Online Appendix

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

Edoardo Briganti Holt Dwyer Ricardo Gabriel Victor Sellemi

December 2024

Replication Instructions. The results of the paper are derived following the replication package guidelines suggested by Professor Ingar Haaland in his post on X.

- All results are replicated by running with one click the `master_dofile.do`, which runs multiple do-files stored in a `code` directory.
- All raw data are stored in a `raw_data` directory which is used to create the datasets used in the analysis, stored in a `data` directory.
- All output is stored in an `output` directory organized in a figures and tables directory.
- Tables' contents is automatically uploaded from `Stata` to `Overleaf`, also following his code.

*Emails: ebrigant@ucsd.edu, bdwyer@ucsd.edu, ricardo.f.duquegabriel@frb.gov, vsellemi@ucsd.edu.

Contents

A More On Competition Procedure

A.1 Solicitations

Contracts awarded competitively are solicited on a government website, Federal Bizz Opportunities, or FBO, now migrated to SAM.gov. Contracts 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 steps sealed bidding) or, when the nature of the product is more complex, written proposals (i.e. contract by negotiations). Once the offer periods expires, awardee are competitively selected. All pre-award notices are gathered daily on SAM.gov. Following Gonzalez-Lira, Carril, and Walker (2021) 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. Figure 1 summarizes the competitive procurement timeline process.

Figure 1: Timeline of Competed Contracts

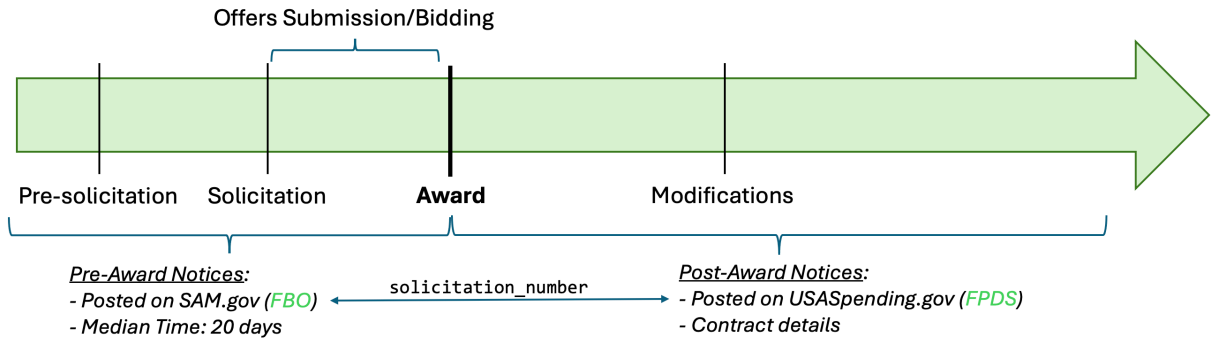

Notes: Once the contract is awarded, all detailed contract information is recorded in FPDS by the responsible federal contracting officer. Several "*post-award* actions follow the award, known as contract-modifications. Frequent examples of contract modifications are options to buy more from the government, extra-costs for extra work, appropriations of extra funds and contracts termination.

We keep all award histories from fiscal year 2006 to fiscal year 2019 to be consistent with the sample choice of the paper and then we analyze the number of days from the oldest pre-solicitation to the award notice, dropping solicitations which either (i) lack an award notice or (ii) consist only of a single notice. Figure 2 shows the box-plot of the (unweighted) number of days from the oldest pre-award notice to the award notice.

Figure 2: Box-Plot of Number of Days from Oldest Pre-Award Notice to Award

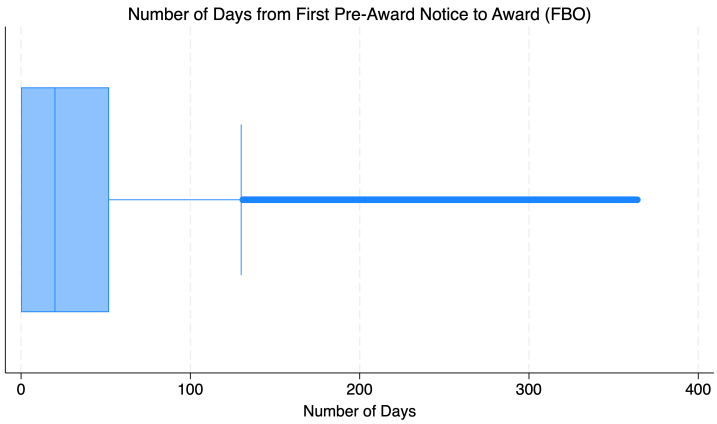

Notes: Distribution is not weighted by the value of a contract. Data source is the universe of federal procurement solicitation from FBO (Federal-Bizz-Opportunities.gov), now migrated to SAM.gov.

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, that is, well below the quarterly frequency used in the paper.

$\rightarrow$ In light of the short time period between pre-solicitations and award date, we use the award date available from FPDS, a much more complete and comprehensive dataset than FBO, to identify the timing of the award. We address potential anticipation effects owing to the pre-award solicitations period by carrying out anticipation tests in the main body of the paper.

A.2 Institutional Knowledge: Publicizing Requirements and Contract Notices

Although public procurement contracts are awarded at a highly decentralized level (i.e., by over 69 federal agencies, 209 sub-agencies), all contracting officers are required to abide by the guidelines proposed in the Federal Acquisition Regulation (FAR). FAR Part 5 (*Publicizing Contract Actions*) requires that contracting officers publicize contract opportunities with the goal of increasing competition, broadening industry participation, and assisting small businesses in obtaining contracts.

Since October 1, 2001, contract actions with an expected value of over $25,000 must be publicized in an online and easy-to-access government platform. Contract actions below the threshold might still be posted to increase visibility. On the other hand, FAR allows for exemptions to the requirement above the threshold when the posting might "compromise national security" or when the posting is "not in the government's interest". The result is that many contracts which are awarded are never solicited. When the regulation applies, Contract Opportunity notices are posted publicly at `beta.sam.gov` and include award notices such as solicitations, pre-solicitations, or other pre-award and post-award actions. We describe the most frequent types of contract notices below.[1]

Special Notice: Agencies use Special Notices to announce important pre-award events such as business fairs, long-run procurement forecasts, or pre-award conferences and meetings. Special Notices might also refer to "Requests for Information" (RFI) or draft solicitations.

Sources Sought: Agencies post Sources Sought Notices in order to seek possible sources for a project. As discussed in FAR 7.3, the Sources Sought notice is not a solicitation for work or a request for proposal. Agencies typically use Sources Sought notices to collect industry feedback on key contracting strategy decisions and to perform market research on firm capabilities.

[1] Gonzalez-Lira, Carril, and Walker (2021) also provides a useful description and analysis of the publicizing requirements for Federal Procurement and the effects of information diffusion via public notices. We thank Andres Gonzales-Lira for directing us to the General Services Administration Technical Documentation for the FedBizOpps (FBO) website, whose information is now migrated to Contract Opportunities.

Pre-Solicitation: Agencies post a pre-solicitation to notify vendors that a solicitation may follow. Potential vendors might then express interest in the contract by adding themselves to the Interested Vendors List. Government agencies use pre-solicitations to determine the number of qualified vendors to perform the desired work. Contracting officers can also use pre-solicitations to gather information on interested suppliers and determine if a set-aside for a small business might be applicable.

Solicitation: Agencies post a solicitation to clearly define government requirements for a potential contract so that businesses can submit competitive bids. A "Request for Proposal" (RFP) is the most common type of solicitation used by federal agencies. The solicitation also sets conditions and requirements for contractor proposals and includes the government's plan for evaluating submissions for potential award.

Combined Synopsis/Solicitation Agencies post a combined synopsis/solicitation when a contract is open for bids from eligible vendors. The Synopsis/Solicitation includes specifications for the product or service requested and a due date for the proposal, as well as the bidding procedures associated with the solicitation.

Award Notice Agencies post an award notice when they award a contract in response to a solicitation. Federal agencies may choose to upload a notice of the award to make aware other interested vendors of the winning bid. Note that the requirement guidelines for posting the award notice vary based on the agency and the solicitation.

B Unpredictable Contracts: Descriptive Stats

B.1 Sample Coverage

To summarize, our FPDS sample is made of all firms, identified using recipient DUNS, which have received at least one highly competed contract. A highly competed contract is defined as a (i) newly awarded, (ii) definitive, (iii) competed contract (iv) with at least 2 offers received. What fraction of federal procurement spending meets these four conditions?

We answer this question by collapsing the dataset by fiscal year and merge it with the time series of all federal contracts in FPDS to calculate the fraction of federal procurement spending, measured either in dollar value or number of new contracts, which meets conditions (i) through (iv). Figure 3 shows the results and also break down the statistics by type of spending: either goods (middle panels) or service (bottom panels).

Figure 3: Fraction of Contracts Meeting "Unpredictability" Conditions

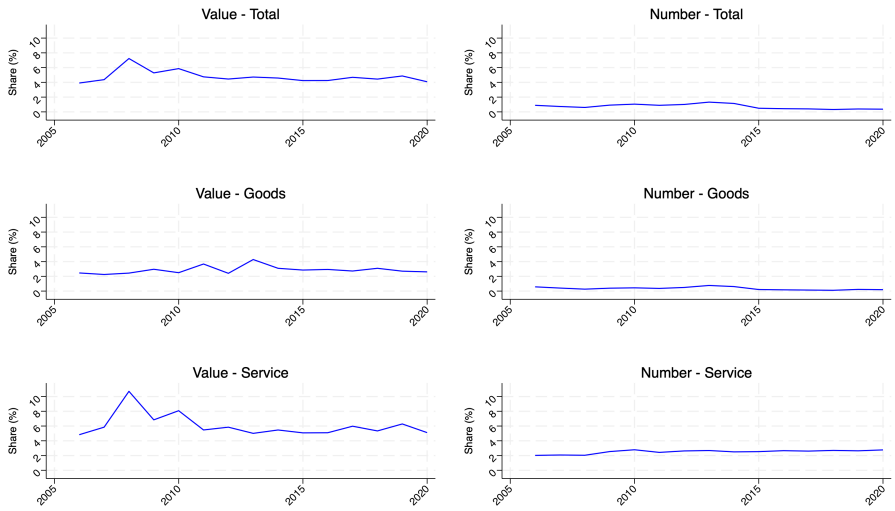

Notice that only 6% of contract value (left panels) from FPDS meets the firm-level "unpredictability" conditions, and only 1.5% of new contracts (right panels) are represented by highly competed contracts. Results are consistent when we look at a breakdown of spending into services and goods, with the former exhibiting slightly larger shares.

In-Sample Fraction of Unpredictable Contracts: Do firms who win unpredictable contracts only win potentially predictable contracts? This question is relevant for at least two reasons. First, if firms who win unpredictable contracts also win a large fraction of potentially predictable contracts, the effects of the former might be confounded with the effects of the latter. If this is the case, we should control for potentially predictable contracts to avoid getting an omitted variable bias. Second, if firms who win unpredictable contracts never win potentially predictable ones, there might be something special about these firms which might cast some doubt about the external validity of the results.

Therefore, we re-propose in Figure 4 the same descriptive statistics of Figure 3 but this time the ratio is taken over the aggregate value of all contracts won by firms in our sample that is, firms who win at least one unpredictable contract.

Figure 4: In-Sample Fraction of Contracts Meeting "Unpredictability" Conditions

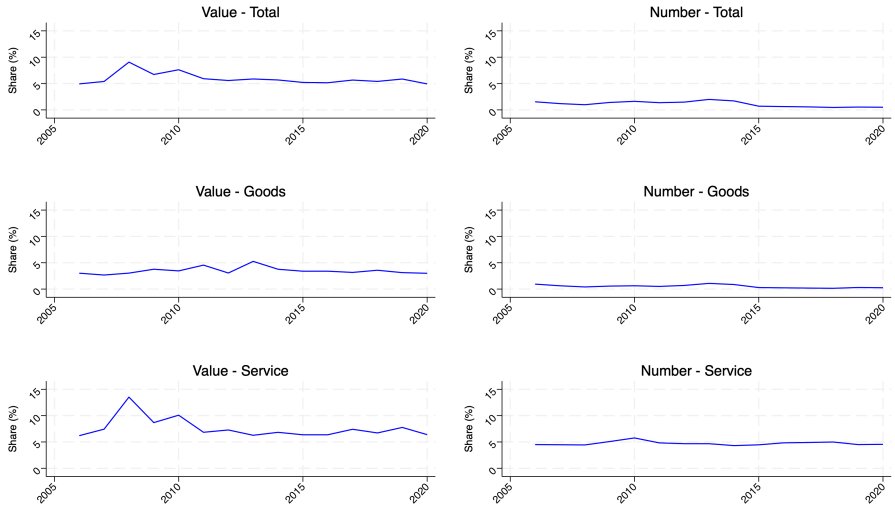

The results are clear: firms who win unpredictable contracts also win potentially predictable ones. The fractions are slightly higher compared to those ones in Figure 3 but still around 7-10%. If this is reassuring for the external validity of the results, it suggests to control for contemporaneous and lagged values of potentially predictable contracts to control for confounding effects.

B.2 Products Purchased via Unpredictable Contracts

What kind of products are bought via unpredictable contracts? We answer this question by calculating the average fraction of each product category out of total spending via unpredictable contracts in a given fiscal year. Average is taken over fiscal years. Following Muratori, Juarros, and Valderrama (2023) we use the four-digit product category to distinguish between goods and service and aggregate products at 2-digits.

Figure 5 shows the top ten services purchased and their average fraction of total spending in the sample.

Figure 5: Top 10 Services - Fraction of Unpredictable Contracts

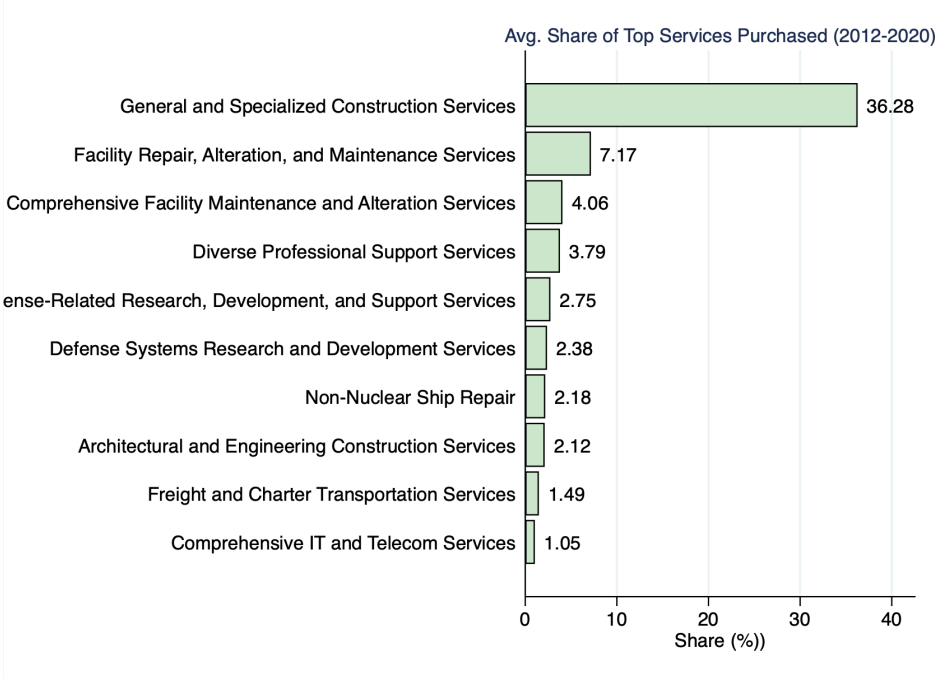

Notice that the first three services are construction services; in fact, almost half of all spending via unpredictable contracts are represented by construction-related services. Moreover, more than 5% of spending originates from defense related R&D services.

Figure 6 shows the fraction of unpredictable contracts spent on the top ten goods categories.

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

Figure 6: Top 10 Goods - Fraction of Unpredictable Contracts

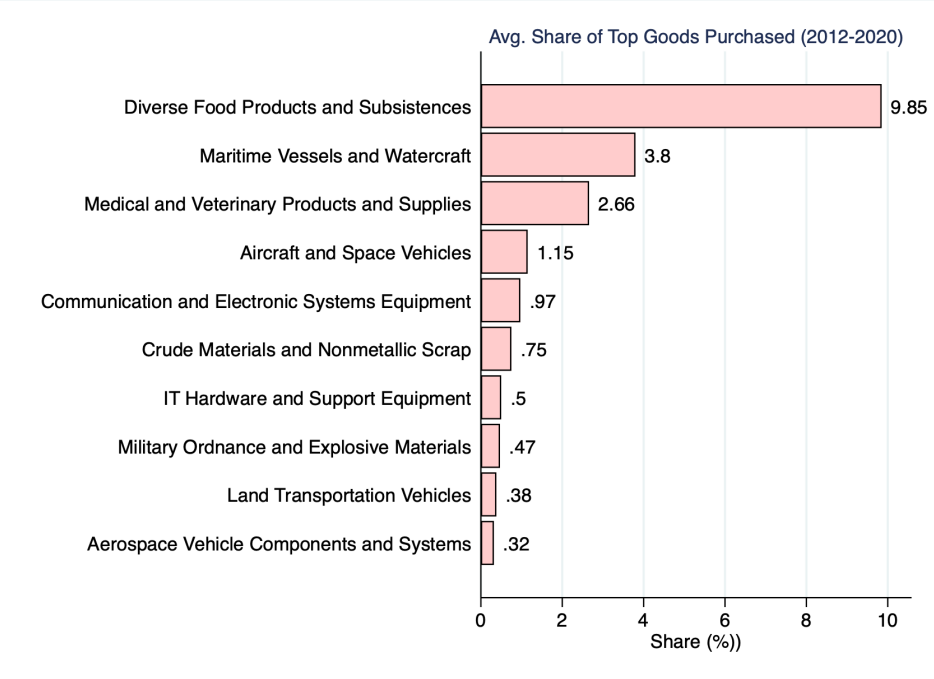

spending via unpredictable contracts: maritime vessels and watercraft (3.8%), aircraft and space vehicles (1.15%), communications and electronic equipment (0.97%), military ordnance and explosive materials (0.47%), land vehicles (0.38%) and aerospace vehicles components and systems (0.32%).

Lastly, we calculate the share of total contracts awarded for services: $G^s_{i,t}/G_{i,t}$. If the share is equal to or close to one, it means that the firm specializes as a service provider to the government, whereas if the share is close to zero, it means that the firm specializes as a goods supplier to the government. Figure 7 shows the distribution of average shares for services, where the average is taken within firms and over fiscal years.

Notice that the distribution is highly bimodal, with peaks at both zero and one, indicating that firms specialize as either goods suppliers or service providers. Notice also that the peak is higher for service providers, meaning that there are more service providers than goods suppliers in the sample, which is consistent with a large fraction of procurement spending being awarded for services rather than goods, as depicted in Figure 6 and Figure 5.

Figure 7: Top 10 Goods - Fraction of Unpredictable Contracts

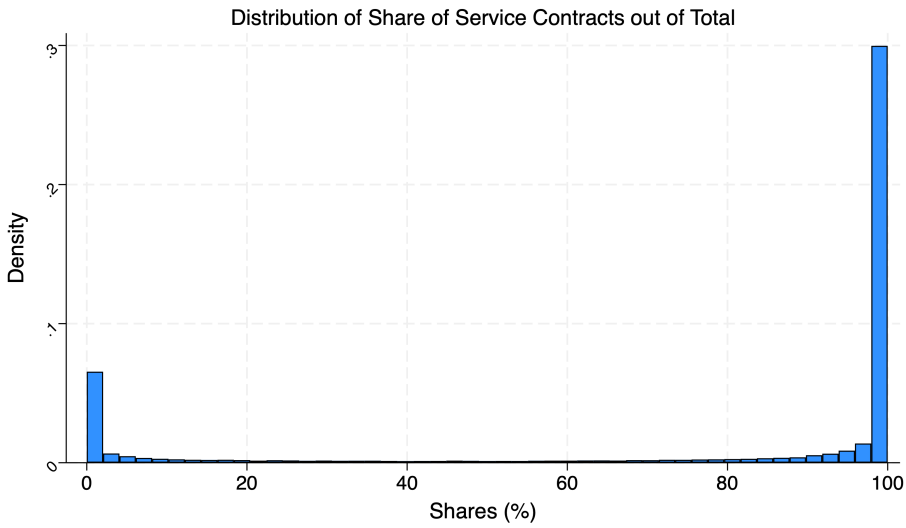

B.3 Duration of Unpredictable Contracts

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.

Figure 8 shows 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. These numbers are much larger than the unweighted contract duration reported in Cox et al. (2023) of 25 days, suggesting that unpredictable contracts have longer duration than more standard task/delivery orders, or non-competed contracts.

Muratori, Juarros, and Valderrama (2023) find that service spending has much large multipliers than goods spending, suggesting that labor intensity might explain the differential effects of the two. The much longer duration of service contracts might also explain why service multipliers are larger, since it is very likely that contracts with longer duration have stronger and more lasting

Figure 8: In-Sample (Unweighted) Distribution of Contracts' Duration

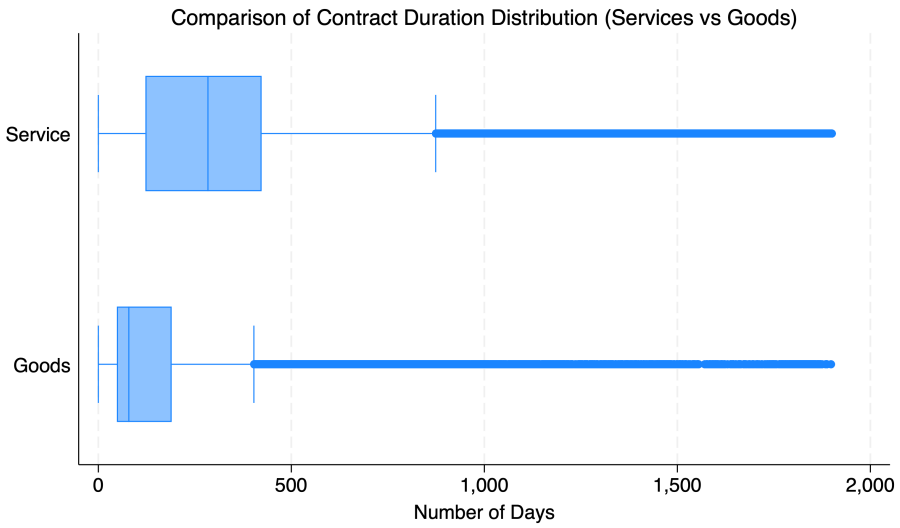

effects than short-lived ones.

B.4 Spatial Distribution of Unpredictable Contracts

Figure 9 shows the spatial distribution of the average number of firms/DUNS which receive at least one unpredictable contract by counties in a fiscal year. The average is taken over FY2016 through FY2018. Similarly, Figure 10 shows the spatial distribution of the average number of unpredictable contracts by counties by fiscal year.

Figure 9: Spatial Distribution - Number of Recipients of Unpredictable Contracts

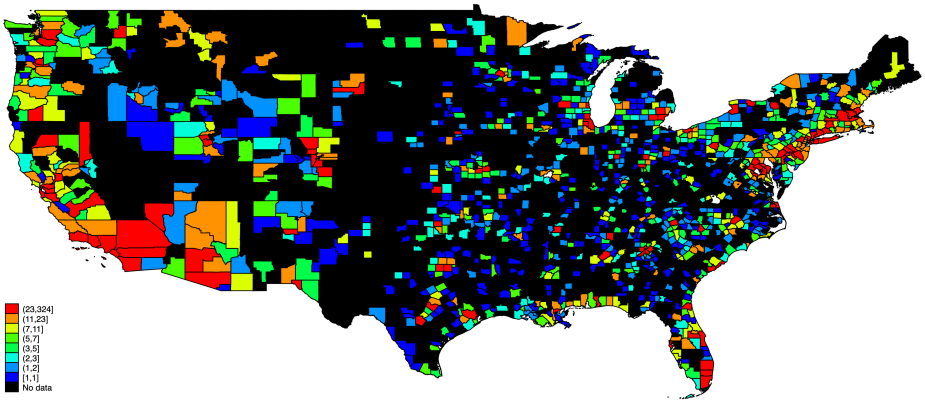

Notes: Data is aggregated by fiscal year. Results show geographic average of fiscal years 2016 through 2018. Geographic unit: county.

Figure 10: Spatial Distribution - Number of Unpredictable Contracts

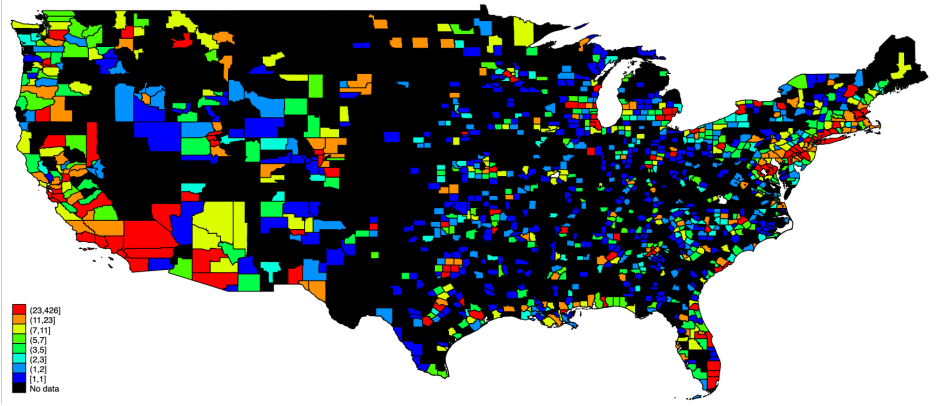

Notes: Data is aggregated by fiscal year. Results show geographic average of fiscal years 2016 through 2018. Geographic unit: county.

Unpredictable contracts and their recipients tend to be spatially located close to military installments as well as the DC area. Top recipient regions remain the Forth Worth-Dallas area (TX) and

San Diego county. This is consistent with previous findings from Auerbach, Gorodnichenko, and Murphy (2020), suggesting that unpredictable contracts have common destinations with the rest of contracts.

B.5 Extra: Fraction of Procurement Spending Awarded Domestically

In Section II of the paper we argue that more than 90% of procurement spending after the 2000 is awarded domestically. Meaning that 90% of spending is carried out with a place of performance in the U.S.

Figure 11 shows the distribution of defense (left panel) and non-defense (right panel) procurement spending by fiscal year. Contracts spending is categorized as either abroad, US or US Territories based on the primary place of performance.

Figure 11: Geographic Distribution of Procurement Spending After the 2000

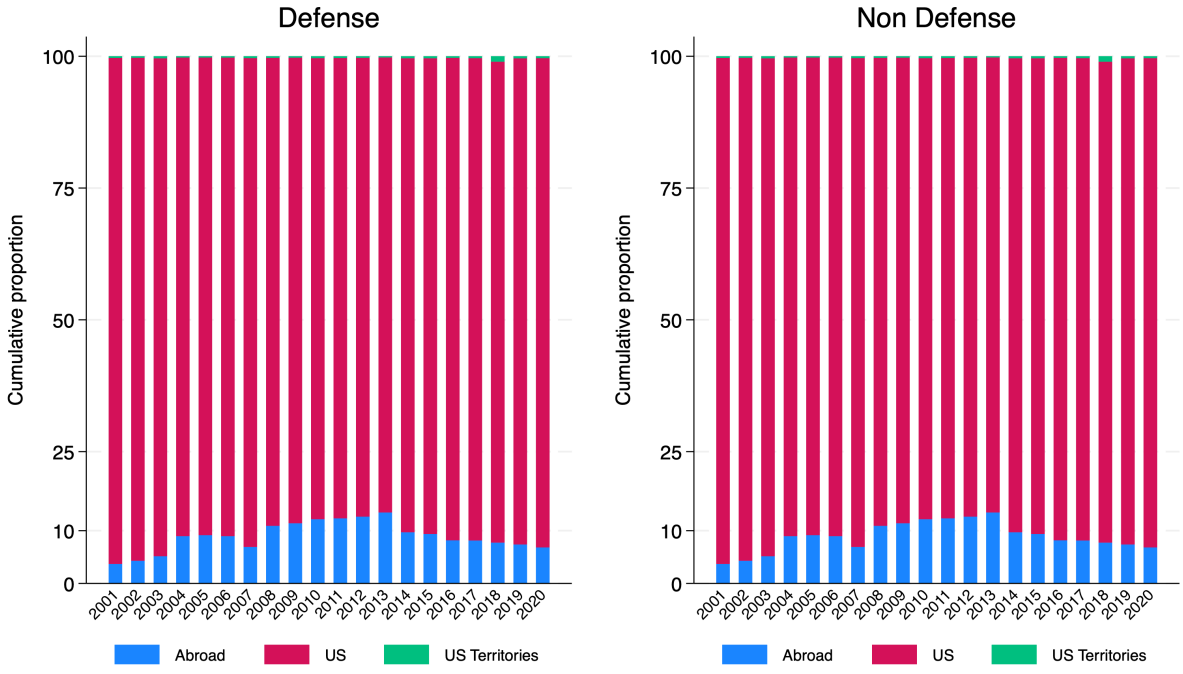

On average, less than 10% of procurement contracts have a primary place of performance different from a US based location.

C Descriptive Statistics on Matched-Sample

We merge contractors who receive at least one unpredictable contract with firm-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 almost never missing, to assign a geographic location to a contractor for a given year. We then use an official zipcode-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.

Matched Sample Descriptive Statistics: We were able to match 13,662 firms between 2000:4 and 2020:3. We then cleaned the sample by removing matches that only appear after 2020 (the COVID year) and before 2006, as data in the FPDS appears more complete and stable from 2006 onwards. The data cleaning process also 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; and (v) removing firms with fewer than one employee on average.

The resulting dataset is an unbalanced panel with $N = 5349$ firms observed from 2006:1 to 2019:4, $T = 56$. Figure 12 presents a screenshot from the log file of the `xtset` and `xtdescribe` commands.

Figure 12: Matched (Polished) Sample

```
96. xtset ein_masked_id qdate
        panel variable:  ein_masked_id (unbalanced)
         time variable:  qdate, 2006q1 to 2019q4
                 delta:  1 quarter

97. xtdescribe

  ein_masked_id:  3, 5, ..., 13660                        n =       5349
     qdate:  2006q1, 2006q2, ..., 2019q4                  T =         56
             Delta(qdate) = 1 quarter
             Span(qdate)  = 56 periods
             (ein_masked_id*qdate uniquely identifies each observation)

  Distribution of T_i:   min      5%     25%       50%       75%     95%      max
                          13      26      49        56        56      56       56

       Freq.  Percent    Cum. |  Pattern
  ----------------------------+-----------------------------------------------------
        3576    66.85   66.85 |  11111111111111111111111111111111111111111111111111111111
> 1
          82     1.53   68.39 |  ..11111111111111111111111111111111111111111111111111111
> 1
          49     0.92   69.30 |  .111111111111111111111111111111111111111111111111111111
> 1
          45     0.84   70.14 |  1111111111111111111111111111111111.....................
> .
          40     0.75   70.89 |  1111111111111111111111111111111111111111111111111.....
> .
          40     0.75   71.64 |  1111111111111111111111111111111111111111111111111111...
> .
          37     0.69   72.33 |  ....11111111111111111111111111111111111111111111111111
> 1
          34     0.64   72.97 |  1111111111111111111111111111111111111111..............
> .
          31     0.58   73.55 |  ..........11111111111111111111111111111111111111111111
> 1
        1415    26.45  100.00 |  (other patterns)
  ----------------------------+-----------------------------------------------------
```

Notes: The EIN is masked using a simple `egen masked_ein_id = group(ein)` command for security reasons: the BLS did not clear the output of the log file when EINs were not masked.

Next, we calculate the distribution of unpredictable contract sizes in our final sample. Figure 13 shows the output of the `sum, d` command in Stata, taken from the log file.

Figure 13: In-Sample Distribution of Contracts/Shocks Size

```
100 * Analyze shock distribution of matched firms:
101 sum EPSit if EPSit>0, d

                            (sum) EPSit
-------------------------------------------------------------
      Percentiles      Smallest
 1%        2263.2            .01
 5%          4660          37.32
10%          7416             49       Obs                10,843
25%         25000           49.9       Sum of Wgt.        10,843

50%        114900                      Mean              691194.9
                        Largest        Std. Dev.          2728904
75%      399578.1       6.10e+07
90%       1293665       6.19e+07       Variance          7.45e+12
95%       2894197       6.33e+07       Skewness          15.44963
99%      1.03e+07       1.14e+08       Kurtosis          410.5306
```

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. (2023).

Lastly, we split the sample into small and large firms. Small firms are defined as those whose average number of employees recorded before receiving their first contract (pre-treatment level of employment) is fewer than 150 employees. The size distribution of firms in the sample is summarized in Figure 14, which also includes a screenshot from the log file showing the output of the `sum, d` command in Stata.

Figure 14: In-Sample Distribution of Number of Employees - Small Firms

```
180         replace large_firms_oecd = 1 if emp_pre_first_contract > 150
  (175 real changes made)

181         sum emp_pre_first_contract if large_firms_oecd == 0, d

                    emp_pre_first_contract
-------------------------------------------------------------
      Percentiles      Smallest
 1%             1              1
 5%             2              1
10%             3              1       Obs               5,142
25%      5.666667              1       Sum of Wgt.       5,142

50%            13                      Mean           21.63049
                        Largest         Std. Dev.      24.66049
75%            27       148.3333
90%            53       148.3333       Variance       608.1399
95%            76       148.6667       Skewness       2.234772
99%           120       148.6667       Kurtosis       8.522913

182
```

Firms in the first quartile have a pre-treatment number of employees between 1 and 5.7; the median small firm in our sample has 13 employees, while firms in the third quartile have between 13 and 27 employees. The upper quartile is much more dispersed, with firms ranging from 27 to 148.7 employees.

C.1 External Validity

In this section, we show how firms in our sample compare with the rest of the US firms. To address the dimensionality problem, we consider a subsample of QCEW data covering the years from 2010 to 2015. Firms operating in multiple states are treated as separate units, with the firm identifier defined as a combination of EIN and state. Firms specialize in a single two-digit NAICS code, which is considered the sector of the firm. Therefore, a firm is characterized by a unique EIN-state pair and a unique major sector.

We aggregate the data by years and firms to construct an unbalanced panel. To further reduce dimensionality, variables are averaged across years (2010 through 2015), creating a cross-sectional database of all firms in the US located in one of our 43 signatory states. We create a dummy variable to indicate whether a firm is part of the sample of contractors analyzed in this paper. Specifically, the dummy D takes a value of one if a firm is in the sample.

We then estimate the following equation:

$$Y_{i,\ell,s} = \lambda_\ell + \alpha_s + \beta \cdot D_{i,\ell,s}^{\text{In the Sample}} + u_{i,\ell,s}$$

where $Y_{i,\ell,s}$ is an outcome variable, such as employment, for firm i, located in state ℓ, operating in sector s (NAICS code). $D_{i,\ell,s}^{\text{In the Sample}}$ is a dummy that takes a value of one when a firm belongs to the sample of firms we analyze. The regression is run separately for small firms (those with an average employment of fewer than 150 employees) and large firms.

The coefficient of interest, β, represents the average effect of being a firm in our sample, controlling for state and sector fixed effects. OLS estimates of β are presented in Table 1.

The results suggest that, on average, a small firm in our sample has 32 more employees than a firm outside of our sample. However, small firms in the sample do not appear to pay statistically different wages per employee, suggesting no significant differences in productivity levels. On the other hand, large firms in the sample do not show statistically significant differences compared to their out-of-sample counterparts.

Table 1: External Validity

Outcome	Employment		Total Wages		Average Wage	
Firm Size	Small	Large	Small	Large	Small	Large
(In the Sample)	32.19^{***}	-47.55	$468,375.6^{***}$	$-1,263,486$	645.43	-484.41
	(0.36)	(118.99)	$(6,938.2)$	$(1,711,001)$	$(1,586.65)$	(345.36)
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
N	7,814,549	371,935	7,814,549	371,935	7,814,549	371,935

Notes: Standard errors are reported in parentheses below the OLS estimates. *** denotes 99% significance level. Small firms have an average number of employees below 150. Data are sample averages for years 2010 through 2015. Firms in the sample represent the universe of US private firms with addresses in one of the 43 signatory states we have access to.

D One-Time Shock

If unpredictable contracts are either (i) highly serially correlated or (ii) highly correlated with potentially predictable contracts, our results could be confounded by these two channels, leading to an overestimation of the multipliers. Therefore, we conduct a test to demonstrate that unpredictable contracts function as a one-time shock.

First, the total quarterly value of contracts for each recipient firm i at time t, denoted by $G_{i,t}$, can be decomposed into two components:

$$G_{i,t} = \underbrace{\tilde{G}_{i,t}}_{\text{Potentially Predictable}} + \underbrace{\varepsilon_{i,t}}_{\text{Unpredictable}}$$

where $\tilde{G}_{i,t}$ is the quarterly value of potentially predictable contracts, and $\varepsilon_{i,t}$ represents the remainder, i.e., the quarterly value of unpredictable contracts. We estimate the effect of $\varepsilon^G_{i,t}$ on total contracts $G_{i,t+h}$; in particular, we estimate the following equation:

$$G_{i,t+h} = \gamma^h_0 \cdot \varepsilon_{i,t} + \delta^h_0 \cdot \tilde{G}_{i,t} + \text{Lags} + \underbrace{\lambda^h_{\ell,t} + \alpha^h_i + \mu^h_{s,t}}_{\text{Fixed Effects}} + u_{i,t} \quad \text{for } h = -8, ..., 8$$

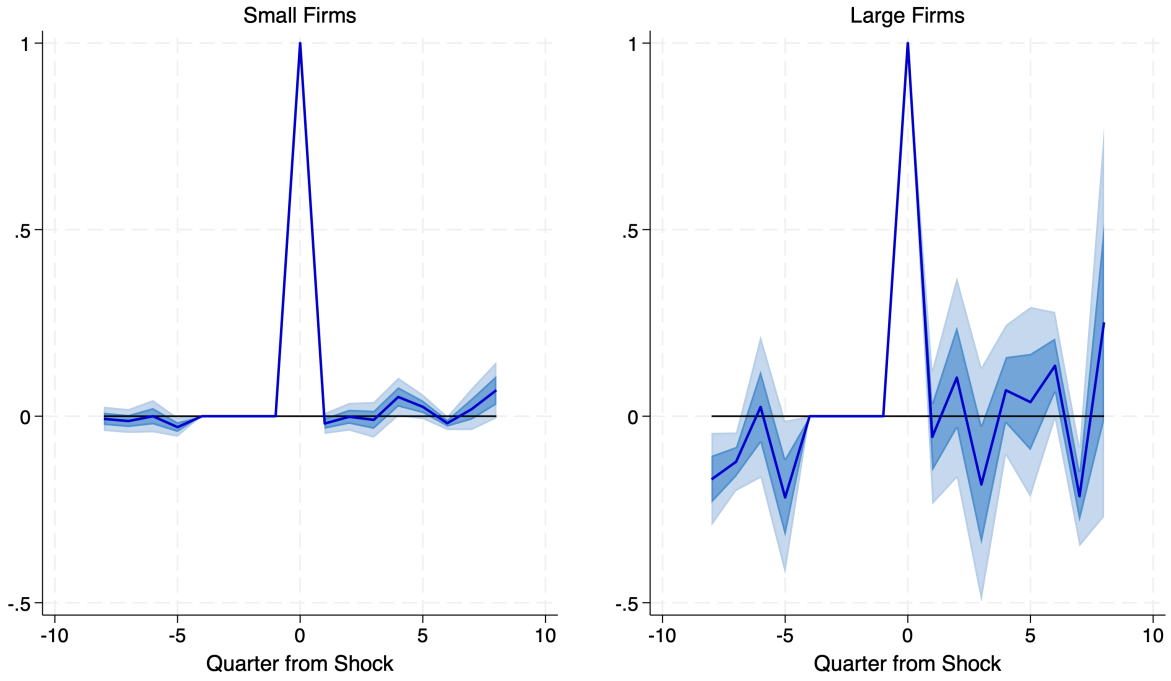

Figure 15: Mostly a One-Time Shock

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

where $\lambda_{\ell,t}^h$ is a location-time fixed effect (state-quarter), α_i^h is a firm fixed effect, and $\mu_{s,t}^h$ is an industry-time fixed effect (two-digit NAICS code-quarter). Lags include four lags of unanticipated contracts and four lags of potentially anticipated contracts. The coefficient γ^h measures the effect of winning an unpredictable contract at time t on the total value of contracts at time $t + h$. Figure 15 shows the OLS estimates of γ_0^h by horizon for both small and large firms.

There are two key observations from the figure. First, at horizon 0, the estimate of γ^0 is exactly one, meaning that winning an unpredictable contract does not systematically lead to winning other contemporaneous potentially predictable contracts. If this were the case, the estimate of γ_0^h would be greater than one.

Second, the estimates of γ_0^h for both negative and positive horizons are not statistically different from zero.[2] This indicates two things. First, the lack of an effect on past contract winnings suggests

[2]Results when $h = -1, ..., -4$ are flat because the specification already controls for four lags of the dependent

that winning a contract does not result from a history of winning more contracts in the previous two years, supporting the claim that these procurement contracts are unpredictable. Second, the null effect on future contracts shows that winning an unpredictable contract today does not make it more likely to win additional contracts in the near future (within the two-year horizon).

As the figure suggests, unpredictable contracts appear to act primarily as a one-time shock, which helps identify their impact on firms' outcomes.

variable ε^{G}_{t+h}.

E Baseline Results: Robustness

In this section of the Online Appendix we report robustness checks as well as extra results of the firm-level analysis.

E.1 Anticipation and Placebo Test by Quartile of Small Firms

Figure 16: Employment - Quartile 1 (Small Firms) - Anticipation and Placebo Tests

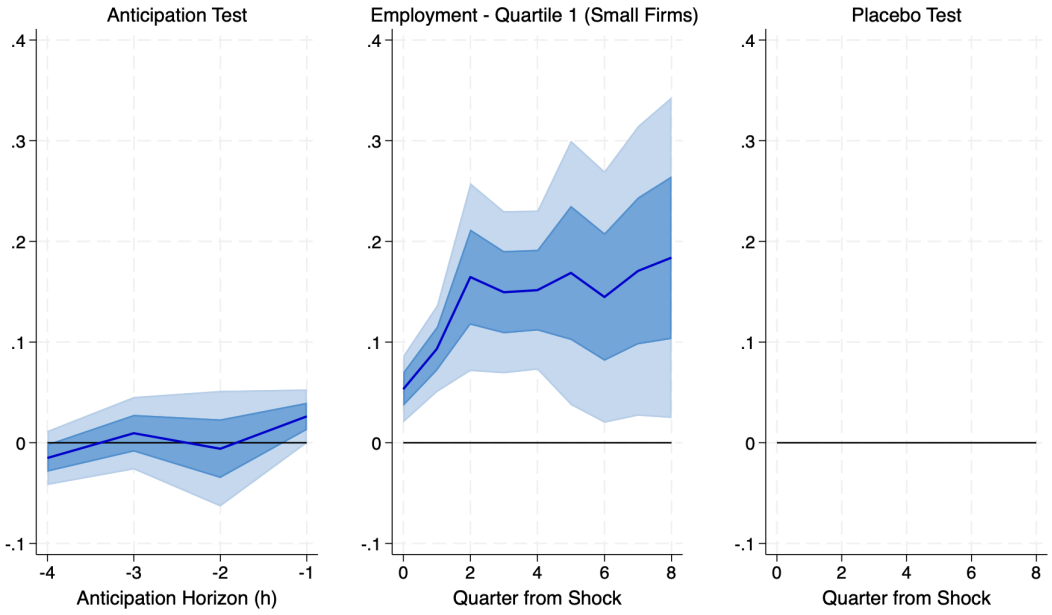

AWAITING BLS CLEARANCE OF PLACEBO TEST

Figure 17: Employment - Quartile 2 (Small Firms) - Anticipation and Placebo Tests

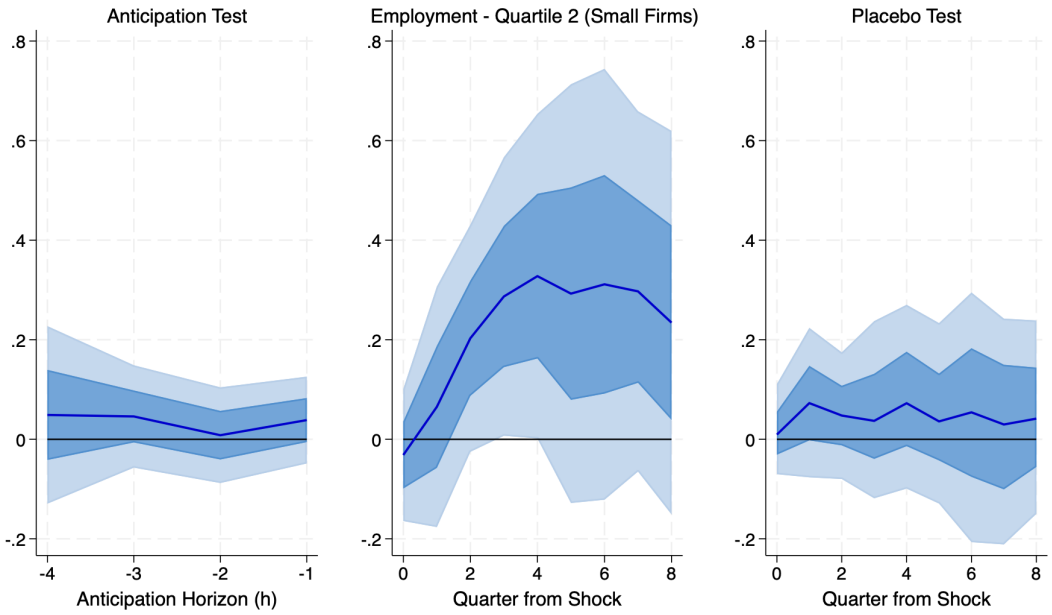

Figure 18: Employment - Quartile 3 (Small Firms) - Anticipation and Placebo Tests

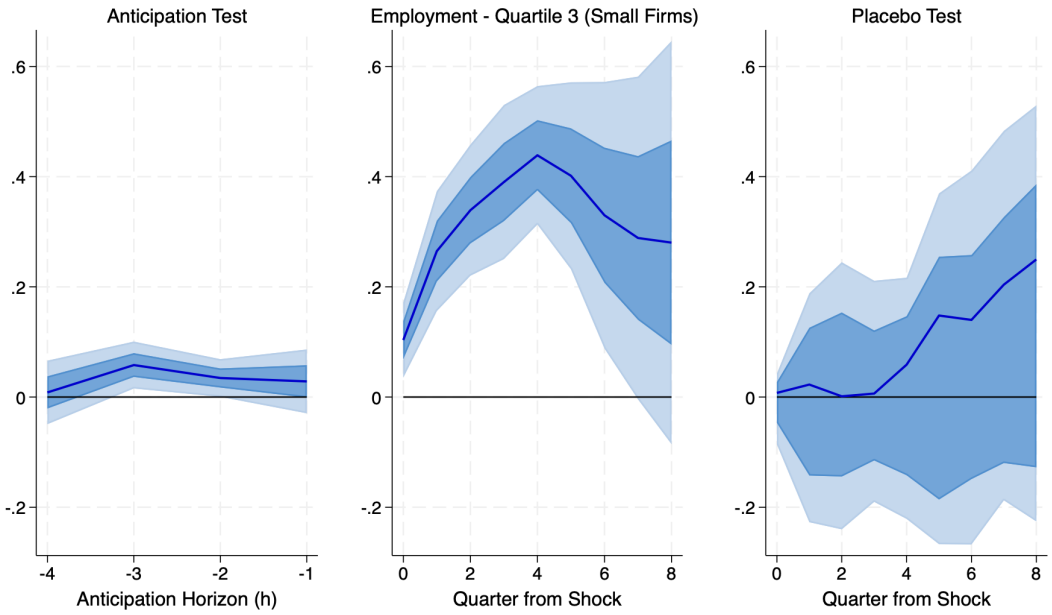

Figure 19: Employment - Quartile 4 (Small Firms) - Anticipation and Placebo Tests

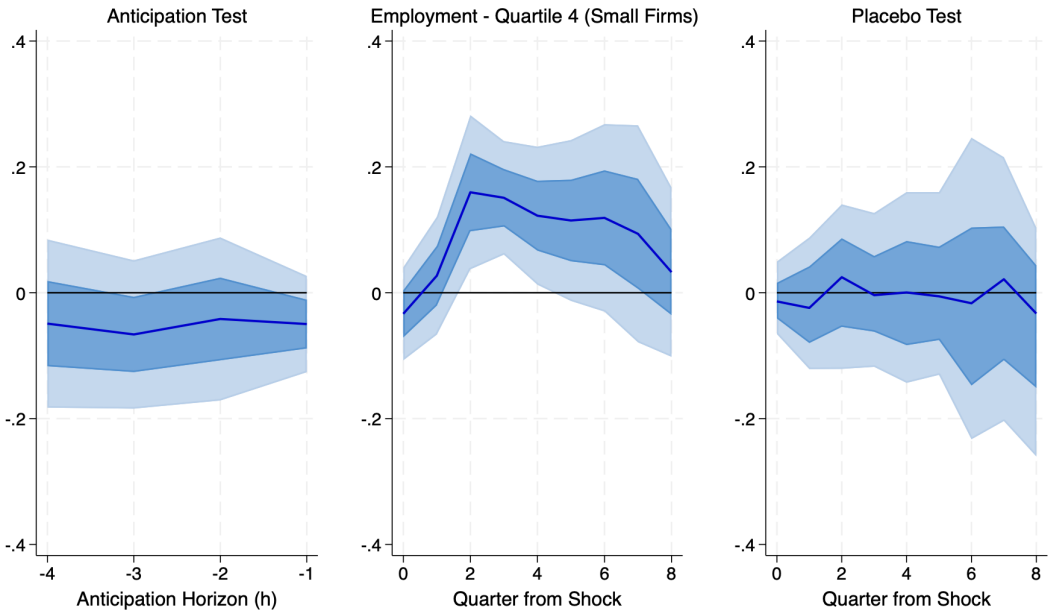

E.2 Controlling for Lags of Wage-per-Worker

We augment our baseline specification with four lags of the wage-per-worker, constructed by dividing quarterly total wages by quarterly employment. This is done to control for time-varying firm-level productivity shocks which cannot be addressed using firm fixed effects.

Figure 20: Employment - Small Firms - Lags of Wage-per-Worker

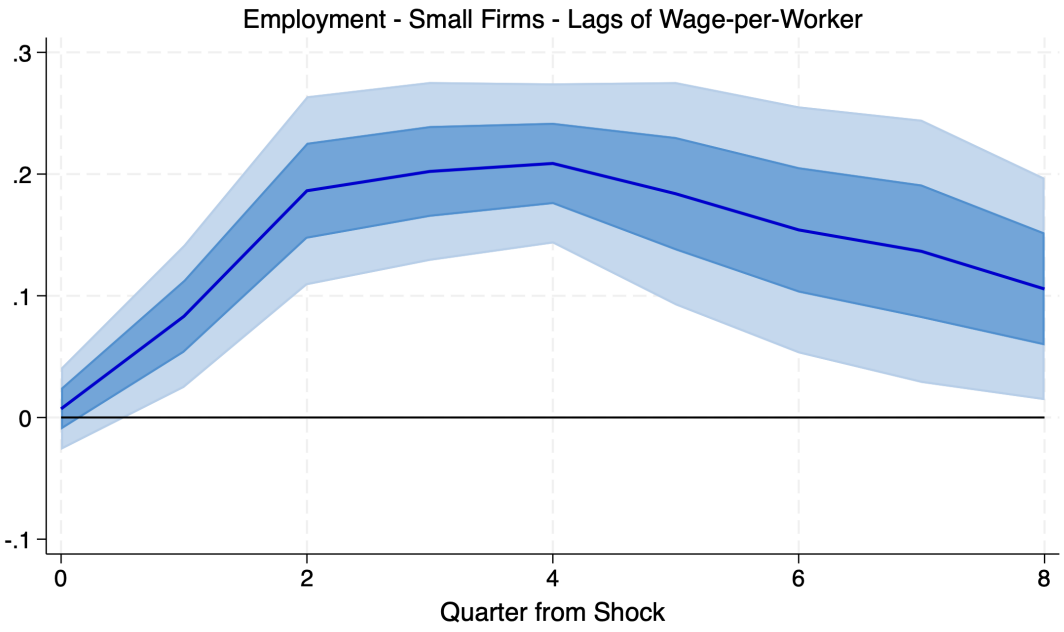

The positive effect of unanticipated contracts on small firms employment is unaffected by augmenting the specification by lags of wage-per-worker.

E.3 Response of Total and Average Wage

We also studied the response of the total and average wage.

Average Wage. The average wage is constructed as the (average) wage-per-worker: we divide the quarterly total wages by the quarterly employment values. We then estimated our baseline equation using as outcome variable the average wage. Results are reported in Figure 21. From the figure it appears that there is no effect on the average wage.

Figure 21: Average Wage - Small vs Large Firms

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

Total Wages. In light of no effect of unanticipated contracts on the average wage and the positive response on employment, we expect to see positive effects on total wages as well. Therefore, we estimate our baseline equation using total wages as outcome variable. Results are reported in the figure below:

Figure 22: Total Wage - Small vs Large Firms

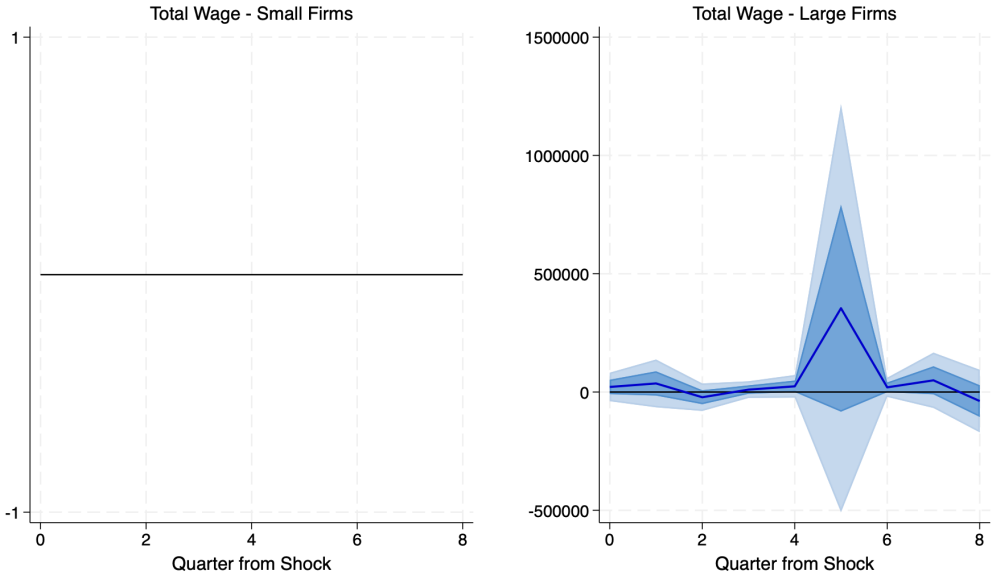

AWAITING BLS CLEARANCE OF TOT WAGE FOR SMALL FIRMS (NON-WINSORIZED)
Notes: Firms are observed from 2006:1 to 2019:4, i.e. $T = 56$. Number of small firms is $N = 5,142$, while the number of large firms is $N = 175$. Standard errors are clustered at the state level. Small bands are 68% confidence. Large bands are 95% confidence.

Notice that the results are severely affected by some outlier values. We did find some exceptionally large values in the distribution of total wages. Therefore, we Winsorize the results and drop firms which contain a value of total wages outside of the 1-99% range of the distribution. Then, we repeated the analysis by estimating the same regression. Results are reported in Figure 23.

Figure 23: Total Wage - Small vs Large Firms - Winsorized

Results are positive and significant for small firms only while for large firms the results are insignificant, in line with the employment results.

Lastly, we carry out placebo and anticipation tests also in this case, for small firms. Results are shown in Figure 24.

Figure 24: Total Wage - Small Firms - Anticipation and Placebo Test - Winsorized