and Spain (di Giovanni et al., 2026) generally finds positive effects of procurement contracts on employment growth. For the United States, however, firm-level evidence on employment effects remains limited: existing work focuses primarily on the investment response of publicly traded firms (Hebous and Zimmermann, 2020; Budrys, 2022), on the financial performance of non-contractors (Juarros, 2022), or on the employment effects of changes in payment timing among small business contractors (Barrot and Nanda, 2020). To our knowledge, no prior U.S. study has focused on firm-level employment growth using administrative data at the quarterly frequency. Contemporaneous work by Park, Zhou, and Zubairy (2025) studies the effects of winning a contract and/or a subcontract on employment growth using annual data from the National Establishment Time Series (NETS), and finds positive and persistent employment effects.

[1] To the best of our knowledge, only Hager and Huber (2025) find negative long-run employment multipliers, using German data.

We contribute to this strand of the literature by exploiting institutional features of the U.S. federal procurement system to identify a novel set of plausibly unanticipated contracts, accounting for roughly 5% of Department of Defense spending, that can be used as establishment-level demand shocks. Using quarterly administrative data from the LDBE, we find positive and persistent employment effects at recipient establishments, but these are quantitatively modest. By comparing the establishment-level response to the regional contractor response, we provide the first estimate of the share of contractor employment gains that occur directly at recipient establishments.

The remainder of the paper is organized as follows. Section II discusses the institutional background of federal procurement in the United States. Section III describes the regional-level data and identification strategy and presents our baseline estimates of the employment multiplier. Section IV examines heterogeneity by firm size. Section V uses restricted-access LDBE data to break down the multiplier into contractor and non-contractor responses. Section VI presents the establishment-level analysis. Section VII concludes.

II. Institutional Background

Procurement spending refers to the purchase of goods and services by the government from private entities. Figure 1 plots federal procurement spending by fiscal year, as measured in the National Income and Product Accounts (NIPA). The left panel reports procurement as a share of total government spending, while the right panel shows its share of GDP. On average, procurement spending constitutes about 16% of total government spending and roughly 3% of GDP. Given its magnitude and its direct effect on U.S. private firms, federal procurement is an important channel through which the government can stimulate economic activity.

Since fiscal year 2001, the full universe of federal procurement contracts has been publicly available through `USASpending.gov`. These data are drawn from the FPDS, the platform used by federal contracting officers to record every federal contracting action. These detailed administrative microdata have been used previously by, among others, Demyanyk, Loutskina, and Murphy (2019), Auerbach, Gorodnichenko, and Murphy (2020), and Cox et al. (2024).

Figure 1 shows FPDS data, aggregated by fiscal year, in red. The FPDS series aligns closely with national accounts, offering an exceptionally detailed micro-origin of federal procurement spending.[2] The richness of FPDS data enables research at highly disaggregated levels, including sub-industries (six-digit NAICS codes), regions (MSAs and counties), firms, and even establishments. In this paper, we exploit this granularity to study the effects of government purchases on employment using a top-down approach, moving from the MSA level down to the level of recipient establishments.

Breakdown of Federal Contracting. As noted in Auerbach, Gorodnichenko, and Murphy (2020), behind each government contract lies a long history of transactions, with considerable heterogeneity in

[2]Differences in timing between the NIPA and FPDS series reflect how spending is recorded (Briganti, Brunet, and Sellemi, 2025). FPDS records obligations at the contract award date—when firms are most likely to begin responding to unexpected contracts—whereas NIPA appears to incorporate some military contracts with delay.

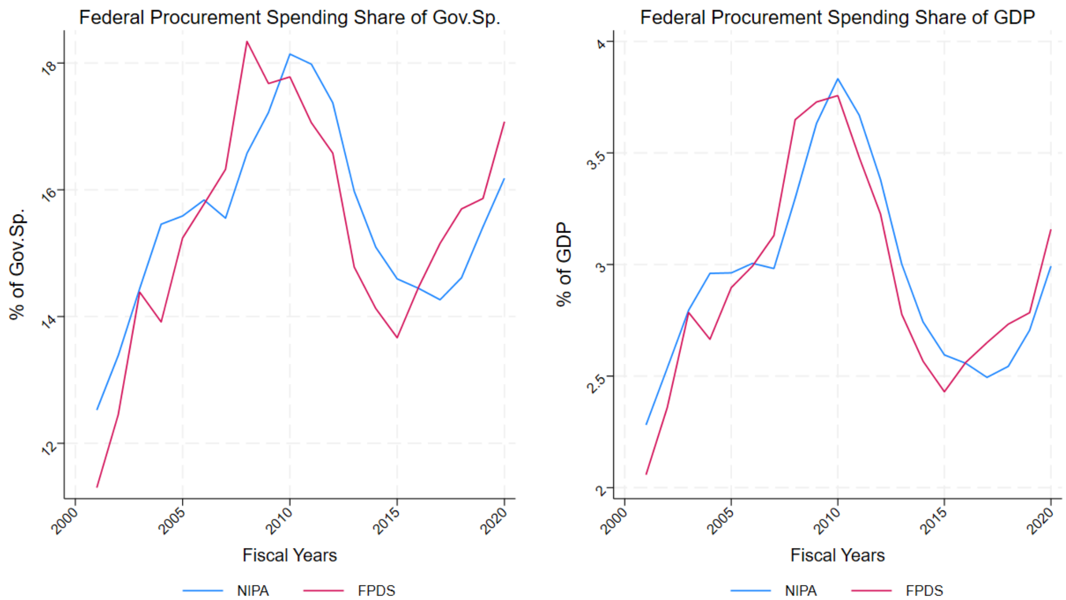

FIGURE 1 — FEDERAL PROCUREMENT SPENDING SHARES. LEFT PANEL: SHARE OF TOTAL GOVERNMENT SPENDING. RIGHT PANEL: SHARE OF GDP

Notes: Federal procurement spending is calculated by adding up (i) NIPA federal government intermediate goods and services purchased and (ii) NIPA federal government gross investment in structures, equipment, and software (see Cox et al. (2024) and Briganti and Sellemi (2023)). Federal spending is the sum of defense and non-defense spending. In both panels, the red series reports FPDS obligations aggregated by fiscal year from USASpending.gov and scaled by the same denominator as the corresponding NIPA series. Sample: FY2001–FY2019.

the types of contracts awarded. We present here a novel breakdown of federal contracting to highlight its complex and highly heterogeneous composition. Figure 2 reports the distribution of contracts across the most common categories.[3]

Two-thirds of all contracts (by value) are awarded by the Department of Defense. Regional-level analyses have traditionally focused on defense spending because its time variation is more plausibly exogenous. Consistent with this approach, we also concentrate on this component of spending in what follows.

Second, only 46.6% of all defense transactions in FPDS are newly awarded contracts. The remainder are contract modifications, such as options, extra work, or administrative actions, all related to an existing contract. Third, not all newly awarded contracts are necessarily "new." In fact, 75.2% of all "new" contracts are task orders (for services) and delivery orders (for goods) issued under a pre-existing parent contract, called an indefinite delivery vehicle (IDV).[4] Conversely, only 24.8% of new defense contracts

[3] We are grateful to a federal government contracting officer, who preferred to remain anonymous, for clarifying the details of each contract type.

[4] IDVs are regulated by the Federal Acquisition Regulation 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 benefits.

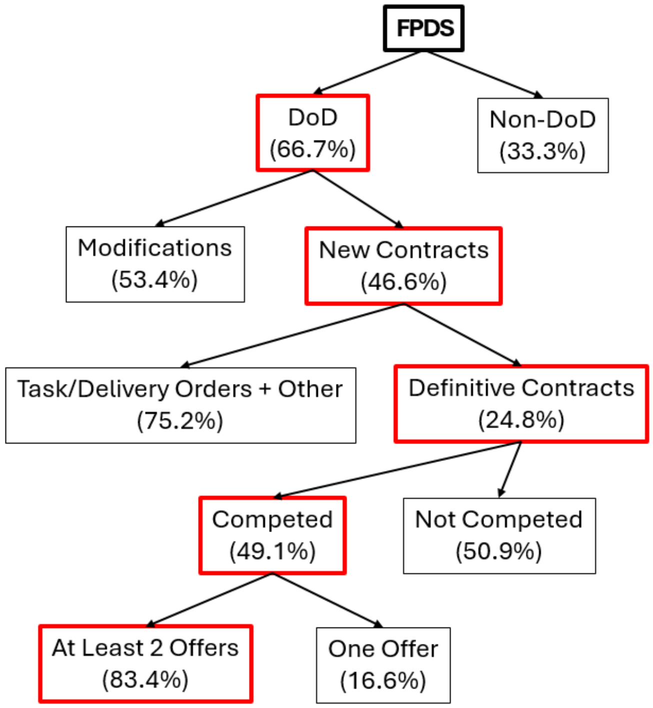

FIGURE 2 — BREAKDOWN OF FPDS CONTRACTS

Notes: Data refer to averages of fiscal-year shares, calculated using contract values rather than the number of contracts. Sample: FY2001–FY2019.

represent standalone contracts, which are not part of any ongoing relationship with the government. These are technically referred to as "definitive contracts." Fourth, 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 (see Federal Acquisition Regulation 6.302-1-a). Figure 2 shows that 50.9% of newly awarded definitive defense contracts fall into this non-competed category. FPDS reports the competition details and bid counts from a parent IDV to its subsequent task or delivery orders, even if these orders were not competed. As a result, many contracts seemingly reported as "new" and "competed" in FPDS are neither newly awarded (since their parent IDV may have been awarded months or even years earlier) nor individually competed, as they represent additional, potentially anticipated government orders placed under the terms of the originally competed IDV. Lastly, 83.4% of new, definitive, and competed defense contracts receive at least two offers, indicating the existence of genuine competition.

Overall, only about 5% of Department of Defense procurement spending is accounted for by new, definitive, and competed contracts with at least two offers.[5] Comparable shares hold for federal non-

[5]The U.S. defense industry has consolidated substantially, from 51 to 5 major prime contractors since the 1990s, and regulatory compliance requirements are widely recognized as barriers to entry for smaller firms (U.S. Department of Defense, 2022). Carril, Gonzalez-Lira, and Walker (2026) show that this consolidation reduced competition in DoD procurement, increasing

defense contracts, which make up the remaining one-third of total FPDS contracts.

Competitively awarded contracts undergo a public solicitation process designed to foster competition.[6] In Appendix A, we use publicly available data on the universe of federal contract notices from fiscal years 2006 to 2019 to reconstruct the full pre-award to award timeline of competed contracts from the earliest pre-solicitation date to the award notice date. We find that the median time from pre-solicitation to award notice is 20 days, while 75% of solicited contracts are awarded within 52 days. Thus, analyses conducted at the quarterly or annual level are unlikely to be affected by firms acting in response to favorable solicitations rather than actual awards, as these typically take place within the same time period.

III. The Regional Employment Multiplier

We begin by estimating the effect of defense procurement on regional employment. This aggregate estimate provides a useful benchmark, capturing both the response of contractors and spillovers to the broader local economy. We then use this baseline to study the mechanisms underlying the overall employment effect.

Data. We collect MSA-level data from different sources, summarized in the top panel of Table 1. The data are then merged into two harmonized datasets, illustrated in the bottom panel of the Table.

TABLE 1 — OUTCOME VARIABLES DATA SOURCES

Clean (Maximum Sample Size) Datasets:					
Source	*Institution*	*Availability*	*MSAs*	*Sample*	*Tables*
Quarterly Census of Employment and Wages (QCEW)	BEA	Public (Discontinued)	380	2001-2019	Tables B1
Longitudinal Database of Establishments (LDBE)	BLS	Restricted (Access Discontinued)	262	2006-2019	Tables E1, E2
Business Dynamics Statistics (BDS)	Census	Public	373	2001-2019	Tables D2, D3
Local Area Unemployment Statistics (LAUS)	BLS	Public	366	2001-2019	Table B9
Harmonized Merged Datasets:					
QCEW+BDS+LAUS (Baseline Large Sample)	-	-	358	2001-2019	Tables 2, 4, 5, B8, B10, D1
QCEW+BDS+LAUS+LDBE (Baseline Small Sample)	-	-	254	2006-2019	Tables 6, 7, B2, D4, D5

Notes: The QCEW public data table from the BEA is called `CAINC4_ALL_AREAS_1969_2022.csv`. The table has now been discontinued by the BEA due to budget cuts but it is still available for download from the BEA archive. The number of MSAs is obtained after merging the datasets with a common zip code, to county, to CBSA crosswalk available from `www.huduser.gov`, which is used to merge with FPDS contracts data.

the share of spending awarded without competition or via single-bid solicitations.

[6]See Federal Acquisition Regulation 5, *Publicizing Contract Actions*. Since October 1, 2001, contract actions with an expected value above $25,000 must be publicized on the government platform `sam.gov` (SAM). Contract actions below this threshold may still be posted to increase visibility.

First, we use the public version of the QCEW provided by the BEA to measure total employment, wages, personal income, and population at annual frequency from 2001 to 2019 for 380 MSAs. Second, we obtained access to firm-level microdata covering 42 states and District of Columbia from the LDBE, the BLS microdata used to produce the public QCEW.[7] After matching the universe of establishments with the universe of defense contractors from FPDS, we aggregate private employment data at an annual frequency from 2006 to 2019 for 262 MSAs. This allows us to break down the regional time series of employment into contractors (i.e., matched firms) and non-contractors (i.e., not matched firms). Third, we use the BDS from the Census to break down private employment and the total number of firms by firm size. These data are available at an annual frequency from 2001 to 2019 for 373 MSAs. Lastly, we use labor force and unemployment data from LAUS, provided by the BLS. We collect data at an annual frequency from 2001 to 2019 for 366 MSAs.

We then harmonize samples by merging the LAUS, BDS, and public QCEW data. This merged dataset includes 358 MSAs observed from 2001 to 2019. Our baseline results reported in the paper rely on this dataset.

The restricted QCEW data from LDBE impose a more severe sample reduction: the harmonized merged dataset from all four sources goes from 2006 to 2019 and includes 254 MSAs. Results that break down the employment multiplier into contractor and non-contractor responses use this database.

Estimation of the Regional Employment Multiplier. Following the literature on regional multipliers, we estimate the following equation:

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

where $E_{\ell,t}$ represents employment in region ℓ and year t, and $Y_{\ell,t-1}$ is annual regional wages and salaries.[8] The terms $\alpha_{\ell,h}$ and $\lambda_{t,h}$ are location and time fixed effects, respectively (specific to the horizon of the estimate h). Our focus on employment is motivated by the fact that variation in employment (rather than hours per employee) typically accounts for more than 80% of fluctuations in aggregate hours (Chetty et al., 2011).

The government spending measure $G_{\ell,t}$ represents defense contracts from FPDS, aggregated by region-year. Defense contracts are proxied by contracts awarded by the Department of Defense (agency code 97). Location is identified by means of the primary place of performance zip code and, when missing, the recipient zip code. Zip code to MSA cross-walks are used to identify the final location. Unlike Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020), who spread Department of Defense contracts over their duration, we assign contracts to their award date us-

[7] Access to a state's microdata is automatic upon BLS approval for certain Cooperative Agreement Signatory States, but requires separate state approval from non-signatory states. The states for which we do not have access are Florida, Kentucky, Massachusetts, Mississippi, New York, North Carolina, Rhode Island, and Vermont. The restricted access to the researchers' program has been discontinued by the BLS.

[8] Results are robust if we use personal income as a normalizing weight, as in Muratori, Juarros, and Valderrama (2023). Using wages and salaries is more similar to the use of earnings as the normalization factor in Auerbach, Gorodnichenko, and Murphy (2020).

ing FPDS data. As we show in Appendix B.2 and discuss in later sections, this assumption is innocuous for our main results.

The estimand of interest, β_h, measures the percentage increase in regional employment in response to a 1% increase in defense spending relative to wages and salaries.

Identification. Two main concerns arise in studying the impact of federal procurement spending on employment. First, the geographical distribution of contracts may be partly endogenous due to political factors (Mintz, 1992). Second, large federal contractors may forecast future regional government demand and increase production in anticipation of future contracts, meaning that it is necessary to isolate a shock that is not only exogenous to the state of the economy but also unanticipated (Ramey, 2011; Auerbach, Gorodnichenko, and Murphy, 2020). This is particularly true for federal procurement spending, because most of it is awarded through long-term agreements, as illustrated in Figure 2.

In other words, the main regressor in Equation (1) contains an endogenous and potentially anticipated numerator. To address these concerns, Auerbach, Gorodnichenko, and Murphy (2020) construct an instrument that replaces the endogenous and potentially anticipated numerator, $G_{\ell,t+h} - G_{\ell,t-1}$, with an exogenous and unanticipated counterpart. Specifically, the regional change in defense contracts is replaced by the national change in defense contracts, $G_t := \sum_\ell G_{\ell,t}$, reallocated across regions using the long-run exposure of contracts flowing to each region, $\exp_\ell$:

$$\exp_\ell := \frac{1}{19} \sum_{t=2001}^{2019} \frac{G_{\ell,t}}{G_t}.$$

The resulting instrument is:

$$Z_{\ell,t+h} := \frac{\exp_\ell \cdot (G_{t+h} - G_{t-1})}{Y_{\ell,t-1}},$$

which corresponds to the original main regressor of Equation (1), but with the numerator purged of potentially endogenous and anticipated components.

Papers in this literature have referred to this instrument as a Bartik-type instrument (Bartik, 1991) or a shift-share instrument. However, it has been noted that this instrument, $Z_{\ell,t+h}$, does not constitute an exact shift-share setup.[9] While there is nothing problematic about this instrument, which remains widely used, employing an instrument that corresponds exactly to the shift-share framework is a useful robustness exercise because the econometric properties of such instruments are well understood (Goldsmith-Pinkham, Sorkin, and Swift, 2020; Borusyak, Hull, and Jaravel, 2022). In Appendix C, we demonstrate that baseline estimates of the employment multiplier using an exact shift-share instrumental variable are consistent with our baseline estimates.

Instrument Time-Variation. Figure 3 shows the real value of defense contracts directed to MSAs from 2001 to 2019. Differences in the level of this variable are used to create time variation in the

[9] We thank Gabriel Chodorow-Reich for raising this point when discussing the paper at the 2025 NBER Conference: Fiscal Dynamics of State and Local Governments.

numerator of the instrument.

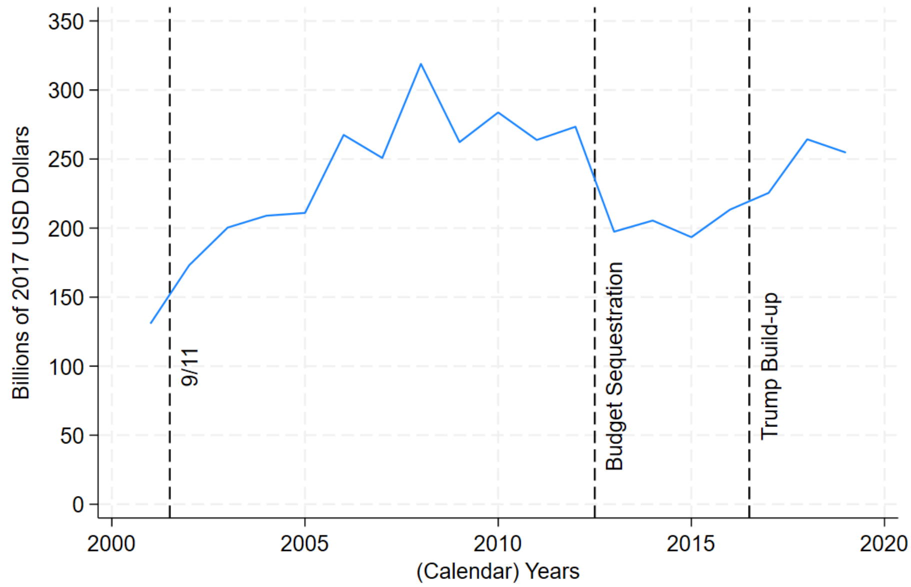

FIGURE 3 — DEFENSE CONTRACTS IN MSAS BY YEAR

National changes in defense expenditure are largely driven by exogenous geopolitical and fiscal events. In the early 2000s, defense spending increased rapidly following the 9/11 terrorist attacks as the U.S. embarked on wars in Afghanistan and Iraq. In the aftermath of the financial crisis, concerns about the worsening government deficit led the Obama administration to announce budget cuts in 2011, which were repeatedly delayed until March 2013, when the unanticipated failure of budget negotiations led to immediate spending cuts ("budget sequestration"). The rise of ISIS and Russia's annexation of Crimea in 2014 intensified concerns in Washington regarding U.S. military readiness after several years of constrained budgets. These developments contributed to bipartisan support for the Bipartisan Budget Act of 2015, which raised discretionary spending caps, and for the National Defense Authorization Act for Fiscal Year 2016. Furthermore, the establishment of unified Republican control of Congress and the Presidency after President Trump's election in 2016 triggered a substantial increase in defense procurement spending, essentially restoring pre-sequester spending levels. The fiscal policy literature has labeled these events as exogenous to output variation.[10]

These examples illustrate that, although the short time series ($T = 19$) may raise concerns about coincidental correlations—such as the overlap of the 9/11 build-up with the Dotcom crash—there is no systematic countercyclical pattern between shifts in national defense spending and national economic growth. Indeed, defense spending cuts in 2013 occurred despite modest growth (nominal GDP grew by about 2%), while the 2016 military build-up coincided with strong growth.

[10] Ramey and Zubairy (2018) record large positive shocks after 9/11 and a sharp negative shock in 2013 due to sequestration in their defense news shock series. Alesina, Favero, and Giavazzi (2014) list the Budget Sequestration Act as an exogenous expenditure-based fiscal consolidation. Briganti, Brunet, and Sellemi (2025) and Amodeo and Briganti (2025) provide a full narrative of major events driving the path of U.S. defense spending in the post-2000 sample, including those discussed here.

Moreover, the inclusion of time fixed effects ($\lambda_{t,h}$) in our baseline specification ensures that any remaining aggregate variation, potentially correlated with national shifts in defense spending and employment growth rates, is fully absorbed. This benefit comes at the cost of interpreting our employment multiplier not as a national-level multiplier but rather as a cross-sectional multiplier. Cross-sectional multipliers can be viewed as a lower bound of deficit-financed, no-monetary-policy, closed-economy national multipliers (Chodorow-Reich, 2019).

Instrument Cross-Sectional Variation. The instrument relies on the long-run regional exposure to defense contracts, $\exp_\ell$, to distribute national-level shifts across regions. Given our relatively short time dimension compared with the cross-sectional size of the sample ($N \gg T$), we are particularly focused on verifying the plausible exogeneity of the exposures. Figure 4 illustrates their geographic distribution.

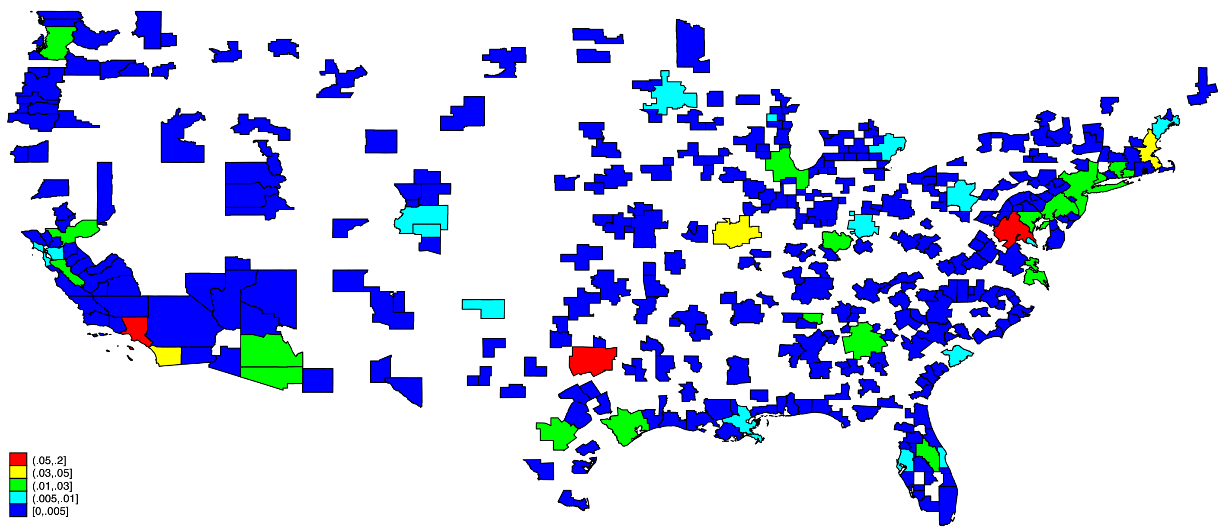

FIGURE 4 — MSA DISTRIBUTION OF LONG-RUN EXPOSURE OF DEFENSE CONTRACTS ($\exp_\ell$)

Notes: The figure shows the value of the long-run exposure to national-level shocks to defense contracts, used in the numerator of the instrument ($\exp_\ell := \frac{1}{T}\sum_{t=2001}^{2019} \frac{G_{\ell,t}}{\sum_\ell G_{\ell,t}}$). Red means long-run exposure larger than 5%. Yellow is from 3% to 5%; green is from 1% to 3%; cyan is from 0.5% to 1%; blue is less than 0.5%. The figure omits Hawaii and Alaska for graphical purposes. Sample goes from 2001 to 2019 and uses all 380 MSAs.

The MSA with the largest share of contracts is Washington–Arlington–Alexandria, accounting for about 12% of Department of Defense spending. This concentration reflects the presence of the Pentagon and numerous other military installations, whose location was determined by proximity to the capital. More broadly, MSAs with high defense contract exposure are typically characterized by long-standing military activities whose location was determined by geostrategic, rather than economic, considerations well before the start of our sample period.[11] Thus, the geographic allocation of national funds across regions is plausibly pre-determined relative to current economic conditions.

[11]Other large MSAs include Fort Worth-Dallas (about 7%, Naval Air Station Joint Reserve Base Fort Worth, 1942), Los Angeles (5%, Los Angeles Air Force Base, 1962), San Diego-Carlsbad (3%, Naval Base San Diego, 1918), and St. Louis (3.1%, Scott Air Force Base, 1917). The locations of these installations, established long before the sample period, reflect strategic rather than economic considerations.

In Appendix C, we show that our regional exposures, plotted in Figure 4, are correlated with the shares of a special case shift-share instrument, constructed as long-run averages of the ratio of regional defense contracts to regional wages and salaries: $w_\ell := \frac{1}{19}\sum_{t=2001}^{2019} \frac{G_{\ell,t}}{Y_{\ell,t}}$. We also show that there is no systematic difference in the time-average annual employment growth rates, LHS of Equation (1) with $h = 0$, between regions with low and high shares. Moreover, even if such differences were to exist, our baseline regression includes region fixed effects ($\alpha_{\ell,h}$), which absorb any time-invariant heterogeneity in employment growth across MSAs.

Baseline Estimates. We calculate the regional employment multiplier by estimating Equation (1) with 2SLS, using $Z_{\ell,t+h}$ as an instrument for the regional change in defense contracts. The left panel of Table 2 reports the 2SLS estimates of Equation (1).

TABLE 2 — REGIONAL EMPLOYMENT MULTIPLIERS - BASELINE ESTIMATES

Response of Total Employment from (Public) BEA Data					
Horizon	*Coefficient (β_h)*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.034	0.021	29.232	1.111	$899,756
	(0.015)			(0.481)	($389,098)
1 year	0.104	0.006	53.120	3.440	$290,733
	(0.038)			(1.244)	($105,111)
2 years	0.099	0.016	27.603	3.275	$305,332
	(0.041)			(1.351)	($125,994)
3 years	0.107	0.026	21.063	3.524	$283,746
	(0.048)			(1.581)	($127,316)

Notes: Sample: 2001-2019; 358 MSAs (QCEW+BDS+LAUS Harmonized Sample). GDP price deflator from BEA, base year 2008. Robust standard errors in parentheses, clustered at the MSA level. Results using either two-way clustered standard errors by state and year or using the Driscoll and Kraay (1998) standard errors are of similar magnitudes. Montiel Olea and Pflueger (2013) effective F is calculated with `weakivtest`. Number of Job-Years refers to one million $. Standard errors of cost-per-job are obtained with the Δ-method.

First, values of the F-statistics are generally above the suggested critical value of 23, indicating no weak-instrument problem (Montiel Olea and Pflueger, 2013). Second, multipliers on impact are estimated precisely, but their magnitude is economically small.

The three-year employment multiplier is approximately 0.1, meaning that increasing regional defense contracts by 1% of regional wages and salaries, results in a 0.1% increase in regional employment. To gain interpretability, we construct the implied number of job-years as follows:

$$\text{job-years}_h := \frac{\beta_h}{N\,T}\sum_{\ell=1}^{N}\sum_{t=2001+1}^{2019}\frac{\$1,000,000}{Y_{\ell,t-1}}\cdot E_{\ell,t-1},$$

where β_h represents the estimated employment multipliers in the second column of Table 2. The reciprocal of the number of job-years gives us the estimate of the cost-per-job consistent with Chodorow-Reich (2019) and Muratori, Juarros, and Valderrama (2023).

Our two-year employment multiplier implies a cost-per-job of approximately $305,000 in 2008 dollars, substantially higher than benchmark estimates of the employment impacts of, for example, fiscal transfers. We discuss this comparison in detail, alongside other recent estimates, in subsection III.2. The main message of this first set of empirical results is:

Fact 1. *Defense procurement in the 21st century raises regional employment at a high cost per job.*

Robustness Checks We evaluate the sensitivity of our estimates along several dimensions. Our estimates are stable across samples: we repeat the baseline analysis using either the largest or smallest available samples, finding similar magnitudes and precision (Appendix B.1). Our baseline assigns the full contract value to the award year, which may overstate the short-run fiscal cost. Spreading contracts over their duration (Appendix B.2) lowers the cost-per-job to approximately $330,000 on impact, with estimates at longer horizons within the same range, confirming that the high cost-per-job is not an artifact of this timing choice. Another concern is that a small number of highly exposed MSAs may drive the results. Excluding the top 5% of MSAs by exposure (top 18 MSAs) produces qualitatively similar, though less precise, estimates (Appendix B.3), indicating that our findings do not depend on a handful of defense-intensive regions. Appendix B.5 reports two additional checks supporting our baseline identification. First, we report bin-scatter plots and first-stage regression output, indicating that our instrument is strong at all horizons. Second, the instrument is not predicted by lagged employment growth, which alleviates concerns that more exposed regions were already on differential pre-trends prior to the spending shock. Finally, the results are also robust to how we construct the instrument. Replacing our baseline instrument with the exact shift-share instrument yields nearly identical estimates, as does using initial (2001) rather than time-averaged exposure shares (Appendix C).

III.1. Origin of the Positive Multiplier.

The positive MSA-level employment multiplier indicates that, when additional government contracts are awarded in a region, employment is not merely reallocated from other firms, since such reallocation would yield a zero regional multiplier. Instead, extra workers must come from one of three sources: (i) the pool of unemployed, (ii) new entrants to the labor force from the existing regional population, or (iii) other regions, through migration or cross-MSA commuting.

Regarding channels (i) and (ii), in Appendix B.6 we use Local Area Unemployment Statistics (LAUS) data to show that increases in the regional labor force and reductions in unemployment explain the bulk of the additional employment, with labor force changes playing a major role. Moreover, in Appendix B.7 we also show, using Business Dynamics Statistics (BDS) data, that the creation of new firms (the extensive firm margin) is negligible over the horizon we study. This indicates that existing firms, rather than newly created firms, are hiring individuals from the labor force and the pool of unemployed, generating the positive employment multiplier. Both results are in line with Auerbach, Gorodnichenko, and Murphy (2024) who also find significant and positive effects on the labor force and reductions in unemployment rate using data from the American Community Survey (ACS) and find insignificant effects on the number of establishments using data from U.S. Census.

Regarding channel (iii), we find that migration plays a limited role: we find no effect on population three years after the shock, a result in line with Foschi, House, and Proebsting (2025) who also show that MSA-level population does not respond to regional shocks to defense procurement up to the five-year horizon.[12] This is consistent with the temporary nature of the shocks and the fact that MSAs approximate self-contained labor markets with limited cross-regional commuting.

III.2. Discussion and Comparative Analysis.

Table 3 presents a range of cost-per-job estimates from recent empirical studies of fiscal interventions, with a particular focus on defense procurement and ARRA transfers.[13]

TABLE 3 — COST-PER-JOB - REVIEW OF ESTIMATES

Study	*Type of G*	*Sample*	*Geography*	*Job-Years*	*Cost-per-Job*
Nakamura and Steinsson (2014)	Defense Contracts	1966-2006	US States	2.23	$44,836
Dupor and Guerrero (2017)	Defense Contracts	1951-2014	US States	0.05	$1,910,092
Demyanyk, Loutskina, and Murphy (2019)	Defense Contracts	2007-2009	828 US CBSAs	0.67	$148,711
Auerbach, Gorodnichenko, and Murphy (2020)	Defense Contracts	2001-2016	383 US MSAs	0.42	$236,822
Muratori, Juarros, and Valderrama (2023)	Defense Contracts	1979-2019	US MSAs	1.43	$69,817
Park, Zhou, and Zubairy (2025)	Defense Contracts	2011-2020	828 US Counties	0.29	$343,659
Wilson (2012)	ARRA Transfers	2009-2010	US States	0.80	$123,839
Conley and Dupor (2013)	ARRA Highway Funding	2009-2011	US States	0.76	$131,578
Serrato and Wingender (2016)	Population Revisions	1980, 1990, 2000	US Counties	3.25	$30,785
Dupor and Mehkari (2016)	ARRA Subcomponents	2008-2010	US Commuting Reg.	0.95	$104,931
Adelino, Cunha, and Ferreira (2017)	Local Spending	2007-2013	US Municipalities	-	$25,000
Dupor and McCrory (2018)	ARRA Subcomponents	2008-2010	US Commuting Reg.	1.85	$54,054
Chodorow-Reich (2019)	ARRA Transfers	2008-2010	US States	2.01	$49,750

Notes: Source is Chodorow-Reich (2019) and authors' calculations using estimates from listed papers. Job-Years per $100,000 calculated at two-year or closest available horizon and using 2008 dollars. For Nakamura and Steinsson (2014) results are from Table 3, Row 1, which uses the Bartik instrument for prime military contracts; for Park, Zhou, and Zubairy, 2025 we use estimates unadjusted for subcontracting for comparability. Geography refers to main estimates. Estimates from Dupor and Guerrero, 2017 depend heavily on inclusion of years 1953-1954 in the sample, as noted in that paper.

Using Chodorow-Reich (2019) as a benchmark is useful because it surveys the cost-per-job estimates commonly emphasized in the recent fiscal-policy literature. As can be seen in Table 3, our cost-per-job estimates for defense spending are not outliers; other contemporary estimates of the cost per job from defense procurement are similarly high. For example, the estimates of Auerbach, Gorodnichenko, and Murphy (2020) correspond to a cost-per-job at the two-year horizon of roughly $237,000 in 2008 dollars. These figures are broadly consistent with our estimate of $305,000 at the same horizon. Moreover, in contemporaneous work, Park, Zhou, and Zubairy (2025) use quarterly defense contracts data from 2011 to 2024 and estimate a cost-per-job of about $343,000 in 2008 dollars at the county level. This evidence suggests that military procurement in the twenty-first century has become significantly less

[12] Auerbach, Gorodnichenko, and Murphy (2024) find insignificant population effects at one year but significant effects at three years. Our null result at longer horizons may reflect differences in sample composition and time coverage.

[13] Other estimates from non-U.S. data include: Corbi, Papaioannou, and Surico (2019), who estimate $8,000 per year using Brazilian municipal transfers; Buchheim and Watzinger (2023), who find $24,000 per year from German public investment in school energy efficiency; and Gabriel, Klein, and Pessoa (2023), who find €30,000 per year from regional Eurozone data.

cost-effective at generating employment compared with either earlier waves of defense spending (Nakamura and Steinsson, 2014; Muratori, Juarros, and Valderrama, 2023) or other forms of public stimulus, such as transfers (Chodorow-Reich et al., 2012).[14]

A natural question is how to interpret the relatively high cost-per-job estimates. One plausible explanation is a structural shift in the industrial and occupational composition of defense-related activity in the twenty-first century. If the first-order employment effects of defense contracts are concentrated within contractors operating in high-paying, skill-intensive industries, the implied cost-per-job will naturally be higher. Consistent with this mechanism, Section V shows that the initial employment gains are concentrated within contractors. Complementary evidence from Bartal and Becard (2024) document that modern U.S. defense contractors are increasingly concentrated in high-tech sectors such as aerospace, software engineering, cybersecurity, and advanced manufacturing. These sectors employ highly educated professionals, implying a higher wage bill per job created.

Consistent with this interpretation, Appendix B.8 estimates the effect of regional defense procurement shocks on wage-per-worker. We find that a procurement shock equal to 1% of regional wages raises the average wage-per-worker by approximately 0.14% after one year, implying that marginal hires are on average paid more than the pre-existing mean wage.[15] This is consistent with procurement-driven demand generating employment in relatively high-paying firms and sectors, which naturally raises the implied fiscal cost of each additional job.

Alternative Explanations. We also note that the high cost-per-job does not appear to be an artifact of three potential measurement concerns. First, subcontracting across geographic units could bias regional multipliers downward, but Park, Zhou, and Zubairy (2025) show that adjusting for subcontracting reduces the cost-per-job only modestly, from roughly \$344,000 to \$282,000 in 2008 dollars, a value only slightly lower than ours and still well above estimates from non-defense spending. Second, cross-MSA spillovers could cause our estimates to miss employment generated in neighboring regions. Using the outflow estimates of Auerbach, Gorodnichenko, and Murphy (2020), accounting for spillovers would reduce the cost-per-job to roughly \$160,000 in 2008 dollars,[16] which remains above most estimates for ARRA transfers in Table 3. Third, assigning the full contract value to the award year could result in higher estimates of cost per job, as the full value of the contract is front-loaded in the cost-per-job calculation. Our choice is motivated by the need to avoid missing the effects of new contracts on contractors' production and hiring decisions, which in the context of defense procurement typically occur before the advent of later deliveries and payments (see Briganti, Brunet, and Sellemi, 2025).

[14] It should be noted that the impact of transfers may vary significantly with economic conditions. For example, recent research suggests that federal-state transfers during the COVID-19 pandemic produced additional employment at a cost of roughly \$225,000 per job-year in 2019 dollars (Clemens, Hoxie, and Veuger, 2025).

[15] Or, less plausibly, that low-skilled labor is laid off during an expansion.

[16] Let the direct effect be θ and the conversion factor be c, so that the cost-per-job is $1/(\theta \cdot c)$. If the outflow effect is approximately one-half of the direct effect, the total effect becomes $\theta + \frac{1}{2}\theta = \frac{3}{2}\theta$. Hence, the cost-per-job accounting for spillovers is $\frac{1}{\left(\frac{3}{2}\theta \cdot c\right)} = \frac{2}{3} \cdot \frac{1}{\theta \cdot c}$, that is, two-thirds of the direct-effect estimate.

Reassuringly, spreading contracts over their duration yields similar cost-per-job estimates at horizons beyond impact (Appendix B.2), consistent with Demyanyk, Loutskina, and Murphy (2019) and Auerbach, Gorodnichenko, and Murphy (2020), who also find that this choice does not materially affect estimates.

While the shift in the industrial composition of defense contracting could plausibly explain the higher cost-per-job estimates observed in the post-2000 period, a definitive explanation lies beyond the scope of this paper. Instead, our contribution is to document this fact and to provide an empirical breakdown of the employment multiplier, highlighting where employment gains are primarily concentrated.

IV. Breaking Down the Employment Multiplier by Firm Size

Recent work on fiscal policy shocks transmission emphasizes the role of small firms. In particular, several studies argue that small firms may amplify fiscal multipliers because procurement contracts relax financial constraints and generate larger firm-level responses than for large firms (Hebous and Zimmermann, 2020; Budrys, 2022; Gabriel, 2024). di Giovanni et al. (2026) show that procurement contracts help small firms relax financial constraints, although policies that tilt procurement toward small, rather than large, firms may have non-trivial macroeconomic effects on output. Relatedly, Juarros (2022) shows that regional fiscal multipliers are larger in areas with a higher small-firm employment share. Yet no previous work quantifies how much of the shock-induced increase in defense contracts accrues to small firms, nor how much of the resulting employment growth they account for. Without such estimates, it is difficult to gauge whether the financial-accelerator channel constitutes a first-order driver of the aggregate regional employment response.

We show that regional defense spending shocks primarily raise employment through large firms. Large firms account for nearly all of the measured employment response because they receive the bulk of the incremental contracts induced by regional shocks. This finding does not rule out financially important firm-level responses among small firms. Rather, it shows that, in the aggregate, the first-order employment effects of defense procurement are concentrated in large firms.

Fact 2. *Regional defense spending shocks primarily affect large firms, which receive more than 70% of total contracts and account for over 80% of the estimated employment gains.*

Breakdown by Firm Size: Who Gets the Employment Response? To quantify the contribution of firms of different size, we use the Business Dynamics Statistics (BDS), which report MSA-year employment by firm size. Following the BDS classification, small firms employ fewer than 20 workers, medium-sized firms employ 20–499 workers, and large firms employ at least 500 workers. BDS classifies firm size at the firm level by aggregating employment across establishments within the same firm. If firm X operates a 10-employee establishment in region A and a 20-employee establishment in region B, it is classified as a medium-sized firm because total firm employment is 30. The 10 employees in A and the 20 in B are therefore both counted as medium-firm employment in their respective regions.

We then decompose the left-hand side of Equation (1) into the contribution of each size category:

$$\frac{E_{\ell,t+h} - E_{\ell,t-1}}{E_{\ell,t-1}} \equiv \frac{E^{\text{Small}}_{\ell,t+h} - E^{\text{Small}}_{\ell,t-1}}{E_{\ell,t-1}} + \frac{E^{\text{Medium}}_{\ell,t+h} - E^{\text{Medium}}_{\ell,t-1}}{E_{\ell,t-1}} + \frac{E^{\text{Large}}_{\ell,t+h} - E^{\text{Large}}_{\ell,t-1}}{E_{\ell,t-1}}$$

We then re-estimate Equation (1) for each size group:

$$\frac{E^{\text{Small}}_{\ell,t+h} - E^{\text{Small}}_{\ell,t-1}}{E_{\ell,t-1}} = \beta^{\text{s}}_h \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda^{\text{s}}_{t,h} + \alpha^{\text{s}}_{\ell,h} + u^{\text{s}}_{\ell,t+h}, \tag{2}$$

$$\frac{E^{\text{Medium}}_{\ell,t+h} - E^{\text{Medium}}_{\ell,t-1}}{E_{\ell,t-1}} = \beta^{\text{m}}_h \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda^{\text{m}}_{t,h} + \alpha^{\text{m}}_{\ell,h} + u^{\text{m}}_{\ell,t+h}, \tag{3}$$

$$\frac{E^{\text{Large}}_{\ell,t+h} - E^{\text{Large}}_{\ell,t-1}}{E_{\ell,t-1}} = \beta^{\text{l}}_h \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda^{\text{l}}_{t,h} + \alpha^{\text{l}}_{\ell,h} + u^{\text{l}}_{\ell,t+h}. \tag{4}$$

Estimating Equations (2)–(4) by 2SLS yields the coefficients β^{s}_h, β^{m}_h, and β^{l}_h, which measure the contribution of small, medium-sized, and large firms to the regional employment response. Table 4 reports these estimates. By construction, the size-specific coefficients add up to the total employment multiplier estimated using BDS employment data, $\beta^{\text{s}}_h + \beta^{\text{m}}_h + \beta^{\text{l}}_h = \beta_h$.[17] The BDS-based employment multipliers (β_h)— reported in Appendix D— are qualitatively similar to the baseline estimates obtained using BEA total employment, confirming that BDS data replicate the main findings of Section III.

TABLE 4 — EMPLOYMENT MULTIPLIER IS DRIVEN BY LARGE FIRMS

	Small Firms			*Medium-Sized Firms*			*Large Firms*		
Horizon	*Coefficient* (β^s_h)	*p*	*Fraction (%)*	*Coefficient* (β^m_h)	*p*	*Fraction (%)*	*Coefficient* (β^l_h)	*p*	*Fraction (%)*
impact	-0.003	0.718	-4.1%	0.000	0.972	0.8%	0.064	0.034	103.4%
	(0.007)			(0.013)			(0.030)		
1 year	0.006	0.336	5.2%	0.007	0.520	6.7%	0.095	0.013	88.1%
	(0.006)			(0.011)			(0.038)		
2 years	0.005	0.392	4.5%	0.014	0.155	12.2%	0.095	0.019	83.3%
	(0.006)			(0.010)			(0.041)		
3 years	0.006	0.341	4.6%	0.012	0.224	10.1%	0.103	0.036	85.3%
	(0.006)			(0.010)			(0.049)		

Notes: Sample: 2001-2019 - 373 MSAs (BDS Dataset). Data source: Business Dynamics Statistics (BDS). Estimates of Equations (2) (β^s_h), (3) (β^m_h) and (4) (β^l_h). Fraction is calculated as $\beta^{\text{s}}_h/\beta_h$ for small firms, $\beta^{\text{m}}_h/\beta_h$ for medium-sized firms and $\beta^{\text{l}}_h/\beta_h$ for large firms, using the values of β_h from Table D1 (values reported in green).

Table 4 shows that the employment response is overwhelmingly concentrated in large firms. Across all horizons, the estimated responses of small and medium-sized firms are close to zero and statistically insignificant, whereas the response of large firms is positive and precisely estimated. For example, at the one-year horizon, the decomposition yields $\beta^{\text{s}}_1 = 0.006$, $\beta^{\text{m}}_1 = 0.007$, and $\beta^{\text{l}}_1 = 0.095$, implying that large firms account for nearly 90% of the total employment response. This pattern is robust across alternative samples.[18]

[17] This follows from the linearity of 2SLS, the linearity of local projections, and the fact that the right-hand side is identical across Equations (2)–(4).

[18] Results using the full sample and results for the smaller harmonized sample are presented in Appendix D. The decomposition remains broadly consistent across samples, although there is some evidence of modest medium-sized-firm responses

Breaking Down Government Spending by Firm Size: Who Gets the Spending Shock? We next show that the dominance of large firms in Table 4 primarily reflects the allocation of shock-induced contracts toward large firms. In other words, large firms account for most of the employment response not because defense spending necessarily has much larger per-dollar effects on them, but because regional shocks predominantly increase contracts awarded to large firms.

To assess this mechanism directly, we decompose MSA-level defense contracts, $G_{\ell,t}$, by recipient-firm size and examine how a regional shock changes contract awards across size categories. We do so using the National Establishment Time Series (NETS), a privately maintained census of U.S. establishments. After applying cleaning procedures following Barnatchez, Crane, and Decker (2017) (described in Appendix F.6), we match FPDS contract data to NETS firms using the DUNS number of the recipient or, when necessary, the DUNS number of the recipient's parent company.[19] This procedure matches 97.6% of total FPDS defense contracts by value.

Once each contract is assigned to a NETS firm, we can aggregate defense contracts in each MSA to a firm size bin, following the BDS definitions of small, medium, and large firms:

$$\frac{G_{\ell,t+h} - G_{\ell,t-1}}{G_{\ell,t-1}} \equiv \frac{G^{\text{Small}}_{\ell,t+h} - G^{\text{Small}}_{\ell,t-1}}{G_{\ell,t-1}} + \frac{G^{\text{Medium}}_{\ell,t+h} - G^{\text{Medium}}_{\ell,t-1}}{G_{\ell,t-1}} + \frac{G^{\text{Large}}_{\ell,t+h} - G^{\text{Large}}_{\ell,t-1}}{G_{\ell,t-1}}$$

We then run parallel estimates to Equations (2)-(4) but using spending rather than employment by category as the outcome variable. Results are presented in Table 5.

TABLE 5 — REGIONAL DEFENSE SHOCKS MOSTLY AFFECT CONTRACTS AWARDS TO LARGE FIRMS

	Small Firms ($G^{\text{Small}}_{i,t}$)			*Medium-Sized Firms* ($G^{\text{Medium}}_{i,t}$)			*Large Firms* ($G^{\text{Large}}_{i,t}$)		
Horizon	*Coefficient*	*p-value*	*Fraction*	*Coefficient*	*p-value*	*Fraction*	*Coefficient*	*p-value*	*Fraction*
impact	0.024	0.045	2.4%	0.136	0.151	13.7%	0.832	0.000	83.8%
	(0.012)			(0.095)			(0.102)		
1 year	0.043	0.005	4.3%	0.184	0.021	18.6%	0.763	0.000	77.1%
	(0.015)			(0.079)			(0.084)		
2 years	0.038	0.006	3.9%	0.178	0.037	18.0%	0.773	0.000	78.2%
	(0.014)			(0.085)			(0.092)		
3 years	0.047	0.012	4.7%	0.196	0.043	19.9%	0.745	0.000	75.4%
	(0.018)			(0.097)			(0.107)		

Notes: Sample: 2001-2019 - 373 MSAs (BDS Dataset).

Table 5 shows that roughly 80% of the increase in contracts induced by a regional shock accrues to large firms, and only the large-firm component is precisely estimated at all horizons. This pattern suggests that the dominance of large firms in the employment breakdown is primarily driven by the allocation of shock-induced contracts toward large firms, rather than by sharply larger employment elasticities within that size class.

An additional implication of Tables 4 and 5 is that small and medium-sized firms together receive

in the most restrictive sample.

[19] NETS is fundamentally an establishment-level database. As in Barnatchez, Crane, and Decker (2017), we aggregate establishments to the ultimate headquarters establishment, which we treat as the firm. We thank Joonkyu Choi for sharing his insights on this dataset.

about one-fifth of the shock-induced contracts, yet account for little of the measured employment response. One possible explanation is subcontracting: Park, Zhou, and Zubairy (2025) document that small and medium-sized firms often subcontract work to large firms. If so, part of the employment generated by contracts initially awarded to smaller firms may ultimately materialize within large firms.

V. Breaking Down the Employment Multiplier Between Contractor and Non-Contractor Effects

Government procurement can affect employment through several channels: (1) direct effects on hiring at recipient establishments; (2) indirect effects on hiring at other establishments belonging to contractor firms; and (3) local demand spillovers to firms that never receive contracts (non-contractors). Figure 5 illustrates these three groups of firms.

FIGURE 5 — BREAKDOWN OF REGIONAL EMPLOYMENT: CONTRACTORS, RECIPIENT ESTABLISHMENTS, AND NON-CONTRACTORS

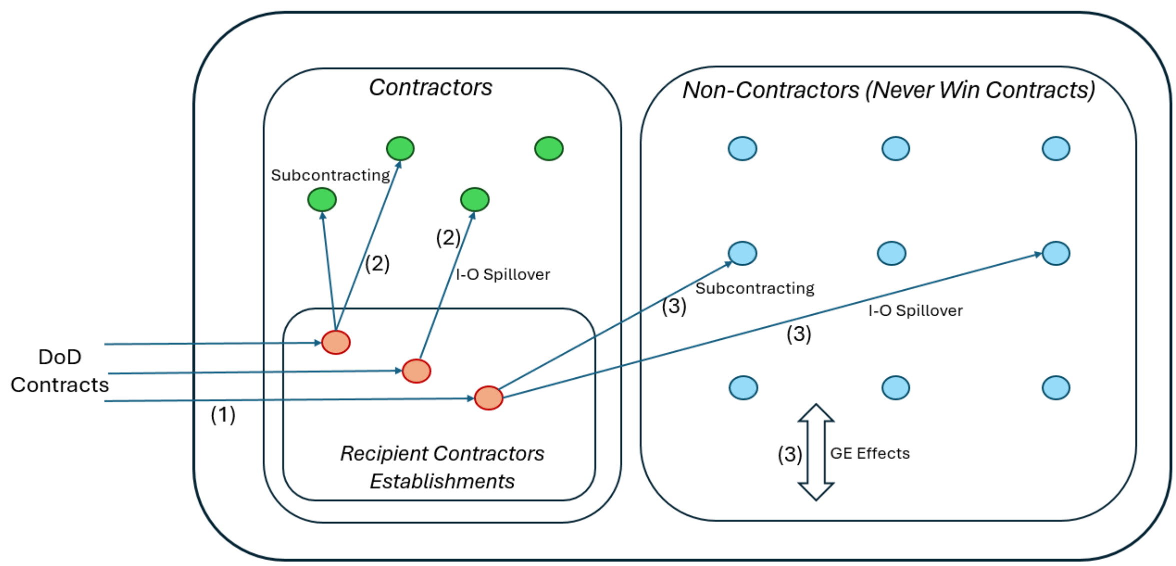

Notes: Red dots denote recipient contractor establishments that directly receive contracts (channel 1). Green dots denote other establishments belonging to contractor firms that do not directly receive contracts (channel 2). Blue dots denote non-contractors, i.e., firms that never receive defense contracts (channel 3). The contractor group comprises both red and green dots; red dots are a subset of this group.

Contractors are firms that receive a defense contract at some point in time (green and red dots in the figure). Among them, a subset of establishments directly receives contracts (red dots), and employment changes at these establishments capture the direct effect of procurement on awardees. Section VI uses establishment-level data to estimate the contribution of these establishments to the overall employment multiplier.

The remaining firms in the figure (blue dots) are non-contractors, defined as firms that never receive defense contracts. Even though these firms do not directly receive procurement awards, their employment may still respond to regional defense spending shocks through local demand spillovers, including general demand stimulus, input–output linkages, or subcontracting relationships.

Policy Relevance. Breaking down employment gains across these groups is important for assessing how broadly the effects of government procurement extend beyond the set of firms that directly interact with the Department of Defense. While some papers focus on specific propagation channels—such as input–output linkages (Auerbach, Gorodnichenko, and Murphy, 2020; Barattieri, Cacciatore, and Traum, 2025) or subcontracting relationships (Park, Zhou, and Zubairy, 2025)—our setting allows for a more comprehensive assessment of how widespread the employment effects of procurement are within the local economy. In particular, by directly observing employment in firms that never receive defense contracts, we can measure whether procurement shocks generate employment gains outside the pool of contractors, providing a direct estimate of local employment spillovers.

This decomposition also helps clarify the mechanisms underlying the aggregate employment multiplier. If employment gains are concentrated within recipient establishments or contractor firms, the multiplier primarily reflects direct hiring effects, supply-chain effects, and subcontracting relationships within the defense production network. By contrast, if employment gains extend to non-contractors, this would point to broader local demand spillovers operating through general equilibrium channels. Distinguishing between these mechanisms is therefore essential for understanding both the incidence of procurement spending and the extent to which its employment effects diffuse through the local economy.

Moreover, understanding the relative importance of these channels has implications for policy design. If contractors effects dominate, policies aimed at stimulating local employment should focus on identifying the types of contracts and firms that generate the largest employment responses. Conversely, if spillovers to non-contractors play a larger role, optimal policy may depend more on the regional allocation of spending and on local economic conditions that amplify demand spillovers (for example, the prevalence of financially constrained firms).

Breakdown by Contractor Status. We leverage our access to the restricted microdata, covering the universe of U.S. establishments, to break down the regional employment time-series into contractor and non-contractor components, separating channels (1)-(2) from channel (3).[20]

In practice, we match the universe of defense contractors from FPDS with the universe of establishments from the restricted QCEW. The matching is restricted to establishments in MSAs located entirely within the 42 signatory states (plus Washington, D.C.) to which we were granted data access, observed between 2006 and 2019. It is carried out using an algorithm described in Appendix F.1.

[20] These data are no longer available through the BLS due to organizational changes within the agency.

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

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

Essentially, we identify the set of establishments, $\mathcal{C}$, that receive at least one government contract over the sample period and aggregate employment for these establishments into the component associated with defense contractors. The residual employment component represents establishments that were never directly involved with defense contracting. Because this definition is based on direct matches between FPDS prime contractors and LDBE establishments, some establishments that participate only indirectly in defense procurement, for example as subcontractors without observed prime awards, may remain in the non-contractor group.

Notice that this contractors versus non-contractors classification uses information from the full sample period, which raises a potential concern about conditioning on a post-treatment outcome. However, using matched NETS-FPDS data, we find that contractor status is highly persistent, mitigating this concern.[21]

Empirical Results. We match almost all defense contractors to LDBE firms; matched contractors account for more than 90% of total defense spending for each state included in the analysis, and typically more than 95%.[22]

We assign a LDBE firm contractor status if they have been matched to an FPDS contractor establishment in this process, allowing us to aggregate up to MSA-year time series of defense contractors' and non-contractors' employment:

$$\frac{E_{\ell,t+h} - E_{\ell,t-1}}{E_{\ell,t-1}} = \frac{E_{\ell,t+h}^{\text{Contractors}} - E_{\ell,t-1}^{\text{Contractors}}}{E_{\ell,t-1}} + \frac{E_{\ell,t+h}^{\text{Non-Contractors}} - E_{\ell,t-1}^{\text{Non-Contractors}}}{E_{\ell,t-1}}.$$

Therefore, we break down the employment multiplier in two components by estimating the following two equations:

$$\frac{E_{\ell,t+h}^{\text{Contractors}} - E_{\ell,t-1}^{\text{Contractors}}}{E_{\ell,t-1}} = \beta_h^{\text{c}} \cdot \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda_{t,h}^{\text{c}} + \alpha_{\ell,h}^{\text{c}} + u_{\ell,t+h}^{\text{c}} \tag{5}$$

$$\frac{E_{\ell,t+h}^{\text{Non-Contractors}} - E_{\ell,t-1}^{\text{Non-Contractors}}}{E_{\ell,t-1}} = \beta_h^{\text{nc}} \cdot \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda_{t,h}^{\text{nc}} + \alpha_{\ell,h}^{\text{nc}} + u_{\ell,t+h}^{\text{nc}}. \tag{6}$$

[21] We find that 94% of contractor employment in any given year is at firms that also hold contracts in the following year, and this share remains above 90% even at a three-year horizon. Entry of new contractors accounts for only 5% of contractor employment annually. This persistence implies that our "ever contractor" classification closely approximates contemporaneous contractor status, limiting the scope for post-treatment conditioning bias. We interpret the breakdown as descriptive, showing where employment gains appear, rather than identifying separate causal effects for ex-ante defined groups.

[22] Confidentiality of the LDBE data prevents us from publishing a sample of name correspondences, but the filtering algorithm (available on request) was extensively and iteratively refined to eliminate false matches (for example, by eliminating matches on short abbreviations or common words, or where the name overlap was not of sufficient length).

TABLE 6 — EMPLOYMENT MULTIPLIER: CONTRACTOR VS. NON-CONTRACTOR BREAKDOWN

	Contractors			*Non Contractors*		
Horizon	*Coefficient* (β_h^c)	p	*Fraction*	*Coefficient* (β_h^{nc})	p	*Fraction*
impact	0.040	0.053	157.6%	-0.015	0.387	-57.6%
	(0.021)			(0.017)		
1 year	0.055	0.055	57.8%	0.040	0.010	42.2%
	(0.029)			(0.016)		
2 years	0.048	0.087	48.1%	0.052	0.054	51.9%
	(0.028)			(0.027)		
3 years	0.049	0.119	43.6%	0.064	0.060	56.4%
	(0.031)			(0.034)		

Notes: Sample: 2006–2019; 254 MSAs (QCEW+BDS+LAUS+LDBE Harmonized Dataset). Estimates of Equations (5) and (6), obtained by 2SLS with the baseline Bartik instrument, year fixed effects, and MSA fixed effects. Standard errors are reported in parentheses and clustered at the MSA level. Fraction is calculated as β_h^c/β_h for contractors and β_h^{nc}/β_h for non-contractors, using the total LDBE multiplier β_h from Table 7.

β_h^c represents the effect of defense contracts on defense contractors, isolating channels (1)-(2) in Figure 5, while β_h^{nc} represents the broader indirect effect of defense contracts on firms that never won a contract, channel (3), e.g., effect on local businesses such as restaurants and shops, subcontractors and input supplier who never directly engage in procurement. It accounts for a broad measure of the strength of the local spillover effect of demand shocks on employment.

Table 6 reports the contractor and non-contractor estimates, β_h^c and β_h^{nc}. On impact, we find evidence of a mild crowding-out of non-contractor employment ($\beta_0^{nc} = -0.015$) that offsets the positive and significant employment response of contractors ($\beta_0^c = 0.040$). The net impact effect, $\beta_0 = 0.040 - 0.015 \approx 0.026$ (reported in Table 7), is statistically insignificant, meaning that the expansion of contractors employment is approximately offset by the employment losses in non-contractors within a year of the shock. Contractors ramp up production but hire partly from other firms, reallocating workers and yielding a near-zero net effect on total employment.[23]

One year after the shock, both contractor and non-contractor employment expand significantly, with nearly 60% of the response coming from contractors. However, the cumulative non-contractor response (impact plus one year) does not overturn the initial crowding-out: it remains statistically indistinguishable from zero. Two years after the shock, both components are positive and significant, with the effect split roughly in half between the two groups. By three years after the shock, the non-contractor response exceeds the contractor response.

Fact 3. *The employment effects of defense procurement are entirely among contractors within the first year and diffuse only gradually to non-contractors.*

This pattern indicates that defense procurement does not immediately generate broad-based local employment gains. Instead, in the short run it primarily reallocates labor toward contractors, while broader spillovers to non-contractors emerge only with delay. The initial decline in non-contractor

[23]In Appendix E, we replicate the results using the larger, non-harmonized LDBE sample with 262 MSAs, where the evidence of crowding-out is stronger in both economic magnitude and statistical significance.

employment is consistent with Barattieri, Cacciatore, and Traum (2025), who show that procurement shocks can crowd out downstream sectors by raising intermediate input costs. It is also consistent with worker reallocation from non-contractors towards contractors. Our results suggest that these negative spillovers are eventually offset by broader employment gains among firms that never receive defense contracts.

Breakdown Validation. We now verify that the employment multiplier estimates used in the contractor/non-contractor breakdown are comparable across data sources and samples.

The BLS's public data on regional employment, the Quarterly Census of Employment and Wages (QCEW), is constructed by aggregating the LDBE microdata as follows:

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

where i denotes an establishment operating in period t in region ℓ, identified by its physical location address. Hence, we aggregate LDBE employment by regions and use it to estimate Equation (1). The left panel of Table 7 presents the estimates of the employment multiplier using the regionally aggregated micro LDBE data in the harmonized QCEW+BDS+LAUS+LDBE sample, which goes from 2006 to 2019 and covers 254 MSAs located in the 42 signatory states. By construction, estimates in Table 7 (β_h), are numerically equivalent to the sum of the contractors (β_h^c) and non-contractors (β_h^{nc}) reponses in Table 6.

TABLE 7 — LDBE ESTIMATES ARE CONSISTENT WITH ESTIMATES FROM PUBLIC DATA

	Private Emp. from LDBE (Restricted QCEW Data)				Total Emp. from Public BEA Data
Horizon	*Coefficient* (β_h)	*p*	*Effective F*	*Job-Years/\$1M*	*Job-Years/\$1M*
impact	0.026	0.143	10.019	0.601	0.762
	(0.017)			(0.409)	(0.632)
1 year	0.096	0.008	29.845	2.235	2.755
	(0.036)			(0.841)	(1.318)
2 years	0.101	0.042	7.532	2.356	2.670
	(0.049)			(1.150)	(1.573)
3 years	0.113	0.063	6.561	2.644	2.989
	(0.061)			(1.419)	(1.867)

Notes: Sample: 2006–2019; 254 MSAs (QCEW+BDS+LAUS+LDBE Harmonized Dataset). Estimates of Equation (1) using (i) restricted data from LDBE on private employment (left panel) and (ii) the public BEA data on total employment (right panel). Both panels are estimated by 2SLS with the baseline Bartik instrument, year fixed effects, and MSA fixed effects. Standard errors are reported in parentheses and clustered at the MSA level. Job-years are expressed per \$1 million of defense contracts, and underlying contract values are deflated to 2008 dollars using the BEA GDP price deflator.

In principle, the estimates of employment multipliers reported in Table 7 may differ from the baseline employment multipliers in Table 2 for two reasons. First, the employment data are different: regionally aggregated microdata on private employment from the LDBE versus publicly available BEA data on total employment. Second, the estimation samples are different: the harmonized QCEW + BDS + LAUS + LDBE sample covers 2006–2019 and 254 MSAs, while the baseline QCEW + BDS +

LAUS sample covers 2001–2019 and 358 MSAs. Hence, we now make sure that employment multiplier estimates used to break down the response between contractors and non-contractors are comparable across data and samples.

To isolate differences arising solely from the employment data, the right panel of Table 7 reports estimates of job-years obtained using BEA total employment but restricting the sample to the smaller harmonized QCEW + BDS + LAUS + LDBE sample. In this specification, any remaining differences reflect only the choice of employment data.

Comparing the estimates obtained using the two employment measures (highlighted in blue in the table) shows that they are quantitatively very similar and increase over time in both cases. The number of job-years is somewhat larger when using BEA total employment than when using LDBE private employment. This difference is expected, as total employment also captures responses in the public sector.[24] Overall, these results indicate that aggregating establishment-level employment from the LDBE to the regional level produces employment multiplier estimates that are closely aligned with those obtained using BEA total employment data.

We next isolate the role of sample differences. Comparing estimates based on BEA total employment in the smaller harmonized sample (right panel of Table 7) with the baseline estimates (Table 2) shows that both the magnitude and precision of the estimates decline somewhat in the smaller sample; the latter is to be expected given the reduced number of observations. Nevertheless, the estimates remain statistically significant and broadly similar to the baseline results. This provides a reassuring robustness check: even with a substantially smaller sample, the estimated employment multipliers remain comparable.

VI. Breaking Down the Employment Multiplier into the Direct and Indirect Effects within Contractors

The regional analysis shows that employment gains from defense spending shocks are concentrated among contractors within the first years from the shock (Table 6). However, this response combines the direct effect of receiving a contract, channel (1) in Figure 5, with indirect effects within contractor firms, channel (2) in Figure 5. In this section we isolate the direct employment response of recipient establishments using restricted administrative establishment-level employment data. By doing so, we also document the direct effect of contracts on establishments, assessing the persistence of the employment gains relative to the median contract duration.

We find that while contract receipt persistently raises employment at recipient establishments, these direct effects account for only a small fraction of the overall contractor response identified in the regional analysis.

[24] Ramey (2013) and Conley and Dupor (2013) suggest that employment responses may partly reflect increases in public-sector employment rather than new private-sector jobs.

Empirical Methodology and Results

Not all contracts can be treated as demand shocks at the establishment-level: many are anticipated or awarded to positively selected firms, creating both foresight and selection biases. To address this, we decompose total government contracts $G_{i,t}$ awarded in quarter t to establishment i into two components:

$$G_{i,t} = \widetilde{G}_{i,t} + \varepsilon^{g}_{i,t}$$

where $\widetilde{G}_{i,t}$ denotes contracts that may be anticipated or otherwise endogenous, and the residual component, $\varepsilon^{g}_{i,t}$, refers to the amount of *unanticipated contracts*, which plausibly behave as unanticipated demand shocks, affecting recipient establishments starting at or after the award date.

Econometric Specification. Using our institutional knowledge of the procurement process, we identify a set of contracts that are plausibly unanticipated at the establishment-level, and aggregate them to construct the unanticipated share of contracts $\varepsilon^{g}_{i,t}$. This allows us to use panel local projections to estimate the effect of $1 of unanticipated contracts on employment (Jordà, 2005). In particular, we estimate via OLS the following baseline equation, expressed in long-differences (Piger and Stockwell, 2025):

$$E_{i,t+h} - E_{i,t-1} = \underbrace{\alpha^{h}_{i} + \alpha^{h}_{s,t} + \alpha^{h}_{\ell,t}}_{\text{Fixed Effects}} + \beta^{h} \cdot \varepsilon^{g}_{i,t} + \gamma^{h}_{0} \cdot \widetilde{G}_{i,t} \tag{7}$$

$$\ldots + \underbrace{\sum_{j=1}^{4} \left\{ \rho^{h}_{j} \cdot \varepsilon^{g}_{i,t-j} + \gamma^{h}_{j} \cdot \widetilde{G}_{i,t-j} + \phi^{h}_{j} \cdot (E_{i,t-j} - E_{i,t-1-j}) \right\}}_{\text{Lags}} + v_{i,t+h} \quad h = 0, 1, \ldots, H$$

where $E_{i,t+h}$ denotes the h-period ahead number of employees; $\varepsilon^{g}_{i,t}$ denotes the dollar value of unanticipated contracts awarded to establishment i in quarter t, while $\widetilde{G}_{i,t}$ indicates the dollar value of potentially anticipated contracts, both are expressed in units of $1 million 2008 dollars. α^{h}_{i} represents an establishment fixed effect, $\alpha^{h}_{s,t}$ is a sector-time fixed effect intended to absorb any sectoral business-cycle effects. Lastly, $\alpha^{h}_{\ell,t}$ represents a state-time fixed effect, capturing regional business-cycle effects within a state.

Before presenting results, we describe the identification strategy and the data. Identification threats, robustness exercises, and external validity are discussed at the end of the section.

Identification. Most government contracts cannot be treated as quasi-random shocks at the establishment level. As discussed in Section II, the majority of procurement spending takes place in the context of long-term agreements (e.g., IDVs) whose award 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 (see Figure 2 in Section II).

Therefore, we identify our demand shock ($\varepsilon^{g}_{i,t}$) using a subset of contracts which we argue are plausibly unanticipated at the establishment-level by nature of the process used to award them. In light of the considerations in Section II, we define a contract as unanticipated ($\varepsilon^{g}_{i,t}$) if it meets four conditions:

1. contracts are awarded through competition[25]
2. with at least two bidders
3. newly awarded (not modifications of existing agreements)
4. standalone contracts, i.e., definitive contracts (not part of an ongoing series of purchases)

Conditions 1 and 2 are similar to those imposed by Hebous and Zimmermann (2020) while condition 3 was introduced in Budrys (2022), both in the context of publicly traded Compustat companies. Condition 4 imposes a novel additional restriction. Descriptive statistics on unanticipated contracts are provided in Appendix F.2.

Data. We study the employment response of recipient establishments using restricted microdata from the Longitudinal Database of Business Establishments (LDBE), compiled by the BLS from the Quarterly Census of Employment and Wages (QCEW). Our analysis covers 42 states over the period 2006–2019.

In the LDBE, we identify establishments using EIN-state pairs. This is a natural simplification because, in the 42 states we study, 96% of firms operate at most one establishment within a state. The restriction still allows large multi-establishment contractors to enter the sample, since subsidiaries and establishments in different states typically file taxes separately and would therefore appear under distinct EINs.

We combine these data with the universe of federal procurement contracts from the Federal Procurement Data System (FPDS). Contracts are aggregated to the recipient establishment-quarter level using Dun & Bradstreet's DUNS identifiers. Because DUNS codes are location-specific, they provide a close proxy for establishments.

Matched Sample. We merge FPDS contracts to LDBE establishments using the matching procedure described in Appendix F.1. We successfully match 13,662 establishments that receive at least one unanticipated contract ($\varepsilon^{g}_{i,t}$) between 2006 and 2019. We then restrict the sample to establishments with sufficiently complete histories for local-projection estimation. Specifically, we exclude establishments with gaps in their time series, establishments whose only contract shock occurs in the first four observed quarters or in the last two observed years, and establishments with average employment below one worker. This yields a sample of 5,317 establishments.

To limit the influence of the right tail of the establishment-size distribution, we further exclude establishments with more than 150 employees, leaving 5,142 establishments in the final analysis sample.

[25] The competition classes are "Full and open competition" and "Full and open competition after exclusion of sources". Contracting officers have indicated that even a single-offer scenario, if open to full competition, is treated as competitive because it may pressure the bidder to refine its proposal in anticipation of additional bids. However, they have also communicated that the number of bids is a good indicator of competitiveness, bolstering our confidence in the unanticipated nature of competed definitive contract awards with multiple bidders.

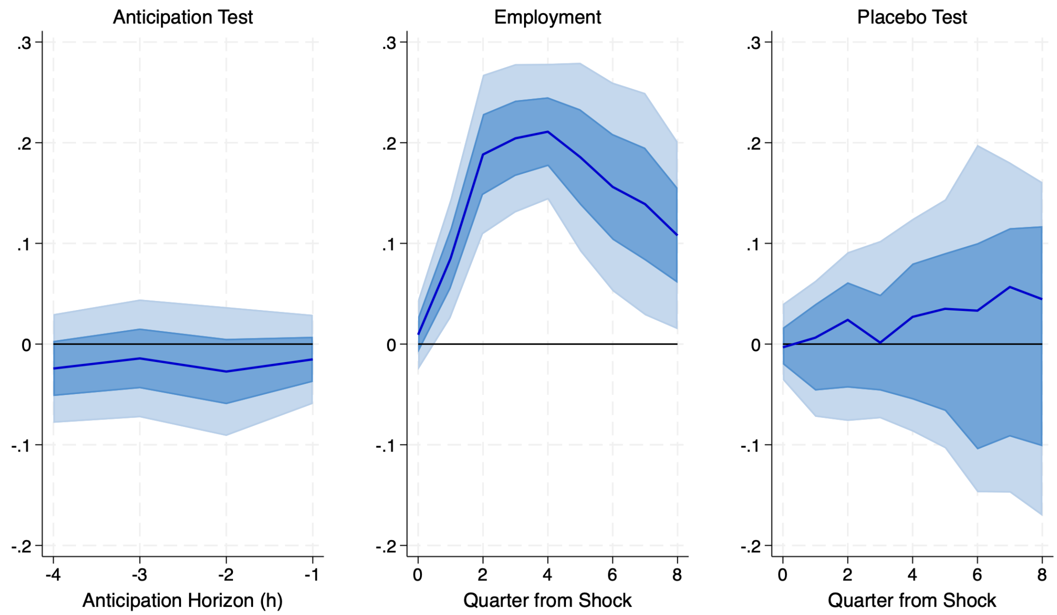

FIGURE 6 — EFFECT OF UNANTICIPATED CONTRACTS ON ESTABLISHMENTS' GROWTH

Notes: Sample: 2006:1 to 2019:4 ($T = 56$); number of establishments $N = 5,142$. Robust standard errors are clustered at the state level. Small bands represent 68% confidence intervals, and large bands represent 95% confidence intervals.

This restriction does not imply that all observations come from small firms: some establishments in the sample belong to medium-sized or large firms, since firm size in the BDS is defined by total employment across all establishments.

Our final sample is composed of 5,142 establishments between 2006q1 and 2019q4.

Results. We estimate Equation (7) by OLS for horizons $h = 0, \ldots, 8$, where the coefficients β^h can be interpreted as impulse responses of establishment-level employment to an additional \$1 million of unanticipated contracts.

The middle panel of Figure 6 shows a positive and statistically significant employment response following contract receipt. The effect builds gradually and peaks at horizon 4. This dynamic pattern is consistent with the regional evidence, where employment also responds only gradually to defense spending shocks.

The response is also persistent. Employment remains elevated even eight quarters after the shock, despite the fact that 75% of unanticipated contracts have a duration shorter than four quarters (Appendix F.2). Thus, the employment effects of procurement contracts appear to outlast the underlying contract itself. This persistence is consistent with previous evidence for Brazil, South Korea, and Portugal in Ferraz, Finan, and Szerman (2021), Lee (2024), and Gabriel (2024), respectively.

At the same time, the direct effect at the establishment-level is quantitatively modest. The peak response at horizon 4 is 0.21, implying that \$1 million in contracts raises employment by 0.21 workers after one year, on average. Following Chodorow-Reich (2019), we summarize the overall employment effect by computing job-years, obtained by cumulating the impulse response function and dividing by four because the data are quarterly. Cumulating the quarterly impulse responses in Figure 6 through horizon 4 yields 0.698 job-quarters per $1 million of contracts. Dividing by four implies 0.174 job-years per $1 million within one year of the award.

Breaking Down the Contractor Response into Direct and Indirect Effects.
The establishment-level estimates calculated in this section are informative not only about the direct effect of contract award at recipient establishments, but also for interpreting the regional contractor response.

Section V showed that, within one year of a regional defense spending shock, essentially all employment gains originate within the pool of contractors; specifically, Table 6 shows that contractors account for 157.6% of regional job-years within one year of the shock (since the impact on non-contractors is negative). Since the average regional response in that sample is 0.601 job-years per $1 million (in 2008 dollars), the total contractor response within one year equals $157.6\% \times 0.601 = 0.947$ job-years per $1 million. By construction, however, this contractor response combines two channels: the direct response of establishments that receive contracts (channel 1 in Figure 5) and the indirect response of other establishments belonging to contractor firms (channel 2).

We use the establishment-level estimates to approximate the contribution of channel 1. In particular, using a back-of-the-envelope calculation, we can combine the regional and establishment-level responses to provide an approximate estimate of the share of contractor employment gains attributable to recipient establishments, within one year from the shock:

$$\frac{\text{Recipient contractor establishments response}}{\text{Total contractor response}} = \frac{\underbrace{0.174\,\frac{\text{job-years}}{\$1\text{M}}}_{\text{Figure 6}}}{\underbrace{157.6\%}_{\text{Table 6}} \times \underbrace{0.601\,\frac{\text{job-years}}{\$1\text{M}}}_{\text{Table 7}}} = 18.4\%.$$

The above calculation provides an approximate decomposition of the contractor response, provided that the establishment-level response to unanticipated contracts is informative about the direct effect of regional defense-spending shocks on recipient contractors' employment.

Fact 4. *Procurement contracts generate persistent but small employment gains at recipient establishments, accounting for only about 18% of the total contractor employment response to defense shocks.*

This finding implies that most employment gains within the contractor pool occur outside the establishments that directly receive contracts. In other words, the bulk of the contractor response reflects employment adjustments at other establishments belonging to contractor firms. This pattern points to substantial propagation within the contractor network—through mechanisms such as internal firm reallocation, input-output linkages, production complementarities across establishments, or subcontracting relationships—rather than employment gains arising solely at awardee establishments.

Interpretation of the Breakdown. A final issue is whether the employment effects identified using unanticipated contracts are informative about the broader contractor response estimated in the regional analysis. We view this as plausible for three reasons.

First, establishments that receive at least one unanticipated contract account for a large share (about 80%) of aggregate federal procurement value over our sample, so the analysis does not focus on a negligible subset of the contractor universe.[26]

Second, while the sample includes contracts from all federal agencies, the Department of Defense accounts for roughly two-thirds of identified unanticipated awards, making the sample highly relevant for interpreting regional defense-spending shocks.

Third, using data from NETS, Figure 7 shows that, although the firm-size distributions for unanticipated versus general defense contracts certainly differ, awards are spread across the firm-size distribution and are not disproportionately tilted toward small firms.

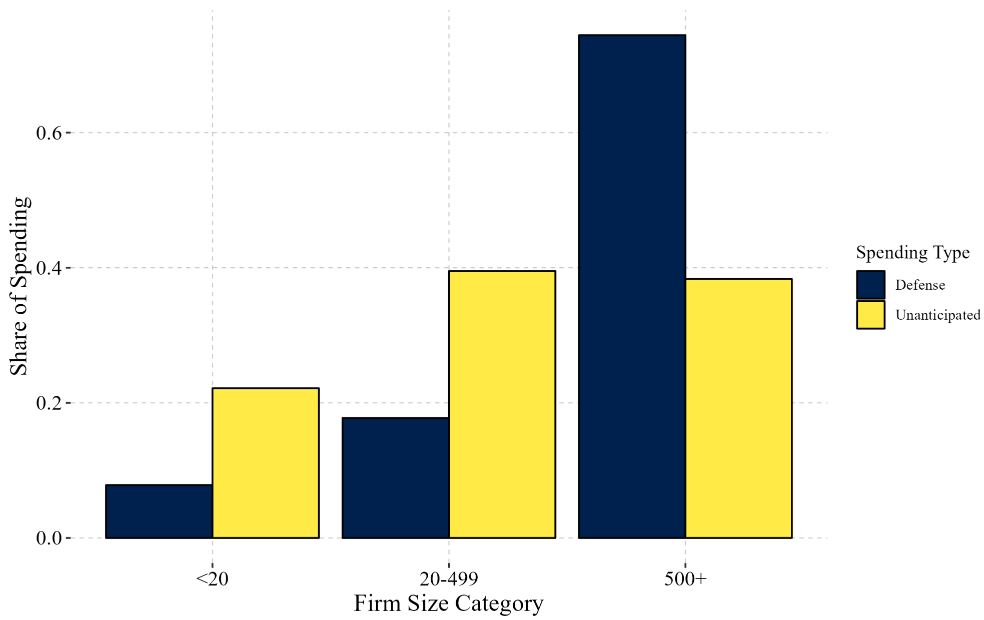

FIGURE 7 — SIZE DISTRIBUTION BY SPENDING TYPE

Notes: Firm sizes are assigned by matching FPDS to NETS using DUNS codes and cumulating employment of establishments to top-level headquarters of different establishments, as described in Section IV. Size bins correspond to Business Dynamics Statistics categories. Firm sizes are measured in the quarter of a contract award. Data span 2000-2019.

For these reasons, we interpret the establishment-level estimates as informative about the direct hiring response at recipient establishments and therefore as a useful approximation to channel (1) in our decomposition. At the same time, this decomposition should be understood as approximate, since the regional analysis captures the full bundle of contracts induced by defense-spending shocks rather than the set of unanticipated contracts used for identification here.

[26]Specifically, for each fiscal year in our sample, we compute the share of total federal procurement contract value awarded to establishments that receive at least one unanticipated contract at some point during the sample period. Averaging this share across fiscal years yields approximately 80%.

Identification Threats and Robustness

This subsection discusses the main threats to identification and summarizes the evidence supporting our interpretation of $\varepsilon^{g}_{i,t}$ as a plausibly exogenous establishment-level demand shock. It concludes with robustness exercises.

Anticipation. A first concern is that establishments may adjust employment before the recorded award date. This would arise if contracts were part of pre-existing procurement relationships, or if firms could confidently anticipate winning a solicitation well in advance. Our definition of "unanticipated" contracts is designed to limit this concern. By construction, unanticipated contracts are newly awarded, standalone, competed, and receive at least two bids. These restrictions exclude modifications of ongoing agreements and contracts embedded in longer procurement sequences, precisely the settings in which anticipation is most likely. In addition, the timing of procurement opportunities suggests limited scope for meaningful anticipation at the quarterly frequency used here. In the FPDS-linked contracts-solicitation data, the median time between the public announcement of a contract opportunity and the award date is only 20 days (Appendix A). Prior evidence for publicly traded contractors also points in the same direction: Hebous and Zimmermann (2020) and Budrys (2022) show that future competed awards do not predict current stock-price movements of eventual awardees. We further assess anticipation directly in the data. Specifically, we re-estimate the impact specification replacing the contemporaneous shock with leads of $\varepsilon^{g}_{i,t}$ at horizons one through four quarters. The left panel of Figure 6 shows no economically meaningful or statistically significant effect of future awards on current employment growth. This absence of pre-trends supports the view that the contracts we classify as unanticipated do not induce adjustment before the award date.

Selection. A second concern is selection bias. In the context of federal purchases, Nekarda and Ramey (2011) highlight that industry technological progress can endogenously drive medium-term changes in industry-level government purchases (Perotti (2007)), i.e., there may be reverse causality in the relationship between government contracts and firm innovation and growth. Indeed, government purchases driven by technological progress not only occur frequently, but they are specifically regulated by Federal Acquisition Regulation 6.302-1-a: sole-source acquisition procedures.

Our empirical strategy addresses this concern in two ways. First, our contract definition is designed to exclude precisely the types of awards most likely to reflect endogenous firm-specific innovation or targeted non-competitive selection. In particular, conditions (1) and (2) ensure that unanticipated contracts are unlikely to be awarded because an establishment introduced a new, innovative product; such acquisitions fall into the non-competed category and are therefore excluded.

Second, by including establishment fixed effects, we purge time-invariant differences in productivity or efficiency across contractors, so our estimates are not mechanically driven by establishments that systematically win more contracts and exhibit higher employment growth.

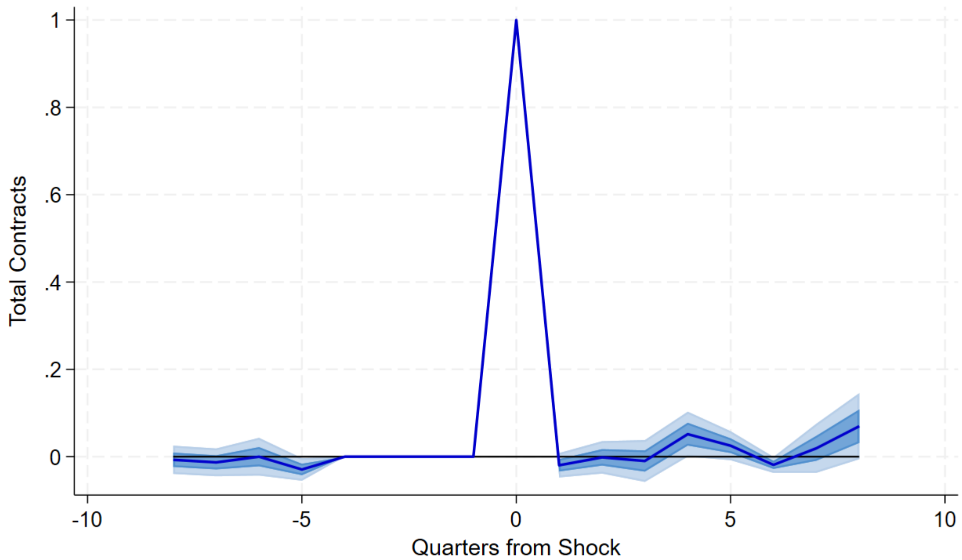

FIGURE 8 — UNANTICIPATED CONTRACTS BEHAVE AS ONE-TIME (DEMAND) SHOCKS

One-time Shock. A related concern is that our identified shocks may not behave like one-time demand innovations. To address this, we provide event-study evidence showing that total contracts jump one-for-one on impact with $\varepsilon^{g}_{i,t}$, with no detectable anticipatory dynamics and little persistence beyond what is absorbed by lagged controls. In particular, we estimate the following equation by OLS:

$$G_{i,t+h} = \lambda^h_i + \lambda^h_{s,t} + \lambda^h_{\ell,t} + \psi^h \cdot \varepsilon^g_{i,t} + \ldots$$
$$\ldots + \sum_{j=1}^{4} \rho^h_j \cdot \varepsilon^g_{i,t-j} + \sum_{j=0}^{4} \gamma^h_j \cdot \widetilde{G}_{i,t-j} + u_{i,t+h}, \qquad h = -8, \ldots, -1, 0, 1, \ldots, 8,$$

We plot the OLS estimates of ψ^h in Figure 8 for both positive and negative horizons.

On impact, the estimated coefficient jumps to 1, indicating that all contemporaneous variation in contracts originates from the unanticipated component, because we control for contemporaneous values of potentially anticipated contracts. Contracts also show little persistence: serial correlation is absorbed by lags of contracts. Before the shock, leads of unanticipated contracts do not predict current contract activity, ruling out anticipatory dynamics. After the shock, unanticipated awards do not trigger additional contracts in subsequent quarters, confirming the transitory nature of the shock.

Robustness

Placebo. We also implement a placebo exercise. We randomly shuffle the timing of contract shocks within establishments and re-estimate the baseline local projections using these synthetic shocks. The right panel of Figure 6 shows no significant employment response. This makes it unlikely that the baseline results are driven by spurious serial correlation or generic differences in establishment growth patterns.

Robustness Across the Size Distribution. Appendix F.3 shows that the establishment-level responses are robust across quartiles of the establishment-size distribution.

Controlling for Lags of Average Wage. Appendix F.4 shows that the establishment-level responses of employment are robust to controlling for lags of the average wage-per-worker, which are meant to absorb any potential time-varying productivity shocks that may predict contract awards.

Response of Wages. Appendix F.5 shows that total wages rise after contract receipt in line with the employment expansion. By contrast, average wages per worker do not display any significant dynamics. This pattern suggests that the response operates primarily through the extensive margin rather than through higher pay within recipient establishments. Consistent with this interpretation, the wage-per-worker effects observed in the regional analysis are more naturally understood as reflecting compositional changes in the broader contractor network than wage increases at directly treated establishments.

Robustness to Staggered Treatment and Alternative Specifications. Finally, Appendix F.6 results reinforce the baseline findings along several dimensions. First, we show that the decomposition is not sensitive to the quarterly frequency of the establishment analysis: estimating analogous effects at annual frequency using NETS data yields similar magnitudes. Second, we address concerns about already-treated controls by implementing a local-projections difference-in-differences design that uses only not-yet-treated and never-treated contractors as controls, following Dube et al. (2025). The resulting estimates exhibit flat pre-trends and confirm positive post-award employment responses. Third, we consider alternative treatment definitions, including specifications that normalize contracts by establishment size and log-log specifications that reduce the influence of extreme contract values; the qualitative conclusions are unchanged.

VII. Conclusion

This paper studies how defense procurement affects employment growth across U.S. local labor markets and firms. Using newly matched federal procurement data and restricted administrative employment records, we show that the employment effects of defense spending are positive but costly, highly concentrated, and slow to diffuse across the local labor market. The implied cost per job-year is roughly $290,000 in 2008 dollars, well above the benchmark range commonly cited in the recent fiscal-policy literature and in line with other recent estimates for modern defense procurement. Employment gains accrue primarily to large firms, which capture the bulk of shock-induced contracts. Non-contractors are initially crowded out, with broader spillovers emerging only in the medium run. And most contractor-side employment growth occurs outside the establishments that directly receive contracts, pointing to substantial propagation through indirect channels.

These findings carry two implications. First, defense procurement is better suited to supporting gradual employment growth than to generating rapid, broad-based stimulus. Second, evaluating procurement solely through the lens of recipient establishments understates its employment content, since indirect channels within contractor networks and delayed spillovers to non-contractors are first-order features of fiscal transmission through procurement.

Our analysis quantifies the aggregate importance of indirect channels but treats the contractor network as a single object. An important direction for future work is to open this box by linking prime contractors to subcontractors and production-network partners, tracing how procurement propagates through firm-to-firm relationships. Such an analysis could clarify whether the high cost per job reflects the sectoral composition of modern defense production or the structure of contractors' networks.

Appendix

A Solicitations

Bids and proposals for competitively awarded contracts are solicited on a government website, Federal Business Opportunities (FedBizOpps or FBO), now migrated to SAM.gov. Contract solicitation allows any potential vendor to view the contract opportunity on the website and participate in the auction or negotiation. Usually, agencies post a "pre-solicitation" notice, informing vendors about the possibility that a contract opportunity may arise. Contracts are then officially solicited on the same website. In this period, contractors can submit offers in the form of (i) bids (i.e., either one or two step sealed bidding) or, when the nature of the product is more complex, written proposals (i.e., contract by negotiations). Once the offer period expires, awardees are competitively selected. All pre-award notices are gathered daily on SAM.gov. Following Carril, Gonzalez-Lira, and Walker (2026) approach, we download all daily solicitations posted on SAM.gov from fiscal year 2006 to fiscal year 2020, and then use information from the (i) solicitation number, (ii) awarding sub-agency name and (iii) fiscal year to identify unique contracts solicitations and reconstruct the entire pre-award sorted history: from the oldest pre-award notice to the award notice.[27] Figure A1 summarizes the competitive procurement timeline process.

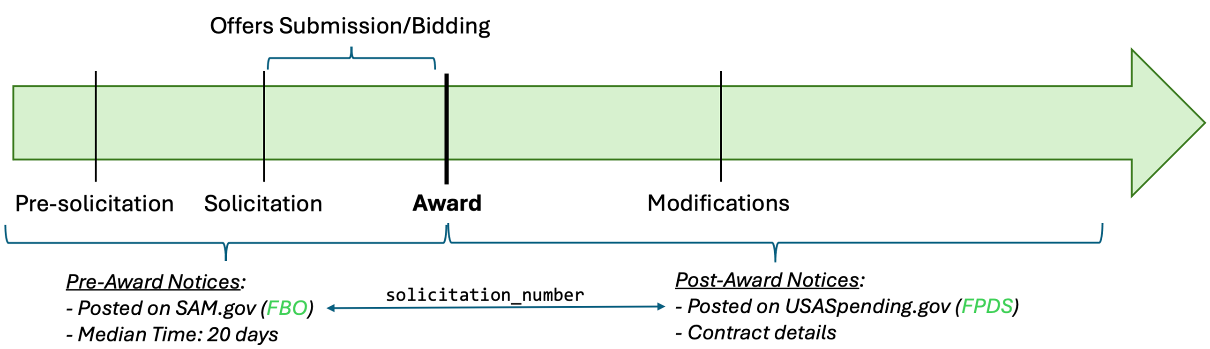

FIGURE A1 — TIMELINE OF COMPETED CONTRACTS

Notes: Once the contract is awarded, all detailed contract information is recorded in FPDS by the responsible federal contracting officer. Several contract modifications, referred to as "*post-award* actions, may follow the award. Common examples of contract modifications are invoking an option to purchase additional quantities, agreeing to extra costs for extra work, appropriation of extra funds, and contract termination.

We keep all award histories from fiscal year 2006 to fiscal year 2019 to be consistent with the sample used in our establishment-level analysis, and analyze the number of days from the first pre-solicitation to the award notice, dropping solicitations that either (i) lack an award notice or (ii) consist only of a single notice. Figure A2 shows the box plot of the (unweighted) number of days from the oldest pre-award notice to the award notice.

[27] We thank Andres Gonzalez-Lira for directing us to the General Services Administration Technical Documentation for the FedBizOpps (FBO) website, whose content has since migrated to Contract Opportunities (SAM).

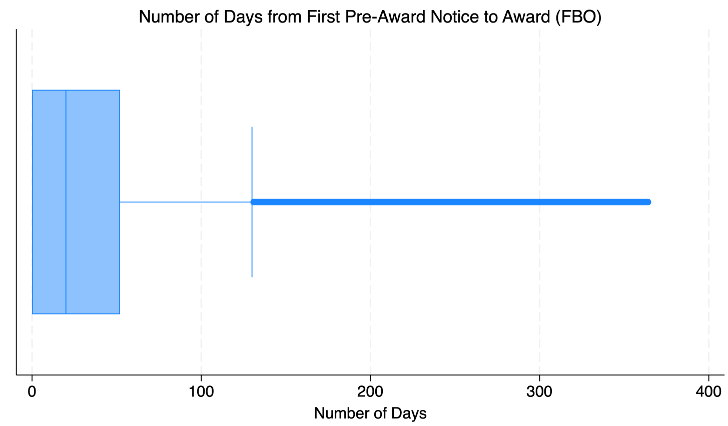

FIGURE A2 — BOX-PLOT OF NUMBER OF DAYS FROM OLDEST PRE-AWARD NOTICE TO AWARD

Notes: Distribution is not weighted by the value of a contract. Data source is the universe of federal procurement solicitations from FBO (Federal-Bizz-Opportunities.gov), now migrated to SAM.gov.

We find that the median time taken from the first '*pre-award*" notice (e.g., pre-solicitation) and the award notice for any competed federal contracts is 20 days, while for 75% of contracts this interval of time is 52 days or fewer. Thus for most contracts it would be difficult for a firm to react to a contract solicitation even a single quarter before the corresponding award date.

In light of the short time period between pre-solicitations and award dates, we use the award date available from FPDS to identify the timing of the award. We address potential anticipation effects owing to the pre-award solicitation period by carrying out an anticipation test in the main body of the paper.

B Extra: Regional Employment Multipliers

B.1. Robustness of Fact 1: Different Samples

TABLE B1 — REGIONAL EMPLOYMENT MULTIPLIERS FROM BEA DATA - LARGEST SAMPLE

BEA Total Employment - Largest Sample: 2001—2019; 380 MSAs					
Horizon	*Coefficient*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.020	0.004	13.643	0.655	$1,527,213
	(0.007)			(0.227)	($529,936)
1 year	0.095	0.000	93.200	3.136	$318,887
	(0.025)			(0.816)	($82,970)
2 years	0.081	0.002	52.473	2.665	$375,188
	(0.026)			(0.851)	($119,836)
3 years	0.114	0.018	10.553	3.750	$266,650
	(0.048)			(1.577)	($112,133)

Notes: Robustness of Table 2 (baseline regional employment multipliers). Largest available sample BEA.

TABLE B2 — REGIONAL EMPLOYMENT MULTIPLIERS FROM BEA DATA - SMALLEST SAMPLE

BEA Employment - Smallest Sample: 2006—2019; 254 MSAs					
Horizon	*Coefficient*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.024	0.229	10.019	0.762	$1,311,855
	(0.020)			(0.632)	($1,088,125)
1 year	0.085	0.037	29.845	2.755	$362,987
	(0.041)			(1.318)	($173,608)
2 years	0.083	0.091	7.532	2.670	$374,548
	(0.049)			(1.573)	($220,603)
3 years	0.092	0.111	6.561	2.989	$334,577
	(0.058)			(1.867)	($208,961)

Notes: Robustness of Table 2 (baseline regional employment multipliers). Sample: Harmonized QCEW+BDS+LAUS+LDBE Dataset.

B.2. Robustness of Fact 1: Spreading Contracts Over Their Duration

For robustness, we also recompute the employment multipliers by spreading the value of contracts over their duration, in line with the approach adopted in several papers in the literature (see, among others, Demyanyk, Loutskina, and Murphy, 2019; Auerbach, Gorodnichenko, and Murphy, 2020; Auerbach, Gorodnichenko, and Murphy, 2024; Auerbach, Gorodnichenko, and Murphy, 2025). Table B3 reports the results.

The estimated cost-per-job on impact is now much closer to the values estimated at longer horizons and substantially lower than the values obtained when using the full contract value at the time of award. This pattern is consistent with the presence of long delays between contract awards and the actual realization of spending.

TABLE B3 — REGIONAL EMP. MULTIPLIERS: SPREADING CONTRACTS OVER DURATION

Response of Total Employment from (Public) BEA Data					
Horizon	*Coefficient (β_h)*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.091	0.003	30.412	3.013	$331,877
	(0.030)			(0.997)	($109,845)
1 year	0.128	0.001	28.242	4.225	$236,697
	(0.040)			(1.314)	($73,617)
2 years	0.134	0.002	27.153	4.421	$226,182
	(0.043)			(1.417)	($72,502)
3 years	0.126	0.002	33.755	4.169	$239,841
	(0.041)			(1.349)	($77,575)

Notes: Sample: 2001-2019; 358 MSAs (QCEW+BDS+LAUS Harmonized Sample). All rest identical to Table 2.

B.3. Robustness of Fact 1: Employment Multiplier Dropping High-Exposure MSAs

One concern raised by Figure 4 is that our results might be driven by a small number of observations which have very high values of the instrument as a result of their unusually high exposure to defense spending. To assess this possibility, we re-estimate our baseline estimate from Table 2 but dropping the MSAs with the top 5% of defense spending shares. As shown in Table B4, this exercise produces very similar estimates to the baseline.

TABLE B4 — REGIONAL EMPLOYMENT MULTIPLIERS - SHIFT-SHARE INST. (DROP TOP 5% SHARES)

Response of Total Employment from (Public) BEA Data					
Horizon	*Coefficient* (β_h)	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.042	0.030	17.86	1.41	$708,343
	(0.019)			(0.65)	($325,895)
1 year	0.103	0.017	84.87	3.44	$290,330
	(0.043)			(1.44)	($121,107)
2 years	0.096	0.035	37.67	3.20	$312,422
	(0.045)			(1.51)	($147,875)
3 years	0.101	0.049	29.08	3.37	$296,411
	(0.051)			(1.71)	($150,377)

Notes: Sample: 2001-2019; 340 MSAs (QCEW+BDS+LAUS Harmonized Sample Less Top 5% MSAs by Long-Run Exposure). All rest identical to Table 2.

B.4. Employment Multiplier After Dropping the Right Tail of the Procurement Distribution

In this section, we examine how the results change when we trim the right tail of the procurement distribution, to assess whether the estimates are driven by MSAs hit by procurement shocks that are large relative to local economic activity. For example, the Oshkosh-Neenah MSA in Wisconsin has the largest average value, with defense procurement accounting for more than 50% of local wages and salaries, likely reflecting the presence of one of the largest defense contractors, Oshkosh Corp.

We therefore exclude the top 5% of MSAs in the distribution of average within-MSA government procurement relative to local wages and salaries in the baseline estimation sample. The resulting sample contains 340 MSAs, and the largest procurement-to-wages-and-salaries ratio falls to 16%. Results are reported in Table B5.

TABLE B5 — REGIONAL EMPLOYMENT MULTIPLIERS – SHIFT-SHARE INSTRUMENT (DROPPING THE RIGHT TAIL OF PROCUREMENT)

Response of Total Employment from (Public) BEA Data					
	IV: Shift-Share Instrument				
Horizon	*Coefficient* (β_h)	p	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
Impact	0.065	0.067	43.964	2.136	$468,072
	(0.035)			(1.162)	($254,633)
1 year	0.171	0.001	132.721	5.668	$176,428
	(0.051)			(1.691)	($52,642)
2 years	0.172	0.002	133.381	5.697	$175,532
	(0.054)			(1.784)	($54,956)
3 years	0.198	0.002	98.690	6.543	$152,827
	(0.063)			(2.084)	($48,682)

Notes: Sample: 2001–2019; 340 MSAs (balanced panel; QCEW+BDS+LAUS harmonized sample). GDP price deflator from BEA, base year 2008. Robust standard errors in parentheses, clustered at the MSA level. The Montiel Olea and Pflueger (2013) effective F-statistic is calculated using `weakivtest`. Job-years are reported per $1 million of spending. Standard errors for cost-per-job are obtained using the Δ-method.

Employment multipliers remain precisely estimated in this smaller sample, and the point estimates are generally larger than those reported in the baseline full-sample table. One possible explanation is that MSAs with especially high exposure to defense spending are more subject to crowding-out effects, with workers reallocating from non-contractors to contractors in response to a defense shock. Nevertheless, the estimated cost-per-job remains well above the upper bound of $125,000 reported by Chodorow-Reich (2019), suggesting that the high cost-per-job in FACT 1 does not depend on the inclusion of MSAs with unusually high reliance on defense procurement.

B.5. Robustness of Fact 1: First Stage and Instrument Validity

In this section we carry out several robustness checks to corroborate the validity of our baseline estimates of employment multipliers reported in Table 2.

First Stage. In this section, we report the results of the first stage of our baseline estimation of employment multipliers (Equation (1)). The first-stage regression is specified as follows:

$$\underbrace{\frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}}}_{\text{RHS of Equation (1)}} = \phi_h \cdot \underbrace{\frac{\exp_\ell \cdot (G_{t+h} - G_{t-1})}{Y_{\ell,t-1}}}_{Z_{\ell,t+h}} + \underbrace{\lambda_{\ell,h} + \alpha_{t,h}}_{\text{FEs}} + e_{\ell,t+h}.$$

Values of the OLS-estimated first-stage coefficients (ϕ_h) are reported in Table B6, along with robust and clustered standard errors. As expected, the instrument is strong, with effective F-statistics exceeding 23, consistent with the baseline results reported in Table 2. Moreover, also as expected, the coefficient on impact is statistically insignificantly different than 1.

TABLE B6 — FIRST STAGE COEFFICIENTS OF EQUATION (1)

Horizon	*Coefficient* (ϕ_h)	p
impact	0.981	0.000
	(0.176)	
1 year	1.506	0.000
	(0.200)	
2 years	1.638	0.000
	(0.301)	
3 years	1.887	0.000
	(0.396)	

Notes: Sample: 2001-2019 - 358 MSAs (QCEW-BDS-LAUS Harmonized Dataset). All rest is identical to Table 2.

Lastly, Figure B1 reports a bin-scatter plot of the residualized first-stage regressions by horizon h. By the Frisch–Waugh–Lovell theorem, the slope of the fitted line coincides with the estimates of ϕ_h reported in Table B6. Given our sample size, each point in the graph represents an average of roughly 21 observations. The strong fit is consistent with the low p-values in Table B6 and the relatively high effective F-statistics reported in the baseline results (Table 2).

Instrument Validity. We check whether our baseline instrument is explained by lagged outcomes at any horizon. This exercise can be viewed as a balance test on pre-treatment trends, which provides evidence on the plausibility of the parallel trends assumption. If regions with higher exposure (i.e., high instrument values) were systematically growing faster than average in the past, then identification would be undermined, as the instrument would also capture endogenous placement of contracts rather than exogenous shocks. In that case, it would also be necessary to control for lagged employment growth in the baseline analysis to mitigate this source of endogeneity.

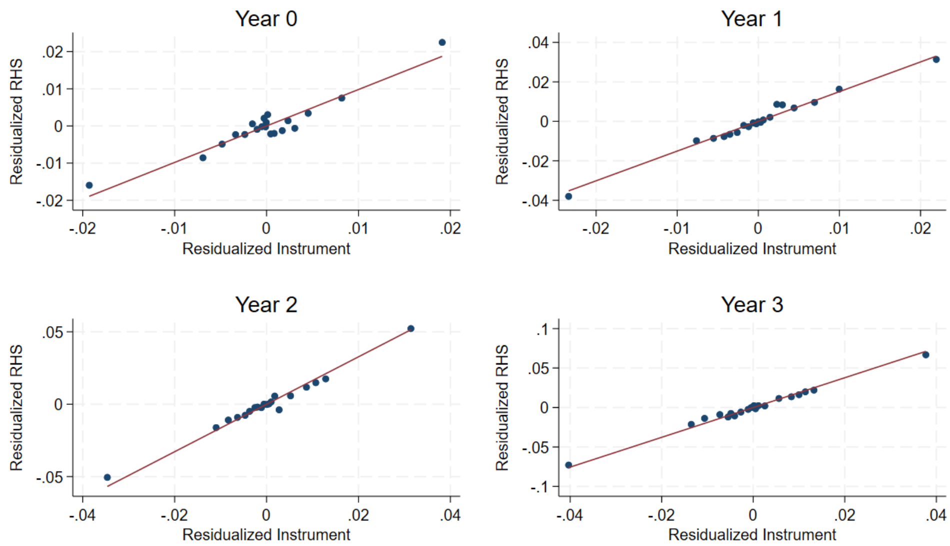

FIGURE B1 — BIN-SCATTER PLOT OF (RESIDUALIZED) FIRST-STAGE REGRESSIONS

In practice, we estimate via OLS using our baseline harmonized BEA-BDS-LAUS sample the following equation:

$$\underbrace{\frac{\exp_{\ell} \cdot (G_{t+h} - G_{t-1})}{Y_{\ell,t-1}}}_{z_{\ell,t+h}} = \rho_h \cdot \underbrace{\frac{E_{\ell,t-1} - E_{\ell,t-2}}{E_{\ell,t-2}}}_{\text{Lagged LHS}} + \underbrace{\omega_{\ell,h} + \pi_{t,h}}_{\text{FEs}} + v_{\ell,t+h}$$

TABLE B7 — IS OUR INSTRUMENT EXPLAINED BY LAGGED EMPLOYMENT GROWTH?

LHS: Instrument - RHS: Lagged employment growth		
Horizon	*Coefficient* (ρ_h)	*pvalue*
impact	0.010	0.182
	(0.008)	
1 year	0.010	0.431
	(0.012)	
2 years	0.003	0.879
	(0.018)	
3 years	-0.019	0.397
	(0.023)	

Notes: Sample: 2001-2019 - 358 MSAs (QCEW-BDS-LAUS Harmonized Dataset). All rest is identical to Table 2.

Results are reported in Table B7. Notice that lagged employment growth—the horizon 0 variable in our baseline regression, Equation (1)—does not predict values of the instrument at any horizon h. This result supports our identification strategy, providing evidence of the absence of pre-trends in regions with high exposure to defense contracts.

B.6. Extra: Effect on LAUS Unemployment versus Labor Force Participation

We re-estimate Equation (1), substituting employment with either unemployment or labor-force levels from the BLS Local Area Unemployment Statistics (LAUS). The results are shown in Table B8.

TABLE B8 — BASELINE ESTIMATES WITH LAUS DATA

	Employment data from LAUS									
	Unemployment					Labor Force				
Horizon	Coefficient	p	Effective F	Job-Years ($1M)	Cost-per-Job ($)	Coefficient	p	Effective F	Job-Years ($1M)	Cost-per-Job ($)
impact	0.212	0.256	29.232	0.371	$2,696,413	0.029	0.143	29.232	0.834	$1,198,829
	(0.186)			(0.326)	($2,368,343)	(0.020)			(0.568)	($815,970)
1 year	-0.201	0.057	53.120	-0.352	$(2,840,907)	0.075	0.038	53.120	2.165	$461,937
	(0.105)			(0.184)	($1,487,100)	(0.036)			(1.037)	($221,317)
2 years	-0.184	0.040	27.603	-0.322	$(3,103,027)	0.068	0.056	27.603	1.963	$509,388
	(0.089)			(0.156)	($1,504,555)	(0.036)			(1.026)	($266,241)
3 years	-0.195	0.078	21.063	-0.342	$(2,925,322)	0.073	0.077	21.063	2.118	$472,238
	(0.111)			(0.194)	($1,656,749)	(0.041)			(1.192)	($265,880)

Notes. Sample: 2001-2019 - 358 MSAs (Harmonized QCEW+BDS+LAUS Sample). Cost-per-Job is constructed as before, replacing employment with unemployment or labor-force statistics. All other details match Table 2.

The left panel of the table shows that unemployment and labor force coefficients are statistically significant after the impact, but insignificant on impact, mirroring the dynamics of the baseline employment results reported in Table 2. Specifically, the left panel shows negative changes in the number of unemployed individuals, indicating that regional shocks to defense purchases reduce unemployment, while the right panel documents a significant and growing increase in the labor force.

To gauge magnitudes, we express the multipliers in job-years and cost-per-job, following the same procedure used for employment but replacing employment with unemployment or labor-force levels in the conversion factor. The corresponding figures appear in the *Job-Years* and *Cost-per-Job* columns of Table B8. Three years after the regional shock, it costs $472,000 (2008 dollars) to add one person to the labor force, while it is much more expensive to reduce the number of unemployed. The positive employment multiplier originates primarily from new entrants into the regional labor force. Notice that the overall magnitudes appear smaller than those reported using high-quality administrative data on employment, which may reflect difficulties in timely tracking of changes in unemployment statistics in the LAUS methodologies.

Robustness. Lastly, for robustness, we repeat the analysis using the full available sample. Results are reported below.

TABLE B9 — UNEMPLOYMENT AND LABOR FORCE FROM LAUS DATA - LARGEST SAMPLE

Employment from LAUS - Largest Sample: 2001-2019; 366 MSAs										
	Unemployment					*Labor Force*				
Horizon	*Coefficient*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*	*Coefficient*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.212	0.256	29.232	0.371	$2,696,413	0.029	0.143	29.232	0.834	$1,198,829
	(0.186)			(0.326)	($2,368,343)	(0.020)			(0.568)	($815,970)
1 year	-0.201	0.057	53.120	-0.352	$(2,840,907)	0.075	0.038	53.120	2.165	$461,937
	(0.105)			(0.184)	($1,487,100)	(0.036)			(1.037)	($221,317)
2 years	-0.184	0.040	27.603	-0.322	$(3,103,027)	0.068	0.056	27.603	1.963	$ 509,388
	(0.089)			(0.156)	($1,504,555)	(0.036)			(1.026)	($266,241)
3 years	-0.195	0.078	21.063	-0.342	$(2,925,322)	0.073	0.077	21.063	2.118	$ 472,238
	(0.111)			(0.194)	($1,656,749)	(0.041)			(1.192)	($265,880)

Notes: Robustness of Table B8 (effect on unemployment and labor force). Largest available sample from LAUS.

B.7. Extra: Effects on Number of Firms

In principle, a regional shock could increase the number of firms in a region if it represents a sufficiently large demand shock. Consequently, the positive employment multiplier might arise from the creation of new businesses (the *extensive margin*) rather than from the expansion of the workforce within existing businesses (the *intensive margin*).

To investigate this possibility, we use data from the Business Dynamics Statistics (BDS), which report the number of firms in each MSA—year pair. We re-estimate Equation (1), replacing employment with the firm count, to assess whether the extensive margin contributes to the positive employment multiplier. The results are presented in Table B10.

TABLE B10 — EXTENSIVE MARGIN: NUMBER OF FIRMS

	Number of Firms from BDS				
Horizon	*Coefficient*	*p*	*Effective F*	*Firms Year ($1M)*	*Cost per Firm ($)*
impact	-0.005	0.756	29.232	-0.006	$(170,298,832)
	(0.017)			(0.019)	($546,542,528)
1 year	0.023	0.382	53.120	0.025	$40,784,724
	(0.026)			(0.028)	($46,605,776)
2 years	0.024	0.363	27.603	0.026	$37,820,644
	(0.027)			(0.029)	($41,520,636)
3 years	0.029	0.318	21.063	0.032	$31,613,240
	(0.029)			(0.032)	($31,616,600)

Notes: Sample: 2001-2019 - 358 MSAs (Harmonized QCEW+BDS+LAUS Sample). Cost per firm-year is constructed as before, replacing employment with the firm count. All other details match Table 2. No cost-per-firm year is calculated for negative point estimates.

We find no meaningful effect of the regional shock on the number of firms and we conclude that additional defense spending does not generate significant job creation through the establishment of new firms: the extensive margin is negligible. Instead, the positive employment multiplier appears to originate from existing firms hiring additional workers from the labor force and the pool of unemployed.

B.8. Extra: Effect on Wage-per-Worker

We construct regional wage-per-worker by dividing total wages and salaries from the BEA by total BEA employment, $\omega_{\ell,t}$. We then use our baseline instrument to estimate the following equation:

$$\frac{\omega_{\ell,t+h} - \omega_{\ell,t-1}}{\omega_{\ell,t-1}} = \beta_h \cdot \frac{G_{\ell,t+h} - G_{\ell,t-1}}{Y_{\ell,t-1}} + \lambda_{t,h} + \alpha_{\ell,h} + u_{\ell,t+h},$$

here, $Y_{\ell,t}$ denotes total regional wages and salaries, as in our baseline equation (7).

To provide context for the magnitude of wages across regions, we compute the average value of real wage-per-worker (in 2008 dollars) for each MSA over the 2001–2019 period, which corresponds to our baseline sample. We then examine the cross-sectional distribution of wage-per-worker across MSAs. The lowest value is $21,389 (McAllen–Edinburg–Mission, TX), the highest is $72,097 (San Jose–Sunnyvale–Santa Clara, CA), while the median MSA has a wage-per-worker of $30,530 (Montgomery, AL).

Figure B2 plots the IV estimates of β_h. We find that an increase in regional defense procurement spending equal to 1% of regional wages raises wage-per-worker by about 0.14% after one year, on average.

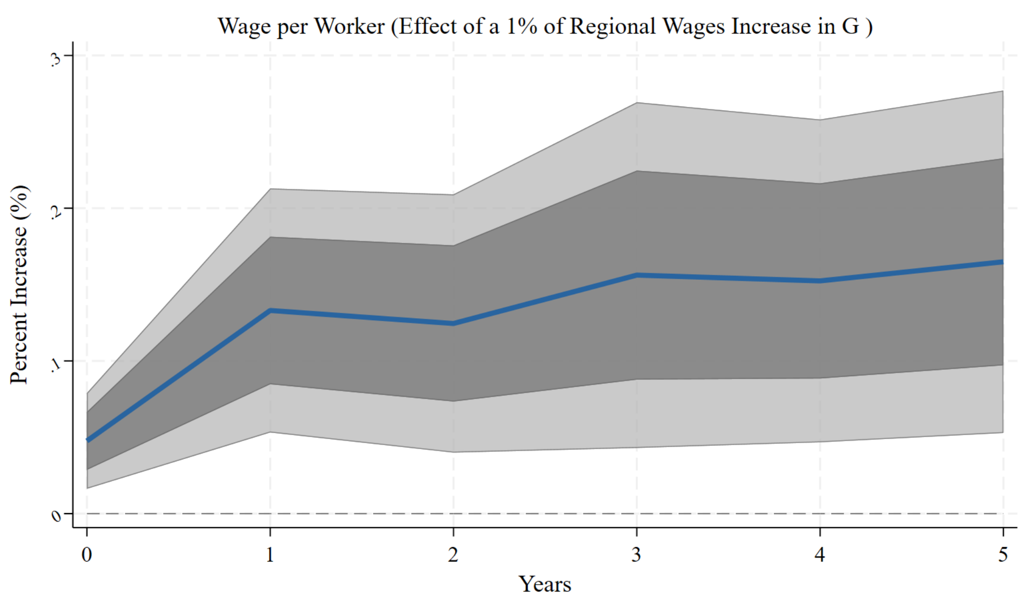

FIGURE B2 — EFFECT OF REGIONAL SHOCKS ON WAGE-PER-WORKER

Notes: Baseline BEA-BDS-LAUS harmonized sample. Bands represent 68% and 90% confidence intervals constructed using heteroskedasticity-robust standard errors.

C Employment Multiplier Using an Exact Shift-Share Instrument

In this section, we show that our baseline estimates of employment multipliers (Section III) remain unchanged when we replace the original instrument with an exact shift-share instrument.

Shift-Share Interpretation of Regional Government Spending Shocks. Following Goldsmith-Pinkham, Sorkin, and Swift (2020) and Borusyak, Hull, and Jaravel (2022), a generic shift-share (Bartik) instrument can be written as:

$$Z_{\ell,t} \; = \; \sum_{s=1}^{S} w_{\ell,s} \cdot g_{s,t},$$

where $w_{\ell,s}$ denotes the exposure (or "share") of region ℓ to sector s, and $g_{s,t}$ is a sector-level shock (the "shift"). In our setting, variation is driven by a single sector ($S = 1$), namely government spending. The instrument then simplifies to:

$$Z_{\ell,t} \; = \; w_{\ell} \cdot g_t,$$

where g_t is the aggregate government spending shock and w_ℓ is a fixed measure of region ℓ's baseline exposure to government spending.

This representation highlights that our design is a special case shift-share instrument: all regions face the same aggregate shock g_t, and cross-sectional variation arises solely from differences in w_ℓ. In contrast to the canonical Bartik setting, there are not multiple quasi-random shocks across sectors. The credibility of the design thus hinges entirely on the exogeneity of g_t and the predetermined nature of the exposure measure w_ℓ.

The exact shift-share instrument in this setting can be expressed in terms of government purchases over output:[28]

$$\tilde{Z}_{\ell,t+h} := \underbrace{\left(\frac{1}{19} \sum_{\tau=2001}^{2019} \frac{G_{\ell,\tau}}{Y_{\ell,\tau}} \right)}_{\text{Share}(w_\ell)} \cdot \underbrace{\frac{G_{t+h} - G_{t-1}}{Y_{t-1}}}_{\text{Shift}(g_t)} .$$

where Y_{t-1} denotes the aggregate value of the normalizing regional activity variable—-in our case, wages and salaries, defined as $Y_t := \sum_\ell Y_{\ell,t}$.

Identification. Identification requires the exogeneity of both the shifts and the shares. National shifts in defense expenditure are largely driven by geopolitical factors, as discussed in the paper. One potential concern is small-sample bias, since the time series dimension is relatively short ($T = 19$). For instance, the early 2000s were characterized by a military build-up following the 9/11 Terrorist Attacks, which coincided with the economic slowdown after the Dotcom crash. However, this coincidence does not hold systematically: defense spending cuts in 2013 due to sequestration occurred despite modest growth (nominal GDP grew by about 2%), while the military build-up of the first Trump presidency

[28] We thank Gabriel Chodorow-Reich for raising this point when discussing the paper at the 2025 NBER Conference: Fiscal Dynamics of State and Local Governments.

overlapped with strong economic growth. Thus, there is no clear countercyclical pattern between national defense spending shifts and national economic growth in recent years. Moreover, the time fixed effects in our baseline specification absorb any remaining aggregate variation, fully addressing these concerns.

The main source of variation in our data is cross-sectional, given the 358 MSAs in our harmonized baseline sample. Hence, the exogeneity of the shares is central for identification. In the canonical shift-share framework, shares are constructed as the time-average of the ratio of regional defense contracts to regional wages and salaries:

$$w_\ell := \frac{1}{19} \sum_{t=2001}^{2019} \frac{G_{\ell,t}}{Y_{\ell,t}}.$$

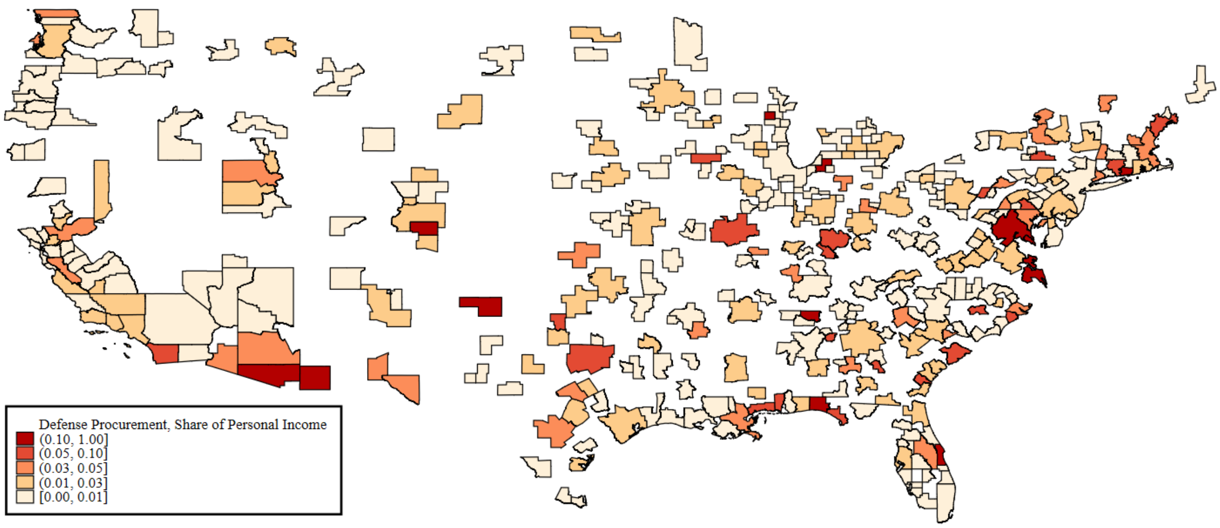

FIGURE C1 — MSA Geographic Distribution of Long-Run Shares of Defense Contracts w_ℓ

Notes: There are 380 MSAs. Long-run shares calculated using averages from 2001 to 2019. The figure omits Hawaii and Alaska.

By contrast, our baseline instrument redistributes national changes in defense spending according to the time-average fraction of defense contracts flowing into each region (i.e, regional exposure):

$$\exp_\ell := \frac{1}{19} \sum_{t=2001}^{2019} \frac{G_{\ell,t}}{G_t}.$$

We first show that these two definitions of shares are highly correlated. In particular, Figure C2 presents a bin-scatter of the log of the baseline analysis exposures ($\exp_\ell$) against the log of canonical shift-share shares (w_ℓ). We take logs to mitigate the influence of the thick right tails in the shares distribution. The figure reveals a strong positive correlation: regions that historically received a large fraction of total defense contracts (high-exposure) are also those with high contract-to-wage ratios (high-shares).

Second, we assess whether high-share regions differ systematically in employment growth. We

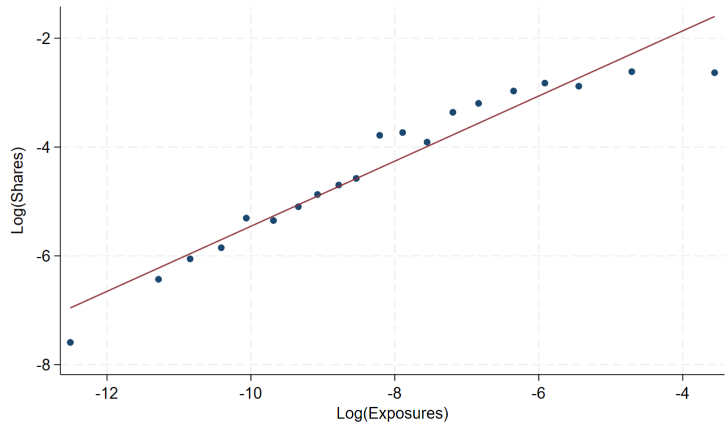

FIGURE C2 — (LOG)EXPOSURE AND (LOG)SHARES ARE HIGHLY CORRELATED

compute the average annual employment growth rate

$$\frac{1}{19}\sum_{t=2001}^{2019}\frac{E_{\ell,t}-E_{\ell,t-1}}{E_{\ell,t-1}}$$

for high- and low-share regions, defining "high share" as above the 75th percentile of the share distribution (0.039, i.e., defense contracts equal to at least 3.9% of wages and salaries; the median share is 1.2%). Figure C3 compares the distributions of average annual employment growth rates for high- and low-share regions. The figure shows no systematic difference. A cross-sectional regression of average

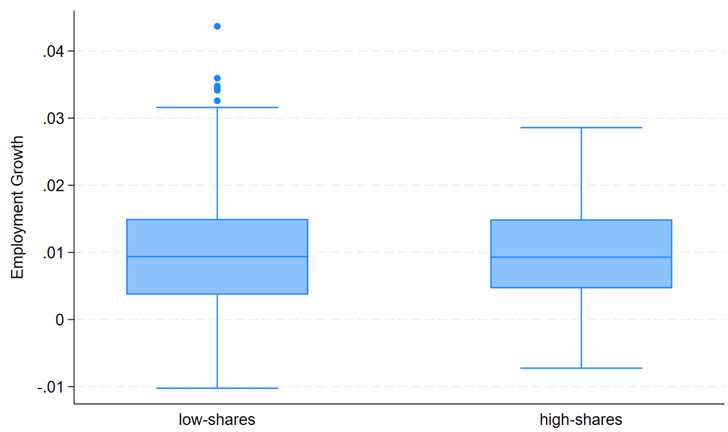

FIGURE C3 — HIGH- AND LOW-SHARE REGIONS HAVE SIMILAR ANNUAL EMP. GROWTH

employment growth on a constant and a high-share dummy confirms the absence of any statistically significant gap.

In summary, regions with higher baseline exposure to defense contracts coincide with high-share regions (Figure C2). Crucially, they do not exhibit systematically different average employment growth rates (Figure C3). Moreover, even if such differences existed, our baseline regression includes region fixed effects, which absorb any time-invariant heterogeneity in employment growth rates across MSAs.

Robustness Estimates. After having clarified the origin of the exogenous variation in our shift-share design, we re-estimate equation (1) by instrumenting the RHS with $\tilde{Z}_{\ell,t+h}$, the exact shift-share instrument, reporting the results in Table C1.

TABLE C1 — REGIONAL EMPLOYMENT MULTIPLIERS - SHIFT-SHARE INSTRUMENTATION

Response of Total Employment from (Public) BEA Data					
	IV: Shift-Share Instrument				
Horizon	*Coefficient* (β_h)	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.029	0.037	25.636	0.971	$1,030,040
	(0.014)			(0.464)	($492,296)
1 year	0.095	0.010	66.461	3.128	$319,701
	(0.036)			(1.205)	($123,147)
2 years	0.091	0.023	32.734	2.988	$334,660
	(0.040)			(1.304)	($146,073)
3 years	0.097	0.036	25.287	3.192	$313,282
	(0.046)			(1.518)	($149,023)

Notes: Sample: 2001-2019; 358 MSAs (QCEW+BDS+LAUS Harmonized Sample). GDP price deflator from BEA, base year 2008. Robust standard errors in parentheses, clustered at the MSA level. Montiel Olea and Pflueger (2013) effective F is calculated with `weakivtest`. Number of Job-Years refers to one million $. Standard errors of cost-per-job are obtained with the Δ-method.

Employment multipliers (β_h) are estimated precisely at all horizons and their values are similar to those ones obtained from of the baseline analysis (Table 2). Similarly, the effective F-statistics are consistent with those ones obtained with the instrument reported in the paper. The conversion factor used to transform multipliers into estimates of job-years is not affected by the choice of the instrument. By consequence, both values of job-years and cost-per-job are similar to the baseline estimates reported in the paper.

Overall, adopting an exact shift-share instrument approach yields identical estimates to our baseline approach based on Auerbach, Gorodnichenko, and Murphy (2020).

C.1. Employment Multiplier using Initial Exposure

Readers may be concerned that the exposure measures (the share of each MSA in total defense spending) used in constructing our baseline instrument are not fully exogenous because they average the defense spending share over the length of the sample. Thus, if the shares are not constant over time, these averages could reflect endogenous reallocation of defense spending shares in response to the shocks we are studying. To address this concern, we re-estimate the baseline employment effects presented in Table 2, but using initial (2001) defense spending shares to construct the instrument in all periods.

As is evident from the Table C2, this instrument is significantly weaker than the standard instrument used in the literature at longer horizons, perhaps because the 2001 defense spending share is a noisy measure of the underlying long-run share. However, estimates are close to those in the baseline,

TABLE C2 — REGIONAL EMPLOYMENT MULTIPLIERS - SHIFT-SHARE INST. (INITIAL SHARES)

Response of Total Employment from (Public) BEA Data					
	IV: Shift-Share Instrument, Starting Exposure				
Horizon	*Coefficient (β_h)*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.014	0.387	25.96	0.46	2,175,813
	(0.016)			(0.53)	(2,510,972)
1 year	0.079	0.055	9.64	2.62	382,198
	(0.041)			(1.36)	(198,483)
2 years	0.076	0.076	7.87	2.51	398,467
	(0.043)			(1.41)	(223,953)
3 years	0.081	0.105	6.22	2.67	374,797
	(0.050)			(1.64)	(230,964)

Notes: Sample: 2001-2019; 358 MSAs (QCEW+BDS+LAUS Harmonized Sample). GDP price deflator from BEA, base year 2008. Robust standard errors in parentheses, clustered at the MSA level. Montiel Olea and Pflueger (2013) effective F is calculated with `weakivtest`. Number of Job-Years refers to one million $. Standard errors of cost-per-job are obtained with the Δ-method.

underlining that the magnitude of the estimates is not an artifact of the time averaging of exposure shares.

D Robustness of Fact 2

D.1. Validation of BDS Employment Data.

As a validation exercise, we re-estimate Equation (1) using private employment from the Business Dynamics Statistics (BDS) instead of BEA total employment, while keeping the same harmonized QCEW+BDS+LAUS baseline sample. The goal is to verify that the aggregate employment multipliers obtained from BDS and BEA data are closely aligned, thereby validating the firm-size decomposition. Table D1 reports the results.

TABLE D1 — ESTIMATES FROM BUSINESS DYNAMICS STATISTICS CONSISTENT WITH BASELINE

	Response of Private Employment from (Public) BDS Data			
Horizon	*Coefficient* (β_h)	*p*	*Effective F*	*Job-Years ($1M)*
impact	0.062	0.056	29.232	1.296
	(0.032)			(0.676)
1 year	0.108	0.008	53.120	2.259
	(0.040)			(0.842)
2 years	0.115	0.015	27.603	2.393
	(0.047)			(0.981)
3 years	0.121	0.027	21.063	2.527
	(0.055)			(1.140)

Notes: Sample: 2001-2019 - 358 MSAs (QCEW+BDS+LAUS Harmonized Dataset). Data source: Business Dynamics Statistics (BDS). All else equal to Table 2.

The BDS-based employment multipliers are qualitatively similar to the baseline estimates obtained using BEA total employment, confirming that BDS data replicate the main findings of Section III. The implied number of job-years is somewhat smaller when using BDS private employment than when using BEA total employment, which is expected because total employment also captures public-sector responses. Ramey (2013) and Conley and Dupor (2013) note that part of the employment response to government spending may operate through public-sector employment rather than newly created private-sector jobs.

These results support the use of BDS employment data for the firm-size decomposition.

D.2. Results From the Largest Sample

TABLE D2 — REGIONAL EMPLOYMENT MULTIPLIERS FROM BDS DATA - LARGEST SAMPLE

BDS PRIVATE EMPLOYMENT - LARGEST SAMPLE: 2001-2019; 373 MSAS					
Horizon	*Coefficient*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.040	0.074	9.843	0.837	$1,194,908
	(0.022)			(0.467)	($667,073)
1 year	0.100	0.003	68.547	2.085	$479,521
	(0.033)			(0.693)	($159,427)
2 years	0.110	0.005	32.324	2.289	$436,902
	(0.039)			(0.808)	($154,168)
3 years	0.134	0.025	9.325	2.775	$360,299
	(0.059)			(1.233)	($160,050)

Notes: Robustness of Table D1 (regional employment multipliers using BDS data). Largest available sample from BDS.

TABLE D3 — MULTIPLIER BREAKDOWN: SMALL VS MEDIUM VS LARGE FROM BDS DATA
LARGEST SAMPLE

Breakdown by Size - Largest Sample: 2001-2019; 373 MSAs									
	Small Firms			*Medium-Sized Firms*			*Large Firms*		
Horizon	*Coefficient*	*p*	*Fraction (%)*	*Coefficient*	*p*	*Fraction (%)*	*Coefficient*	*p*	*Fraction (%)*
impact	-0.001	0.821	-2.1%	-0.003	0.717	-6.6%	0.044	0.035	108.8%
	(0.004)			(0.007)			(0.021)		
1 year	0.005	0.325	4.9%	0.001	0.908	1.3%	0.094	0.004	93.7%
	(0.005)			(0.012)			(0.032)		
2 years	0.004	0.439	3.7%	0.009	0.368	7.9%	0.097	0.005	88.4%
	(0.005)			(0.010)			(0.034)		
3 years	0.005	0.432	3.5%	0.008	0.473	6.1%	0.121	0.034	90.5%
	(0.006)			(0.011)			(0.057)		

Notes: Robustness of Table 4 (breakdown of regional employment multipliers by firm size using BDS data). Largest available sample from BDS.

D.3. Results From the Smallest Sample

TABLE D4 — REGIONAL EMPLOYMENT MULTIPLIERS FROM BDS - SMALLEST SAMPLE

BDS Private Employment - Smallest Sample: 2006—2019; 254 MSAs					
Horizon	*Coefficient*	*p*	*Effective F*	*Job-Years ($1M)*	*Cost-per-Job ($)*
impact	0.046	0.281	10.019	0.921	$1,085,957
	(0.042)			(0.853)	($1,005,478)
1 year	0.074	0.055	29.845	1.500	$666,831
	(0.039)			(0.778)	($346,073)
2 years	0.114	0.006	7.532	2.300	$434,829
	(0.041)			(0.823)	($155,623)
3 years	0.116	0.022	6.561	2.341	$427,197
	(0.051)			(1.018)	($185,696)

Notes: Robustness of Table D1 (regional employment multipliers using BDS data). Sample: Harmonized QCEW+BDS+LAUS+LDBE Dataset.

TABLE D5 — MULTIPLIER BREAKDOWN: SMALL VS MEDIUM VS LARGE - SMALLEST SAMPLE

Breakdown by Size - Smallest Sample: 2006-2019; 254 MSAs									
	Small Firms			*Medium-Sized Firms*			*Large Firms*		
Horizon	*Coefficient*	*p*	*Fraction (%)*	*Coefficient*	*p*	*Fraction (%)*	*Coefficient*	*p*	*Fraction (%)*
impact	0.003	0.717	6.4%	0.014	0.371	31.4%	0.028	0.460	62.2%
	(0.008)			(0.016)			(0.038)		
1 year	0.007	0.145	9.6%	0.016	0.231	21.9%	0.051	0.162	68.6%
	(0.005)			(0.014)			(0.036)		
2 years	0.004	0.393	3.9%	0.026	0.099	22.4%	0.084	0.014	73.8%
	(0.005)			(0.015)			(0.034)		
3 years	0.005	0.290	4.6%	0.027	0.079	23.3%	0.084	0.050	72.1%
	(0.005)			(0.015)			(0.043)		

Notes: Robustness of Table 4 (breakdown of regional employment multipliers by firm size using BDS data). Harmonized QCEW+BDS+LAUS+LDBE Dataset.

E Robustness of Fact 3

TABLE E1 — REGIONAL EMPLOYMENT MULTIPLIERS FROM LDBE DATA - LARGEST SAMPLE

LDBE Employment - Largest Sample: 2006—2019; 262 MSAs						
	Employment from LDBE (Restricted QCEW from BLS)					Employment from Public BEA
Horizon	*Coefficient*	*p*	*Effective F*	*Job-Years/$1M*	*Cost-per-Job ($)*	*Job-Years/$1M*
impact	0.012	0.557	10.836	0.275	$3,634,722	0.181
	(0.020)			(0.468)	($6,183,501)	(0.728)
1 year	0.098	0.003	10.554	2.282	$438,176	2.693
	(0.033)			(0.773)	($148,471)	(1.181)
2 years	0.100	0.019	7.125	2.348	$425,853	2.745
	(0.042)			(0.994)	($180,200)	(1.391)
3 years	0.117	0.037	5.466	2.727	$366,739	3.134
	(0.056)			(1.300)	($174,830)	(1.732)

Notes: Robustness of Table 7 (regional employment multipliers using restricted micro-data aggregated to MSAs). Largest available sample from LDBE.

TABLE E2 — MULTIPLIER BREAKDOWN: CONTRACTORS VS NON-CONTRACTORS FROM LDBE LARGEST SAMPLE

LDBE Employment - Largest Sample: 2006—2019; 262 MSAs						
	Contractors			*Non Contractors*		
Horizon	*Coefficient*	*p*	*Fraction*	*Coefficient*	*p*	*Fraction*
impact	0.036	0.040	306.8%	-0.024	0.205	-206.8%
	(0.017)			(0.019)		
1 year	0.063	0.032	64.4%	0.035	0.028	35.6%
	(0.029)			(0.016)		
2 years	0.052	0.044	51.4%	0.049	0.034	48.6%
	(0.026)			(0.023)		
3 years	0.056	0.074	47.8%	0.061	0.037	52.2%
	(0.031)			(0.029)		

Notes. Robustness of Table 6 (breakdown of regional employment multipliers using restricted micro-data aggregated to MSAs). Largest available sample from LDBE.

F Extra Establishment-level Results

F.1. Matched Sample

We merge contractors who receive at least one unpredictable contract with establishment-level outcomes from the QCEW.

First, we construct a list of contractors that received at least one unpredictable contract in a given year and county. Since the recipient-county field is not highly populated in the FPDS, we use the recipient zip code, which is rarely missing, to assign a geographic location to a contractor for a given year. We then use an official zip code-to-county crosswalk to map zip codes to counties. Second, we split the QCEW into year-county sub-samples, which report all establishment names. Almost all firms, identified by a unique employer identification number (EIN), appear to have a single establishment within a county. Third, we use a string-matching algorithm (`reclink`) to match all firms from our dataset of DUNS numbers that win an unpredictable contract with the universe of firm/EIN names within a given year and county from the QCEW.

String Matching Algorithm. We start by extensively cleaning and standardizing establishment names in both datasets, as well as the doing-business-as names present in both datasets and the parent name available in FPDS. We start matching firms within county and sharing the same first letter, using the *reclink* similarity index based on a pair of names and a pair of initial words to generate scores, and flagging only those with a similarity score above 0.92 as initial candidates. We first look for matches between the FPDS establishment name and LDBE establishment name, then between the FPDS establishment name and the LDBE doing-business-as name, then between the FPDS and LDBE doing-business-as names, then between the LDBE establishment name and FPDS doing-business-as name, and finally between the FPDS parent name and the LDBE establishment name. We then extensively filter flagged matches, accepting them only if they have extremely high match scores (above 0.995) or very small name distances relative to name length (i.e., only minor misspellings), and eliminating name matches based wholly on short initial words or frequently occurring words or names (unless these are unique within the region in both datasets). Establishments still unmatched within the county are then searched for matches within the CBSA; those few unmatched within CBSA are matched within the state. Blocking on initial letter and geography, removing matched establishments at each iteration of the geography loop, and parallelizing over states in multiple Stata instances makes it computationally feasible to match the large set of contract winners against the even larger set of establishments on the BLS's high-performance cluster.

Matched Sample Descriptive Statistics. We were able to match 13,662 establishments. The data cleaning process involved: (i) removing observations with incomplete histories, i.e., time series with gaps in the outcome variables; (ii) excluding firms with fewer than 13 quarters of observations (four quarters of lags, eight quarters for the impulse response function horizon, and one quarter for the shock); (iii) excluding firms whose first unpredictable contract appears before the fifth observation, as we control for four lags; (iv) excluding firms whose first unpredictable contract appears in the last eight quarters observed, as we assess the impulse response function with an eight-quarter horizon; (v) removing firms with fewer than one employee on average; and (vi) removing establishments with more than 150 employees. The resulting dataset is an (unbalanced) panel dataset with $N = 5,142$ firms observed from 2006:1 to 2019:4, $T = 56$. The median contract size is $114,900, while the mean is much larger, around $700,000, indicating a very long right tail in the contract distribution, consistent with the findings in Cox et al. (2024).

F.2. Products Purchased via Unpredictable Contracts

In this Appendix section we provide more details on the types of products purchased via unpredictable contracts.

Product Categories. Following Muratori, Juarros, and Valderrama (2023) we use the four-digit product category available from FPDS to distinguish between goods and service and aggregate products at 2-digits. Figure F1 shows the average fraction of total unpredictable contracts spent on the top ten service categories, where the average is taken over fiscal years. Similarly, Figure F2 shows the fraction of unpredictable contracts spent on the top ten goods categories.

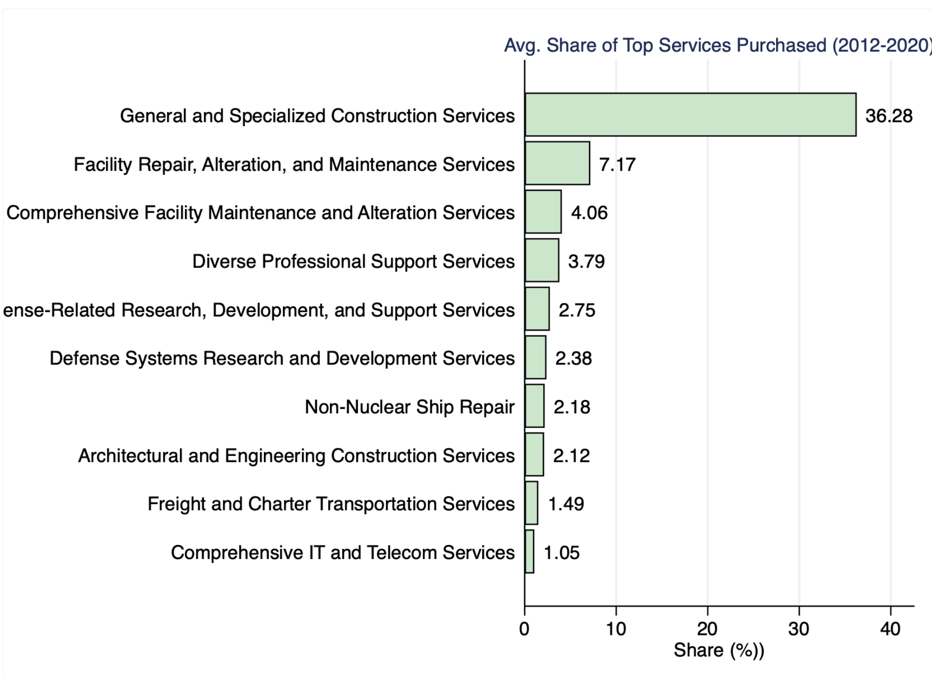

FIGURE F1 — TOP 10 SERVICES - FRACTION OF UNPREDICTABLE CONTRACTS

Almost half of all spending via unpredictable contracts are represented by construction-related services: general and specialized construction services (36.28%), facility repair, alteration and maintenance services (7.17%) and comprehensive facility maintenance and alteration services (4.06%). Moreover, more than 5% of spending originates from defense related R&D services.

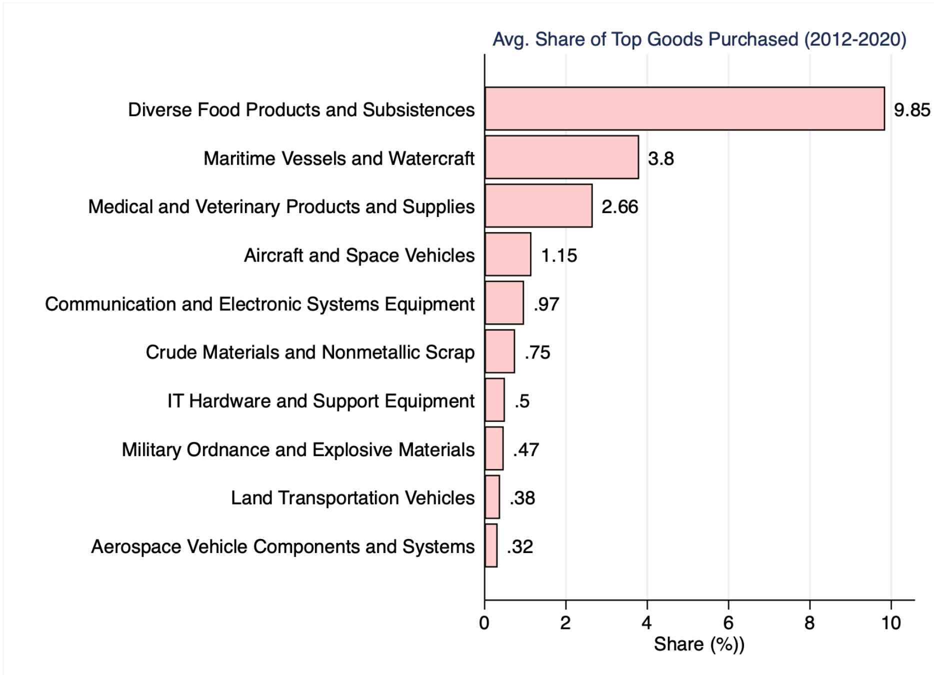

FIGURE F2 — TOP 10 GOODS - FRACTION OF UNPREDICTABLE CONTRACTS

Concerning goods, almost 10% of unpredictable contracts are spent on food products used, for instance, to supply military basis. Manufacturing goods strictly related to defense hardware accounts for about 7% of spending via unpredictable contracts: maritime vessels and watercraft (3.8%), aircraft and space vehicles (1.15%), communications and electronic equipment (0.97%), military ordnance and explosive materials (0.47%), land vehicles (0.38%) and aerospace vehicles components and systems (0.32%).

Duration of Unpredictable Contracts. Every contract in FPDS reports a period of performance start date and a period of performance current end date. We take the difference in days between the two to calculate the duration of all unpredictable contracts. We then plot in Figure F3 the box-whiskers plots of the duration (number of days) of unpredictable contracts by spending category.

Notice that service contracts tend to have a longer duration than contracts for goods. In the case of services the first quartile is 121 days, the median is 283 days, and the third quartile is 423 days. In the case of goods, the first quartile is 48 days, the median is 79 days, and the third quartile is 190 days.

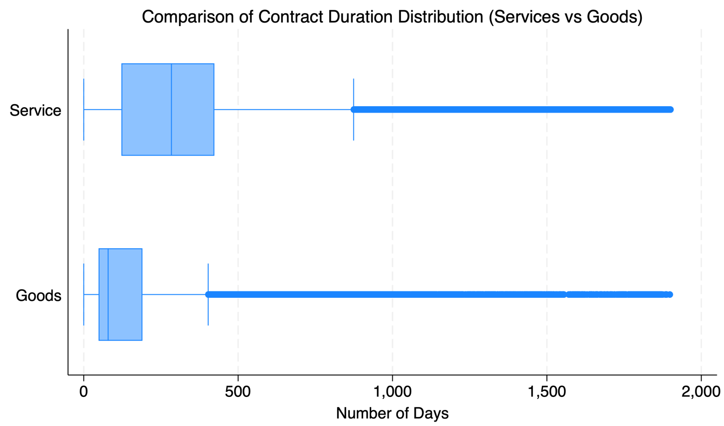

FIGURE F3 — IN-SAMPLE (UNWEIGHTED) DISTRIBUTION OF CONTRACTS' DURATION

F.3. Analysis by Quartile of Establishments

We subdivide the sample of establishments by analyzing each quartile of their size distribution separately. Establishments in the first quartile have between 1 and 6 employees, establishments in the second quartile have between 6 and 13 employees, and establishments 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 establishments.

Therefore, we re-estimate Equation (7) for each quartile of the establishments' size distribution separately to explore the robustness of the result across the sample. Figure F4 shows the IRFs of employment growth for each quartile. Note that the results appear to be robust across all four quartiles of the size distribution.

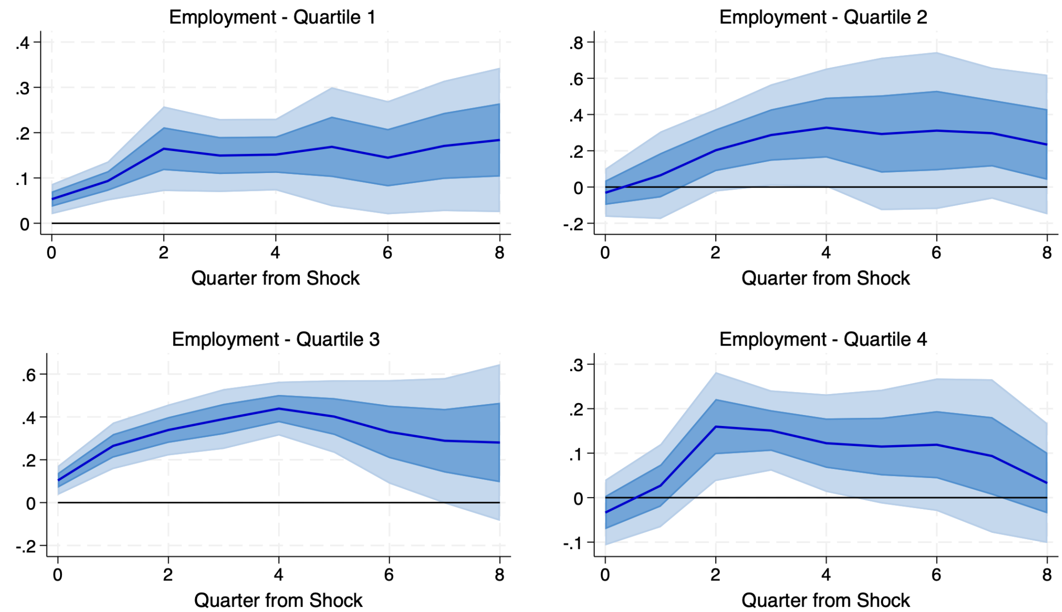

FIGURE F4 — EMPLOYMENT - QUARTILES OF SIZE DISTRIBUTION OF ESTABLISHMENTS

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

F.4. Time-Varying Productivity Shocks

The inclusion of establishment-fixed effects in the baseline equation (7) removes only the effects of systematic differences in productivity levels across establishments, 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 establishments temporarily more productive. Therefore, we are concerned that establishments 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).

To address this concern, we re-estimate Equation (7) 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.

Results are reported in Figure F5, where it is clear that the response of establishments' employment is robust to the inclusion of lags of wage-per-worker in the specification.

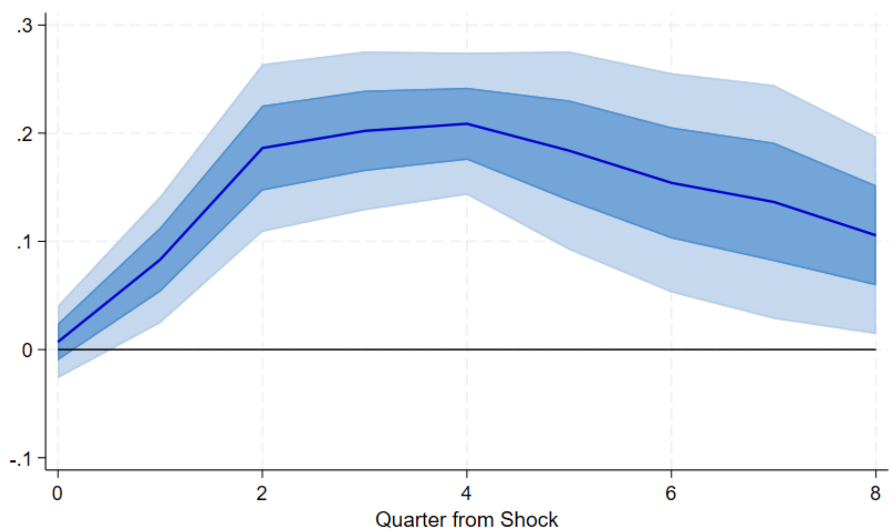

FIGURE F5 — RESPONSE OF ESTABLISHMENT'S EMPLOYMENT CONTROLLING FOR LAGS OF AVERAGE WAGE

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

F.5. Response of Wages

The restricted QCEW data also provide quarterly values of total wages paid by the establishment. We re-estimate Equation (7) using changes in total wages as the outcome variable. We report the results in Figure F6.

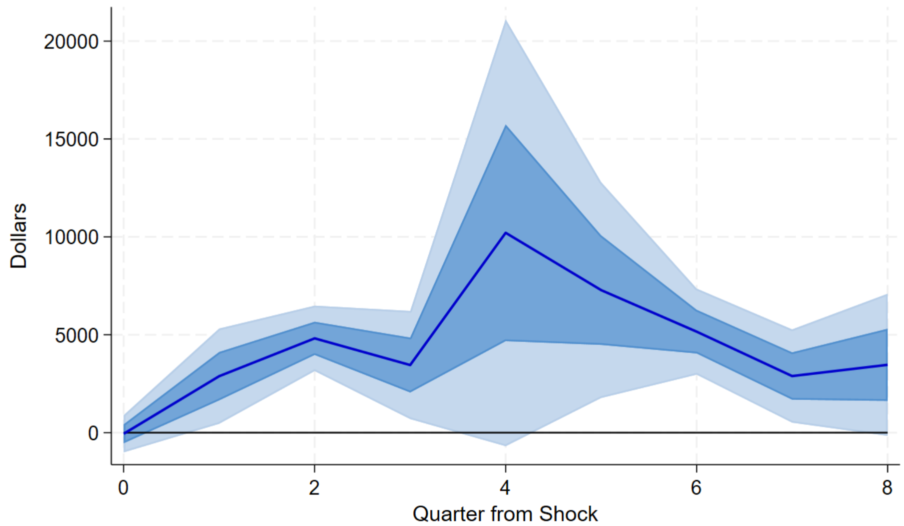

FIGURE F6 — RESPONSE OF TOTAL WAGES

Notes: The unit on the left axis is dollars, in response to a $1 million worth unanticipated contract. Firms are observed from 2006:1 to 2019:4, i.e., $T = 56$. Number of firms is $N = 5,142$. Standard errors are clustered at the state level. Small bands are 68% confidence. Large bands are 95% confidence.

Not surprisingly, unanticipated contracts have positive and significant effects on total wages, consistently with the positive and significant response of employment.

Second, we ask whether unanticipated contracts have any meaningful effect on the average wage paid to employees of winning establishments. In particular, we study the response of the average wage, or wage-per-worker, using the same specification as Equation (7). Results are reported in Figure F7, which displays no significant effect on wage-per-worker.

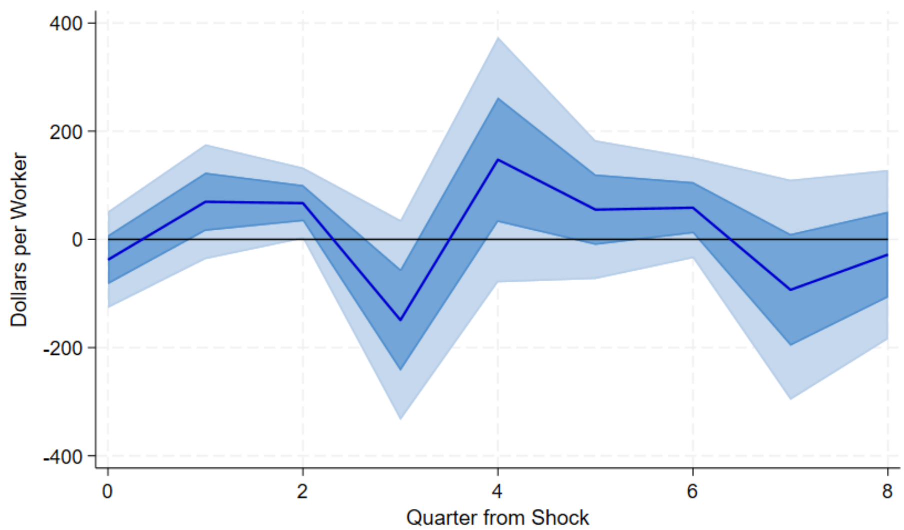

FIGURE F7 — RESPONSE OF AVERAGE WAGE

Our finding that employment, rather than wages, is the primary adjustment margin is consistent with Komarek, Butts, and Wagner (2022), who study procurement cuts under the Budget Control Act and also find that local labor markets primarily respond through employment.

F.6. NETS Analysis

To corroborate the establishment-level employment responses to procurement contracts presented in Section VI, we construct a panel dataset linking federal contract records from the Federal Procurement Data System (FPDS) to detailed employment microdata from the National Establishment Time Series (NETS).[29] This micro-level analysis allows us to directly estimate the effect of procurement inflows on employment using an alternative dataset to the inaccessible BLS dataset we use in our main analysis.

Given well-documented limitations of the NETS dataset, including inflated establishment counts, inconsistencies in employment reporting over time, and potential misclassification of firm ownership and geographic location, we implement a comprehensive series of cleaning and consolidation steps to construct a reliable panel of establishments. Establishments are defined as unique combinations of firm ownership, geographic location (ZIP code), and partial address, following the approaches by Barnatchez, Crane, and Decker (2017), Crane and Decker (2019), and Choi, Penciakova, and Saffie (2023). Firm ownership is traced through a recursive mapping of headquarters identifiers, resolving chains of ownership to identify ultimate parent firms. Observations listing themselves as their own headquarters are adjusted by deducting one employee, in line with standard practice to address over-reporting. After that, we retain only establishments with positive employment. Sectors not covered by the Business Dynamics Statistics (BDS), such as education and public administration and observations with fewer than 10 or more than 1,000 employees are excluded (Barnatchez, Crane, and Decker, 2017). Finally, we exclude from the sample all establishments with at least one imputed observation for employment. The cleaning assumptions are designed to ensure consistency with prior literature while preserving the granularity required for establishment-level analysis.

We use panel local projections to estimate the effect of \$1 of unanticipated contracts on employment (Jordà, 2005). In particular, given that NETS data is of annual frequency, we adapt Equation (7) and estimate via OLS the following equation:

$$E_{i,t+h} - E_{i,t-1} = \beta^h \cdot \varepsilon^g_{i,t} + \gamma^h_0 \cdot \widetilde{G}_{i,t} + \text{Lags} + \underbrace{\alpha^h_i + \alpha^h_{s,t} + \alpha^h_{\ell,t}}_{\text{Fixed Effects}} + \upsilon_{i,t+h} \quad h = 0, 1, 2, 3, \tag{8}$$

where $E_{i,t+h}$ denotes the h-period ahead number of employees; $\varepsilon^g_{i,t}$ denotes the dollar value of unanticipated contracts awarded to establishment i in year t, while $\widetilde{G}_{i,t}$ indicates the dollar value of potentially anticipated contracts. Both are expressed in units of \$1,000,000 of 2008 dollars. Lags $:=$ $\sum_{j=1}^{3}\left\{\rho^h_j \cdot \varepsilon^g_{i,t-j} + \gamma^h_j \cdot \widetilde{G}_{i,t-j} + \phi^h_j \cdot (E_{i,t-j} - E_{i,t-1-j})\right\}$. α^h_i represents an establishment fixed effect, $\alpha^h_{s,t}$ is a sector-time fixed effect intended to absorb any sectoral business-cycle effects. Lastly, $\alpha^h_{\ell,t}$ represents a state-time fixed effect, capturing regional business-cycle effects within a state. Our sample is composed of 28,393 establishments between 2006 and 2019.

The OLS estimates of β^h can be interpreted as impulse response functions (IRF) of the effect of an extra dollar of spending on establishment-level employment. Table F1 presents the results.

[29] NETS data were obtained by Ricardo Duque Gabriel under the purview of the Board of Governors' license agreement with the data provider. The remaining co-authors did not have any unauthorized access to NETS data while working on this paper. We thank Joonkyu Choi and Leland Crane for sharing insights and code to harmonize NETS with BDS.

TABLE F1 — EMPLOYMENT RESPONSE: NETS SAMPLE

Horizon	*Coefficient* (β_h^c)	*p*
impact	0.013	0.531
	(0.020)	
1 year	0.150	0.160
	(0.106)	
2 years	0.173	0.219
	(0.140)	
3 years	0.176	0.220
	(0.143)	

Notes: Sample: 2006—2019. 28,393 establishments. Coefficients are from Equation (8).

To validate the main results using an alternative data source, Figure F8 replicates the anticipation and placebo analyses from Figure 6 using the NETS dataset. While the estimates are notably noisier, consistent with the known limitations of NETS, the magnitude of the effects remains remarkably similar. This alignment reinforces the credibility of the coefficient used in the breakdown between the direct and indirect effect.

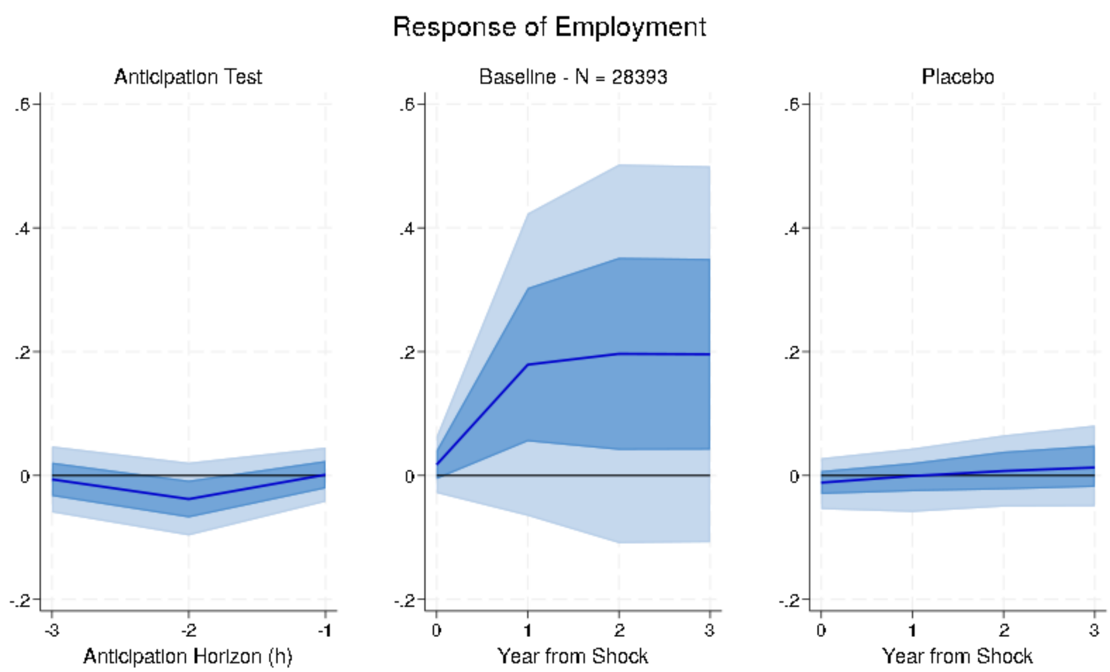

FIGURE F8 — EMPLOYMENT EFFECTS: NETS SAMPLE

Notes: Firms are observed from 2006 to 2019. The number of establishments is $N = 28,393$. Standard errors are clustered at the state level. Small bands represent 68% confidence intervals, and large bands represent 95% confidence intervals.

F.6.1. Robustness to Staggered Treatment Timing

A potential concern with the establishment-level analysis in Equation (8) is that the control group includes already-treated establishments, which could introduce bias under heterogeneous treatment effects. To address this, we implement a local projections difference-in-differences (LPDiD) design using "not-yet-treated" establishments and contractors that never received an unanticipated contract, as controls, following Dube et al. (2025).

Specifically, we restrict attention to establishments receiving their first unanticipated contract and compare their employment trajectories to establishments that will receive their first contract in the future but have not yet been treated and never treated establishments. This design directly addresses concerns about anticipation effects and selection into contractor status: if establishments were expanding *before* winning contracts we would observe differential pre-trends between treated and control groups.

For each horizon h, we estimate:

$$E_{i,t+h} - E_{i,t-1} = \delta^h \cdot D_{i,t} + \gamma^h \cdot \widetilde{G}_{i,t-1:t-2} + \phi^h \cdot E_{i,t-1} + \underbrace{\alpha_f^h + \alpha_{s,t}^h + \alpha_{\ell,t}^h}_{\text{Fixed Effects}} + u_{i,t+h} \tag{9}$$

where $D_{i,t}$ is an indicator equal to one if establishment i receives its first unanticipated contract in year t, and zero otherwise. The sample at each horizon h includes: (i) establishments receiving their first unanticipated contract in year t (treated), and (ii) establishments whose first unanticipated contract occurs more than h periods in the future (not-yet-treated controls). The notation $X_{i,t-1:t-2}$ denotes lags one and two of variable X. We include firm fixed effects α_f^h allowing for within-firm spillovers, alongside sector-time $\alpha_{s,t}^h$ and region-time $\alpha_{\ell,t}^h$ fixed effects. Standard errors are clustered at the CBSA level.

The coefficient δ^h captures the average treatment effect on the treated (ATT) of receiving a first unanticipated contract on establishment employment at horizon h. Unlike the baseline specification, which estimates the marginal effect per dollar of contract value, this design estimates the extensive margin effect of receiving an unanticipated contract for the first time.

Figure F9 presents the results. The pre-treatment coefficients ($h = -3$ to $h = -2$) are precisely estimated and statistically indistinguishable from zero, confirming the absence of differential pre-trends. This finding alleviates concerns that establishments were already on differential growth trajectories prior to receiving contracts, for instance, due to unobserved factors that simultaneously drive both employment growth and the likelihood of winning a contract.

Post-treatment, employment increases by approximately 2.2 workers at impact ($h = 0$) and continues to rise, reaching roughly 3.5 workers by year three. The persistent, gradual increase in employment is consistent with the dynamics observed in our baseline analysis and suggests that the effects of procurement contracts materialize over time as establishments scale up production.

Taken together, the flat pre-trends and positive post-treatment effects validate our baseline identification strategy. The results confirm that the employment responses documented in Section VI are not driven by staggered treatment bias or by pre-existing differential trends among establishments that

FIGURE F9 — LPDID: EMPLOYMENT RESPONSE TO FIRST UNANTICIPATED CONTRACT

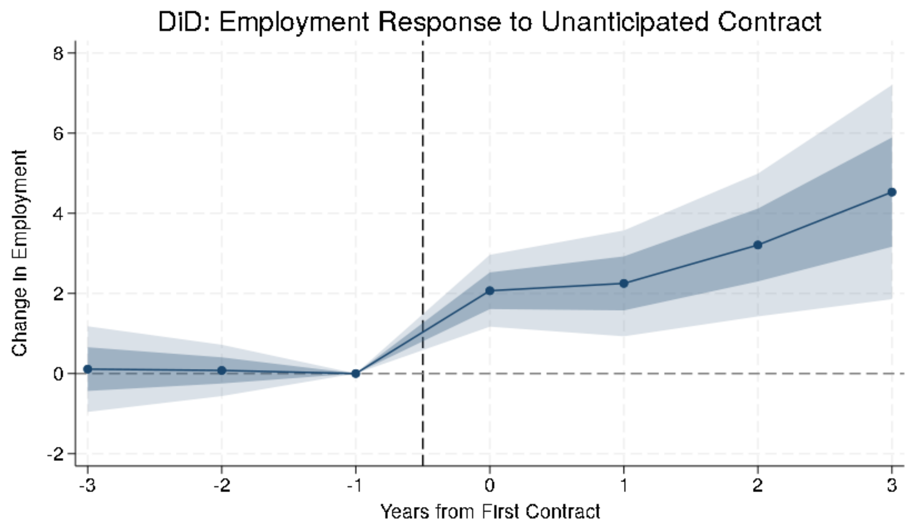

Notes: Sample: 2006–2019. The figure plots LPDiD estimates of Equation (9). Treatment is defined as receiving a first unanticipated contract. Controls are establishments whose first contract is more than h periods away (not-yet-treated). Shaded areas denote 68% (dark) and 95% (light) confidence intervals. Standard errors clustered at the CBSA level.

eventually become contractors.

F.6.2. Intensive Margin and Elasticity Estimates

The previous analysis estimates the extensive margin effect of receiving a first unanticipated contract. We now explore the intensive margin by replacing the binary indicator $D_{i,t}$ in Equation (9) with continuous measures of contract value, allowing us to estimate dose-response effects and elasticities.

Specification 1: Continuous Treatment. We replace $D_{i,t}$ with $\varepsilon^g_{i,t}$, the dollar value of the first unanticipated contract (in millions of 2008 dollars). The coefficient δ^h now captures the employment effect per \$1M of contract value, directly comparable to our baseline NETS estimates in Table F1 but estimated using the cleaner not-yet-treated and never treated contractors as control group. Figure F10 presents the results.

Specification 2: Normalized by Establishment Size. To account for heterogeneity in establishment size, we normalize both the outcome and treatment by lagged employment $E_{i,t-1}$. The outcome becomes employment growth $(E_{i,t+h} - E_{i,t-1})/E_{i,t-1}$, and the treatment becomes contract intensity $\varepsilon^g_{i,t}/E_{i,t-1}$ (contract dollars per baseline employee, in \$M). The coefficient δ^h captures the employment growth response to contract intensity: a contract of \$1M per employee increases employment growth by $\delta^h \times 100$ percentage points through horizon h. This specification gives greater weight to smaller establishments, where a given contract represents a larger relative shock. Figure F11 presents the results.

FIGURE F10 — LPDID: CONTINUOUS TREATMENT

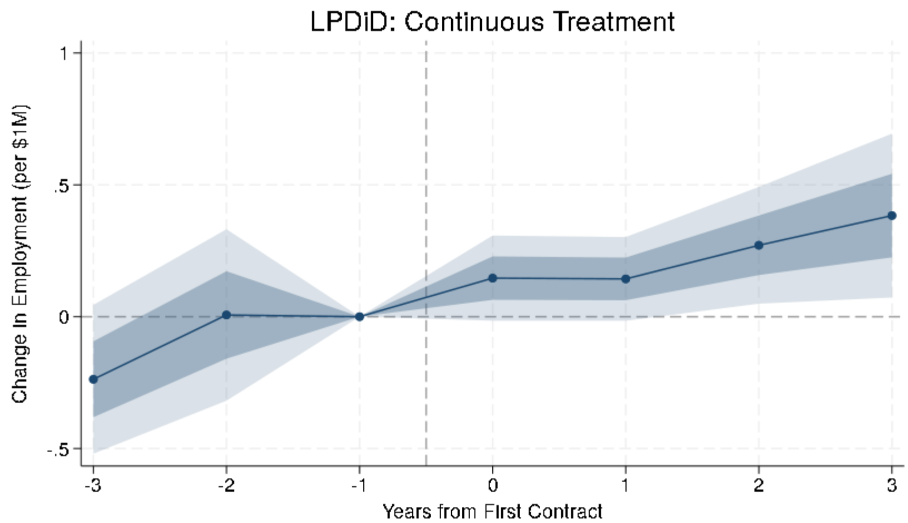

Notes: Sample: 2006–2019. Treatment is the dollar value of the first unanticipated contract ($M). Controls are not-yet-treated and never-treated establishments. Shaded areas denote 68% (dark) and 95% (light) confidence intervals. Standard errors clustered at the CBSA level.

FIGURE F11 — LPDID: NORMALIZED BY ESTABLISHMENT SIZE

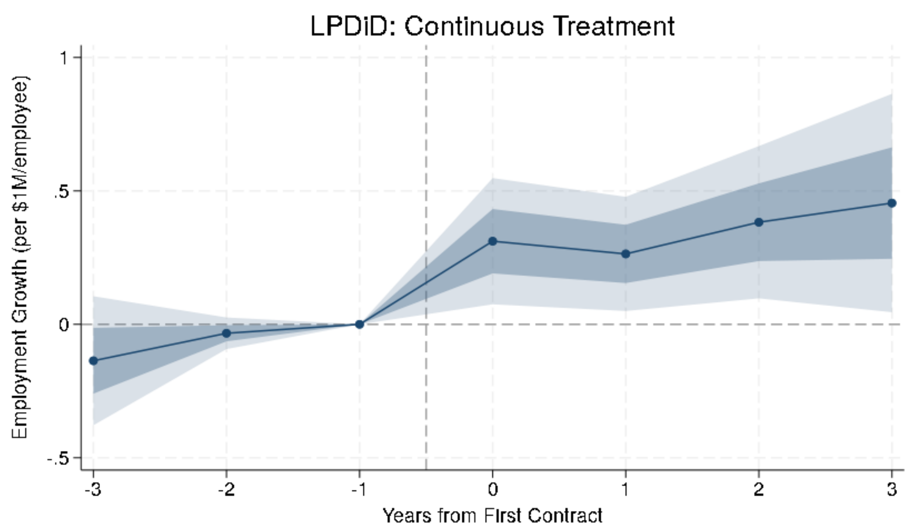

Notes: Sample: 2006–2019. Outcome is employment growth $(E_{i,t+h} - E_{i,t-1})/E_{i,t-1}$. Treatment is contract intensity $\varepsilon^g_{i,t}/E_{i,t-1}$ ($M per employee). Controls are not-yet-treated and never-treated and never treated contractors. Shaded areas denote 68% (dark) and 95% (light) confidence intervals. Standard errors clustered at the CBSA level.

Specification 3: Log-Log Elasticity. Finally, we estimate an elasticity specification with log employment growth $\log(E_{i,t+h}) - \log(E_{i,t-1})$ as the outcome and log contract value $\log(\varepsilon^g_{i,t})$ as the treatment (restricting to $\varepsilon^{g}_{i,t} > 0$). The coefficient δ^h is an elasticity: a 1% increase in contract value increases employment by δ^h% through horizon h. Figure F12 presents the results.

FIGURE F12 — LPDID: ELASTICITY SPECIFICATION

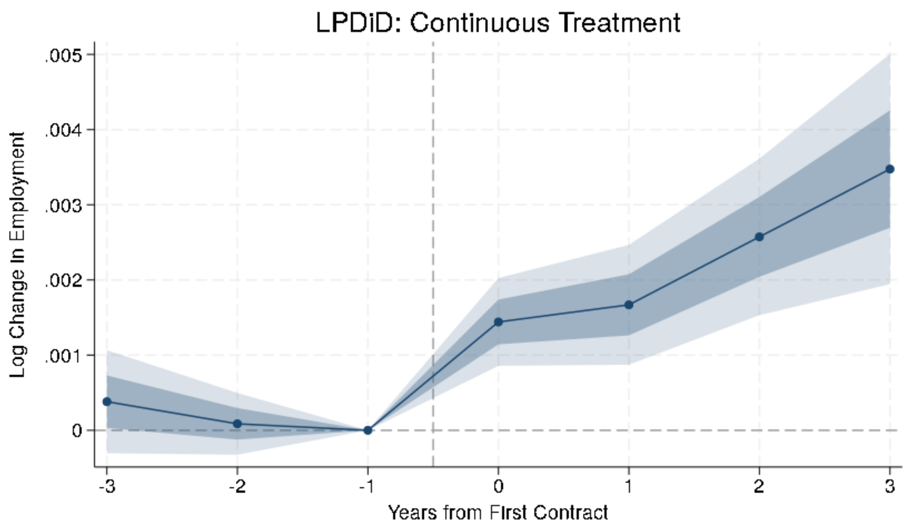

Notes: Sample: 2006–2019. Outcome is log employment growth $\log(E_{i,t+h}) - \log(E_{i,t-1})$. Treatment is $\log(\varepsilon^{g}_{i,t})$. Sample restricted to establishments with positive contract values. Controls are not-yet-treated and never treated contractors. Shaded areas denote 68% (dark) and 95% (light) confidence intervals. Standard errors clustered at the CBSA level.

Across all three specifications, the pre-treatment coefficients remain close to zero, confirming the absence of differential pre-trends regardless of how we measure treatment intensity. The post-treatment dynamics are qualitatively consistent with our baseline estimates. These results address several potential concerns with our identification strategy. The continuous treatment specification validates our baseline estimates under a cleaner identification design that excludes already-treated establishments from the control group. The size-normalized specification accounts for heterogeneity in establishment size, ensuring that our results are not driven by differences in local economic footprint across establishments. The log-log specification addresses the highly skewed distribution of contract values by estimating the elasticity of employment with respect to contract size. Together, these robustness checks confirm that establishments do not exhibit differential pre-trends prior to receiving contracts—alleviating concerns that expanding firms are simply more likely to win contracts—and that the employment effects of procurement are robust across alternative functional form assumptions.

F.7. Do Regional Defense Spending Shocks Include Unanticipated Contracts?

In this section, we test whether the regional defense spending shocks studied in the first part of the paper affect the award of unanticipated contracts studied in the establishment-level analysis. In particular, we break down the first-stage regression of the regional analysis using unanticipated contracts:

$$\frac{\underbrace{\left(\widetilde{G}_{\ell,t+h} + \varepsilon^g_{\ell,t+h}\right)}_{G_{\ell,t+h}} - \underbrace{\left(\widetilde{G}_{\ell,t-1} + \varepsilon^g_{\ell,t-1}\right)}_{G_{\ell,t-1}}}{Y_{\ell,t-1}} = \phi_h \cdot \underbrace{\frac{\exp_\ell \cdot (G_{t+h} - G_{t-1})}{Y_{\ell,t-1}}}_{\text{Regional Defense Spending Shock}} + \underbrace{\lambda_{\ell,h} + \alpha_{t,h}}_{\text{FEs}} + e_{\ell,t+h}.$$

We then study the effect of regional defense spending shocks on both $G_{\ell,t+h}$ (first stage) and the component which accrues to unanticipated contracts, using the harmonized QCEW+BDS+LAUS+BLS sample, which covers the same regions and years used in the establishment-level analysis. In particular, we estimate:

$$\frac{\varepsilon^g_{\ell,t+h} - \varepsilon^g_{\ell,t-1}}{Y_{\ell,t-1}} = \gamma_h \cdot \underbrace{\frac{\exp_\ell \cdot (G_{t+h} - G_{t-1})}{Y_{\ell,t-1}}}_{\text{Regional Defense Spending Shock}} + \underbrace{\lambda_{\ell,h} + \alpha_{t,h}}_{\text{FEs}} + e_{\ell,t+h}.$$

OLS estimates of the first stage coefficients (ϕ_h) and their unanticipated contracts components (γ_h) are reported in Table F2.

TABLE F2 — DO REGIONAL DEFENSE SP. SHOCKS IMPACT UNANTICIPATED CONTRACTS?

	UNANTICIPATED CONTRACTS		TOTAL CONTRACTS	
Horizon	*Coefficient* (γ_h)	*pvalue* (%)	*Coefficient* (ϕ_h)	*pvalue* (%)
impact	0.042	2.13%	0.984	0.10%
	(0.018)		(0.297)	
1 year	0.045	1.31%	1.858	0.00%
	(0.018)		(0.324)	
2 years	0.037	1.03%	1.675	0.40%
	(0.014)		(0.579)	
3 years	0.034	7.33%	2.169	0.70%
	(0.019)		(0.800)	

Notes: Sample: 2006-2019; 254 MSAs (QCEW+BDS+LAUS+BLS Harmonized Sample). GDP price deflator from BEA, base year 2008. Robust standard errors in parentheses, clustered at the MSA level.

A regional defense spending shock also affects the award of unanticipated contracts, $\varepsilon^g_{\ell,t}$, whose estimate is positive and statistically significant at any point in time. The response of unanticipated contracts accounts for about 2-4% of all contract awards in response to a shock, depending on the horizon. This fraction is consistent with the average fraction of unanticipated contracts out of total (see Figure 2).

Bibliography

Acemoglu, **Daron**, **Akcigit**, **Ufuk**, and **Kerr**, **William** (2016). "Networks and the Macroeconomy: An Empirical Exploration". *NBER Macroeconomics Annual* 30.2015, p. 63.

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.

Alesina, **Alberto**, **Favero**, **Carlo**, and **Giavazzi**, **Francesco** (Dec. 2014). "The output effect of fiscal consolidation plans". *Journal of International Economics* 96.2015, S19–S42.

Amodeo, **Francesco** and **Briganti**, **Edoardo** (2025). "High-Frequency Cross-Sectional Identification of Military News Shocks". *Forthcoming Bank of Canada Staff Working Paper*.

Auerbach, **Alan**, **Gorodnichenko**, **Yuriy**, and **Murphy**, **Daniel** (Mar. 2020). "Local Fiscal Multipliers and Fiscal Spillovers in the USA". *IMF Economic Review* 68.1, pp. 195–229. ISSN 2041-4161, 2041-417X.

— (July 2024). "Macroeconomic Frameworks: Reconciling Evidence and Model Predictions from Demand Shocks". *American Economic Journal: Macroeconomics* 16.3, pp. 190–229. ISSN 1945-7707, 1945-7715.

— (2025). "Demand Stimuls as a Social Policy". *Working Paper*.

Barattieri, **Alessandro**, **Cacciatore**, **Matteo**, and **Traum**, **Nora** (July 2025). "Estimating the Effects of Government Spending Through the Production Network". *Working Paper*.

Barnatchez, **Keith**, **Crane**, **Leland D.**, and **Decker**, **Ryan A.** (Nov. 2017). "An Assessment of the National Establishment Time Series (NETS) Database". *Finance and Economics Discussion Series* 2017.0.110. ISSN 1936-2854.

Barrot, **Jean-Noël** and **Nanda**, **Ramana** (Dec. 2020). "The Employment Effects of Faster Payment: Evidence from the Federal Quickpay Reform". *The Journal of Finance* 75.6, pp. 3139–3173. ISSN 0022-1082, 1540-6261.

Bartal, **Mehdi** and **Becard**, **Yvan** (May 2024). "Welfare Multipliers.pdf". *Working Paper*.

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.

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". *Woriking Paper*.

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

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.

Budrys, **Zymantas** (Oct. 2022). "Consumer of Last Resort: Government procurement, firm-level evidence and the macroeconomy". *Working Paper*.

Carril, **Rodrigo**, **Gonzalez-Lira**, **Andres**, and **Walker**, **Michael** (Feb. 2026). "Competition under Incomplete Contracts and the Design of Procurement Policies". *American Economic Review* 116.2, pp. 535–581.

Chetty, **Raj** et al. (2011). "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins". *The American Economic Review* 101.3, pp. 471–475.

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.

Chodorow-Reich, **Gabriel** et al. (Apr. 2012). "Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act". *American Economic Journal: Economic Policy* 4.3, pp. 118–145. ISSN 1945-7731.

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

Clemens, **Jeffrey**, **Hoxie**, **Philip**, and **Veuger**, **Stan** (2025). "Was Pandemic Fiscal Relief Effective Fiscal Stimulus? Evidence from Aid to State and Local Governments". *Journal of Macroeconomics* Forthcoming.

Conley, **Timothy G.** and **Dupor**, **Bill** (July 2013). "The American Recovery and Reinvestment Act: Solely a government jobs program?" *Journal of Monetary Economics* 60.5, pp. 535–549. ISSN 03043932.

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.

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

Crane, **Leland D.** and **Decker**, **Ryan A.** (May 2019). "Business Dynamics in the National Establishment Time Series (NETS)". *Finance and Economics Discussion Series* 2019.0.34. ISSN 1936-2854.

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.

Di Giovanni, **Julian** et al. (2026). "Buy Big or Buy Small? Procurement Policies, Firms' Financing, and the Macroeconomy". *Forthcoming American Economic Review*.

Driscoll, **John C.** and **Kraay**, **Aart C.** (Nov. 1998). "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data". *Review of Economics and Statistics* 80.4, pp. 549–560. ISSN 0034-6535, 1530-9142.

Dube, **Arindrajit** et al. (Nov. 2025). "A Local Projections Approach to Difference-in-Differences". *Journal of Applied Econometrics* 40.7, pp. 741–758.

Dupor, **Bill** and **Guerrero**, **Rodrigo** (Dec. 2017). "Local and aggregate fiscal policy multipliers". *Journal of Monetary Economics* 92, pp. 16–30. ISSN 03043932.

Dupor, **Bill** and **McCrory**, **Peter B.** (June 2018). “A Cup Runneth Over: Fiscal Policy Spillovers from the 2009 Recovery Act”. *The Economic Journal* 128.611, pp. 1476–1508. ISSN 0013-0133, 1468-0297.

Dupor, **Bill** and **Mehkari**, **M. Saif** (June 2016). “The 2009 Recovery Act: Stimulus at the extensive and intensive labor margins”. *European Economic Review* 85, pp. 208–228. ISSN 00142921.

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

Foschi, **Andrea**, **House**, **Christopher L**, and **Proebsting**, **Christian** (Mar. 2025). “Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks.”

Gabriel, **Ricardo Duque** (Oct. 2024). “The Credit Channel of Public Procurement”. *Journal of Monetary Economics*. Monetary Policy challenges for European Macroeconomies 147, p. 103601. ISSN 0304-3932.

Gabriel, **Ricardo Duque**, **Klein**, **Mathias**, and **Pessoa**, **Ana Sofia** (Aug. 2023). “The Effects of Government Spending in the Eurozone”. *Journal of the European Economic Association* 21.4, pp. 1397–1427. ISSN 1542-4766, 1542-4774.

Goldsmith-Pinkham, **Paul**, **Sorkin**, **Isaac**, and **Swift**, **Henry** (Aug. 2020). “Bartik Instruments: What, When, Why, and How”. *American Economic Review* 110.8, pp. 2586–2624. ISSN 0002-8282.

Gugler, **Klaus**, **Weichselbaumer**, **Michael**, and **Zulehner**, **Christine** (Feb. 2020). “Employment behavior and the economic crisis: Evidence from winners and runners-up in procurement auctions”. *Journal of Public Economics* 182, p. 104112. ISSN 00472727.

Hager, **Anselm** and **Huber**, **Kilian** (Apr. 2025). “Big Government and Dynamism Drain”. *Woriking Paper*.

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.

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

Joseph R. Biden, **Jr.** (May 3, 2022). *Remarks by President Biden on the Security Assistance to Ukraine*. Remarks delivered at Lockheed Martin Pike County Operations, Troy, Alabama. Available at: this link. Troy, Alabama: The White House.

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

Komarek, **Timothy M.**, **Butts**, **Kyle**, and **Wagner**, **Gary A.** (Nov. 2022). “Government Contracting, Labor Intensity, and the Local Effects of Fiscal Consolidation: Evidence from the Budget Control Act of 2011”. *Journal of Urban Economics* 132, p. 103506. ISSN 0094-1190.

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.

Montiel Olea, **José Luis** and **Pflueger**, **Carolin** (July 2013). “A Robust Test for Weak Instruments”. *Journal of Business & Economic Statistics* 31.3, pp. 358–369. ISSN 0735-0015, 1537-2707.

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

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.

Park, **Geumbi**, **Zhou**, **Xiaoqing**, and **Zubairy**, **Sarah** (Sept. 2025). “Subcontracting in Federal Spending: Micro and Macro Implications”. *Woriking Paper*.

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.

Piger, **Jeremy** and **Stockwell**, **Thomas** (Aug. 2025). “Differences From Differencing: Should Local Projections With Observed Shocks Be Estimated in Levels or Differences?” *Journal of Applied Econometrics* 40.7, pp. 759–787.

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.

— (2013). “Government Spending and Private Activity”. *Fiscal Policy after the Financial Crisis, edited by Alberto Aleina and Francesco Giavazzi*. University of Chicago Press, pp. 19–62.

Ramey, **Valerie** and **Zubairy**, **Sarah** (Mar. 2018). “Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data”. *journal of political economy*, p. 52.

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

U.S. Department of Defense (Feb. 2022). *State of Competition within the Defense Industrial Base*. Tech. rep.

Wilson, **Daniel J** (Aug. 2012). “Fiscal Spending Jobs Multipliers: Evidence from the 2009 American Recovery and Reinvestment Act”. *American Economic Journal: Economic Policy* 4.3, pp. 251–282. ISSN 1945-7731, 1945-774X.