Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry*

Kory Kroft[†] Yao Luo[‡] Magne Mogstad[§] Bradley Setzler[¶]

March 2020
Preliminary Draft - Do Not Circulate

1 Introduction

Both researchers and policymakers are keenly interested in quantifying the degree of imperfect competition in the U.S. economy and in understanding how it affects the outcomes and behavior of American workers and firms. It is argued that imperfect competition is especially salient in two types of markets. The first is the labor market where it is argued that firms exploit market power to mark-down the wages of workers relative to their productivity, with important implications for earnings inequality and the labor share of gross domestic product. The second is the product market, in which it is argued that firms use their market power to mark-up prices and increase profits, potentially at the cost of consumer welfare, investment, and innovation.

To draw inference about imperfect competition in these two types of markets, it is natural to try to measure the size of rents earned by employers and workers, defined as the excess return over that required to change a decision, as in Rosen (1986). However, these rents are not directly observed, and recovering them from data has proven difficult for several reasons. For example, observationally equivalent workers could be paid differently because of unobserved skill differences, not markdown of wages. Furthermore, profitability may vary across observationally equivalent firms because of unobserved productivity or hard to measure investments (e.g. in intangible capital), not markup of prices.

The primary contribution of our paper is to address these and other empirical challenges in order to draw inference about the amount of rents and imperfect competition in the American construction industry and about how market power affects the outcomes and behaviors of workers and firms in this industry. In contrast to most existing work on impefect competition, we consider the extent of and implications from market power in both the labor market and the product market.

As described in Section 2, our analyses are based on a matched employer-employee panel data set, which is constructed by combining the universe of U.S. business and worker tax records for the period 2001-2015.

*The opinions expressed in this paper are those of the authors alone and do not reflect the views of the Internal Revenue Service or the U.S. Treasury Department. This work is a component of a larger project on income risk in the United States, conducted through the SOI Joint Statistical Research Program. Kroft and Luo acknowledge funding from the SSHRC Insight Grant. Mogstad and Setzler acknowledge funding from NSF Grant SES-1851808.

[†]Department of Economics, University of Toronto; NBER. E-mail: kory.kroft@utoronto.ca.

[‡]Department of Economics, University of Toronto. E-mail: yao.luo@utoronto.ca.

[§]Department of Economics, University of Chicago; Statistics Norway; NBER; IFS. E-mail:magne.mogstad@gmail.com.

[¶]Becker Friedman Institute for Economics, University of Chicago. E-mail: bradley.setzler@gmail.com.

The firm data contain information on sales, profits and inputs as well as industry codes and geographical identifiers. The worker data give us information about the number of (new and incumbent) workers and their wage bill. We merge the employer-employee panel data set with a new data set on U.S. procurement auctions that we have constructed by submitting FOIA requests to state governments and by web-scraping state-specific bidding websites. The resulting data set covers billions of dollars in procurement contracts awarded to thousands of firms. Importantly, we observe the bid of each firm, not only the winner.

In Sections 3 and 4, we use the panel data set to estimate the effects of winning a procurement auction. As described in Section 3, we use a difference-in-differences (DiD) design that compares first-time winners of a procurement auction to the firms that lose the auction, before and after the auction. The inclusion of firm and year fixed effects helps mitigate the natural concern that firms that win are likely to differ from those that lose even in the absence of winning the auction. For example, in the standard model of optimal bidding in first-price auctions with private information on costs, firms bid monotonically in costs, so low-cost firms are more likely to win. This suggests there are likely to be differences in the composition of auction winners and losers, which motivates the DiD design.

Although the DiD design adjusts for differences across firms in levels, one could still be concerned that firms in the treated group may have different underlying trends as compared to firms in the control group. For example, firms with different cost functions may experience different changes in market conditions. We take several steps to make the firms even more comparable and, thus, more likely to satisfy the identifying assumption of common trends. One of these is to restrict the control group to firms that are close to winning the auction in a cardinal sense. These are firms that placed bids within a certain threshold of the winning bid in dollar value. Another is to restrict the control group to firms that are close to winning the auction in an ordinal sense. These are firms that placed bids lower than other bidders in the auction, but still higher than the winner's bid. Though stronger sample restrictions reduce sample size and thus lead to less precise estimates, it is reassuring to find that our results do not materially change across the various specifications.

The estimates from the DiD design are presented in Section 4. Winning a procurement auction – which corresponds to a $2.7 million procurement contract on average – increases sales by 17 percent, expenditure on intermediate inputs by 15 percent, and payment to labor by 10 percent. The total increase in sales is slightly larger than the payments recieved from the government procurement project, suggesting a small crowd-in effect of production in the private market. The 10 percent increase in the wage bill is due to an 8 percent increase in the number of employees and a 2 percent increase in earnings per employee. This finding is consistent with the firm bidding up wages to hire more workers, which would be the case if it faces an upward sloping labor supply curve and, therfore, has wage setting power in the labor market. One potential concern with this interpretation is that firms may engage in skill-upgrading. However, this explanation is at odds with the data. We find no evidence of changes in the quality of the workers as a result of winning a procurement auction. Furthermore, the estimated increase in earnings per worker does not change if we restrict the sample to workers that are employed by the same firm before and after it wins the auction.

Motivated and guided by these findings, we develop in Section 5 a model where multiple construction firms compete with one another for projects in the product market and for workers in the labor market. Importantly, we allow (but do not impose) that firms have market power in the labor market, or the product market, or both. The model serves several purposes. First, it offers an economic interpretion of the estimated treatment effects of winning a procurement auction. Without a model, it is hard to gauge the size of these treatment effects and difficult to understand the responses to the procurement-induced changes in product demand. Second, the model makes explicit the assumptions needed to recover the parameters that govern the behavior of firms and workers in the construction industry. This includes the firm-specific labor supply curve, the technology, and the market structure. Third, the model lets us quantify the incidence of the surplus in the construction industry in general and of the government spending on infrastructure projects in particular. Fourth, the model makes it possible to perform counterfactuals to explore how imperfect competition in the product market or the labor market affects the outcomes and behavior of both workers and firms in the construction industry.

The labor market side of the model builds on work by Rosen (1986), Boal and Ransom (1997), Bhaskar et al. (2002), Manning (2003), Card et al. (2018), and Lamadon et al. (2019). Competitive labor market theory requires firms to be wage takers so that labor supply to the individual firm is perfectly elastic. The evidence that winning a procurement auction causes the firm to bid up wages and hire more workers is at odds with this theory. To allow the firm-specific labor supply curve to be imperfectly elastic, we let workers have heterogeneous preferences over non-wage job characteristics or amenities when making firm choices. Since we allow these amenities to be unobserved to the analyst, they can include a wide range of characteristics, such as distance of the firm from the worker's home, flexibility in the work schedules, the type of tasks performed, the effort required to perform these tasks, the social environment in the workplace, and so on (Hamermesh, 1999; Pierce, 2001; Maestas et al., 2018; Mas and Pallais, 2017; Wiswall and Zafar, 2017). We assume that firms do not observe the idiosyncratic taste for amenities of any given worker. This information asymmetry implies that employers cannot price discriminate with respect to workers' reservation wages. Instead, if a firm faces higher demand for its products and wants to hire more labor, it needs to offer higher wages to all workers. As a result, the equilibrium allocation of workers to firms creates surplus or rents to inframarginal workers and, moreover, some or all the gains from a firm winning a procurement auction may fall on its workers.

The firm side of the model consists of two types of product markets in which the construction firms may choose to participate: projects in the private market and government projects which are procured through auctions. The firm's behavior is specified as a two-stage problem, which we solve backwards. In the first stage, firms submit a bid for a government project that is procured through a first-price sealed-bid auction. The project specifies the amount of output that must be produced within a given time frame. At the end of the first stage, the firm learns the auction outcome. If the firm wins the auction, it receives as revenue the winning bid amount and commences production. In the second stage, the firm chooses inputs to maximize profit from production in the private market, taken as given the outcome in the procurement auction. A firm may earn rents in the private market due to price-setting power and in the government projects because of a limited number of bidders in the auction. Production in both private and procurement projects occur simultaneously at the end of the second stage.

In Section 6, we take the model to the data. We prove identification of the model parameters of interest, before presenting and economically interpreting the parameter estimates. The identification argument forges a direct link between our model in Section 5 and the treatment effects analyses in Section 4. For example, the rents earned by workers can be measured by the elasticity of the labor supply curve to the firm. This elasticity can be recovered from the DiD estimates of the effect of winning a procurement auction on payments to labor compared to the effect on the number of workers. To identify the technology parameters, we use data on the firm's choice of intermediates and labor and, especially, the DiD estimates of how it changes these inputs in response to winning a procurement auction. The identification argument for the product demand curve facing the firm in the private market builds on Ackerberg et al. (2015), who shows the conditions under which one can use the input demand function to control for unobserved productivity across firms.

The estimates of the model parameters yield four key findings. First, firms have significant wage setting power with an estimate of the firms-specific labor supply elasticity of about 4.1. This indicates that, if an American construction firm aims to increase the number of employees by 10%, it can accomplish this by increasing wages by around 2.4%. Second, firms have price setting power in the private market with a product market elasticity of 7.3. This implies that, in order for a firm to increase output by 7.3% in the private market, it must reduce price by 1%. Third, workers in a typical firm earn around 25% of the total rents that accrue from winning a procurement auction. Workers gain, on average, about $1,000 at the typical firm from winning a procurement auction, with nearly all of these additional rents being captured by incumbent workers. By comparison, the typical firm captures about $2,900 per worker in additional rents from winning an auction. Fourth, the estimated return to scale (over capital and labor) is slightly above one, consistent with the findings of Levinsohn and Petrin (2003). It is is important to observe, however, that market power in the product and the labor market is sufficiently strong to offset the incentives to keep increasing production due to an increasing returns to scale. At the same time, the market power is not strong enough to create a crowd-out of private market production if the firm wins a procurement auction.

We use the estimated model to perform counterfactual analyses that allow us to infer how imperfect competition affects the outcomes and behavior of workers and firms in the construction industry. We present two sets of results. The first considers the consequences of increasing a firm's wage setting power in the labor market by rotating the labor supply curve it faces, whereas the second explores the impact of increasing a firm's price setting power in the private market by rotating the firm-specific product demand curve. Consider, for now, a counterfactual economy where the labor supply elasticity of a given firm is reduced by half. In this economy, the firm employs 15% fewer workers and decreases wages by 7%. Capital only decreases by 4%, indicating a shift toward capital-intensive production as the marginal cost of labor rises. While output is reduced by 12%, the firm's profit is 3% higher because it takes advantage of its additional market power in the labor market to increasingly markdown wages. Taken together, these results highlight the scope for increased labor market power to affect the wages that firms pay and the profits that firm accrue.

By way of comparison, the incidence of a procurement auction is far less sensitive to changes in the firm-specific labor supply curve and, as a consequence, these results vary relatively little between the actual and counterfactual economies. The reason is that a steeper labor curve creates two opposing forces for how a firm responds to winning the procurement auction. On the one hand, fewer workers benefit from the procurement auction because the firm chooses to use less labor and cut back production in the private market. On the other hand, the workers that do benefit receive a large increase in their wages, as the firm must bid up wages more to achieve a given increase in employment. Overall, these two effects nearly offset each other, so that the percent increases in the rents of workers and firms due to a procurement win do not depend strongly on the wage setting power of the firm.

This paper contributes to several existing literatures. It contributes to the literature in labor economics that estimates firm-specific labor supply elasticities and uses them to infer labor market power and rents (Berger et al., 2019; Garin and Silverio, 2019; Howell and Brown, 2020; Kline et al., 2019; Lamadon et al., 2019). It differs in several respects. First, it differs in the empirical context and source of exogenous variation used to identify labor market power by comparing construction firms that win and lose procurement auctions. Procurement auctions have previously been studied as a source of exogenous variation in other contexts (Ferraz et al., 2015; Anagol and Fujiwara, 2016), while we leverage procurement auctions to study rents. Second, it differs in that it accounts for imperfect competition in the product market by allowing firms to mark-up output prices when measuring the share of rents captured by firms, which in turn affects the estimated worker share of rents. Third, it differs from the existing literature in that it studies rents in counterfactual economies in which firms have more labor market power or more product market power. This counterfactual analysis requires modeling and estimating the firm technology, product market structure, and other determinants of rents.

This paper also makes several novel contributions to the empirical auction literature. To our knowledge, this is the first empirical paper that estimates optimal bidding as a function of unobserved productivity rather than unobserved costs. By modeling bidding as a function of productivity, we allow for a relationship between the probability of winning the auction and other firm outcomes that depend on productivity, such as employment and output (Foster et al., 2008; Ackerberg et al., 2015). We are able to make this contribution because we observe rich data on the non-auction activities of firms, whereas the data typically available in the empirical auction literature only contains information on auction-related activities (e.g., Guerre et al. 2000). To our knowledge, this is the first empirical auction paper that accounts for heterogeneous outside option value using data on the private market activity of firms that lose auctions. The auction literature almost always normalizes the outside option value to be zero (Athey and Haile, 2007) or models the outside option as (dynamic) participation in future auctions (Jofre-Bonet and Pesendorfer, 2003). To our knowledge, this is also the first empirical auction paper that accounts for the interaction between the auction market, the private market, and the labor market. Even among winners of procurement contracts, most economic activity is in the private market, so it is important to observe private market activity to understand the effects of winning an auction on firms.

The remainder of the paper proceeds as follows. Section 2 describes the institutional context, construction of the procurement auction data, and how we linked auction records to tax records. Section 3 assesses the relevance and validity of our instrument. Section 4 estimates the causal effect of being awarded a procurement auction versus denial on the outcomes of interest, providing a number of robustness checks. Sections 5 and 6 develop, identify, and estimate a model where multiple construction firms compete with one another for projects in the product market and for workers in the labor market, accounting for equilibrium bidding in government procurement auctions. Section 7 uses the estimated model to perform counterfactual analyses that allow us to infer how imperfect competition affects the outcomes and behavior of workers and firms in the construction industry. The final section concludes.

2 Data

Our empirical analyses are based on a matched employer-employee panel data set with information on the characteristics and outcomes of U.S. workers and firms. The employer-employee data covers the years 2001-2015. The data set is constructed by first linking U.S. Treasury corporate tax returns to worker-level tax returns, and then merging this linked data set with procurement auction records. The tax returns cover nearly all firms and workers in the private sector, whereas the procurement auction records cover hundreds of thousands of auctions in 46 states. Below, we briefly describe data sources, sample selection, and key variables. Additional details on the sample and variable definitions are provided in Appendix A.

The corporate tax returns include balance sheet and other information from Forms 1120 (C-corporations), 1120S (S-corporations), and 1065 (partnerships). We then link the corporate tax returns to worker-level W-2 (direct employee) tax returns and 1099 (independent contractor) tax returns, defining the highest-paying firm in a given year as the worker's primary employer. Our baseline set of workers consists of prime-aged W-2 employees with annual earnings from the primary employer greater than the annualized full-time minimum

	Sample Size		Share of the Construction Sector
Number of Firms	7,876		0.9%
Workers per Firm	46		11.7%
	Value Per Firm ($ millions)	Mean of the Log	Share of the Construction Sector (%)
Sales	19.927	15.061	12.1%
EBITD	9.159	14.075	9.6%
Intermediate Costs	14.661	14.719	12.4%
Wage bill	2.737	13.549	13.4%

Table 1: Sample Characteristics

Notes: This table displays 2010 descriptive statistics for the sample of firms that place bids in that year.

wage in the year. Because firms sometimes use part-time workers or contracted labor, we also consider a broader meaure of the workforce that includes any worker to whom the firm reports payments on a W-2 or 1099 tax record.

The key variables that we draw on from the corporate tax returns are sales, expenditures on intermediate inputs, and the NAICS industry code. We define two measures of profits. The first measure is sales minus expenditures on intermediate inputs minus wage bill (hereafter, "sales net of expenditures"). The second measure of profits is earnings before interest, taxes, and depreciation (hereafter, "EBITD"), which we construct following Kline et al. (2019). The key variables we draw from the worker tax returns are the number of employees and their wage bill for the primary sample of workers. We also consider the number of employees and wage bill for the broader sample that includes part-time workers and independent contractors. Using the panel structure of the employer-employee data, we define three measures of mean earnings: mean earnings among all workers; mean earnings among stayers, which we define as workers employed at the bidding firm consistently from 2 years prior to the procurement auction until 2 years after; and the past earnings of new hires at their previous firm, which we define as the mean earnings at $t-1$ of workers who become primarily employed by a new firm at t.

We obtain the new data set on procurement auctions in three ways: webscraping from BidX[1], a company that facilitates the bidding process online for a number of states; webscraping from state-specific bidding websites; and obtaining records directly by submitting FOIA requests to state governments. We recover Department of Transportation records from 46 states during 2001-2015. An observation in this data set is at the level of auction-firm and the variables are the firm's name and location as well as the firm's bid. Importantly, we observe the bid of each firm for a given auction, not only the winner. The resulting data set covers billions of dollars in procurement contracts awarded to thousands of firms through hundreds of thousands of auctions.

To merge the auction data to the tax records, we use a fuzzy matching approach based on the firm's name and location. Five states provided not only the name and address but also the federal employer identification number (EIN) of the firm, allowing us to perform an exact match. We train the algorithm on these five

[1]See https://www.bidx.com/.

states before applying it to the other 41 states. Online Appendix D provides the details on how we obtained the auction records, trained and validated the linking algorithm, and merged the auction records to the tax returns. Furthermore, this appendix includes several robustness checks to the matching procedure. It shows that our procedure performs well in an out-of-sample validation on publicly available pension data tax filings. Furthermore, it shows that our main results do not change materially if we restrict the sample to the the five states where we can an perform exact matches.

Table 1 displays the sample sizes of firms and workers that participate in auctions in 2010. In total, our sample includes almost 8,000 unique firms that generate over $150 billion in annual revenues and employ about 360,000 full-time workers. Nearly all the firms are recorded as being in the construction industry (i.e. the firms have NAICS codes beginning with 23). As a share of the national construction sector (as recorded in the tax records in 2010), our sample of 8,000 firms accounts for 12% of sales, 12% of employment, 10% of EBITD, 12% of intermediate costs, and 13% of wage payments. The state-specific sample size and share of the local economy represented by auction participants linked to tax records is displayed in Online Appendix Table A.5. California, Michigan, and Texas are the states with the most bidding firms, while Iowa, Montanta, and North Dakota are the states for which bidders employ the greatest share of the construction sector.

3 Institutional Setting and Research Design

3.1 Procurement Auctions in the US

The procurement contracts studied in this paper are administered by the Department of Transportation (DOT) in each of 46 states. Our data show that these DOTs allocated $383 billion through 155,768 distinct auctions involving 16,697 unique contractors in 2010 alone. The procurements broadly involve the construction and landscaping of local roads, bridges, and highways. The DOTs are responsible for determining the nature of the project, including the blueprints, detailed list of tasks to be performed or items to be constructed, quality guidelines and standards, and expected or required time to completion. This information is publicly available in the solicitation for bidders posted by the DOT. The auctions are administered through a standard first-price sealed-bid auction, in which a firm submits a bid without observing the bids of other firms or which other firms are bidding.

The awarding of a contract has two steps. The first step is qualification. In order to submit a bid, a contractor must be pre-qualified by the DOT to ensure sufficient experience, equipment, and competence to carry out the tasks involved.[2] Once approved, the contractor is awarded a license to bid. The second step is the auction. In the first-price sealed-bid auction, a qualified firm submits a bid without observing the bidding behavior of other firms, and the contract is awarded to the firm with the lowest bid. Importantly, we observe the bids of every participant in the auction, not only the winner.

3.2 Research Design

To estimate the effects of winning a procurement auction, we use a difference-in-differences (DiD) research design. The idea is to compare first-time winners of a procurement auction to the firms that lose the auction, before and after the auction.

[2] The DOT may choose to further restrict bidder participation. For example, in California, some projects are labeled a "business development initiative," in which case, only local small or disadvantaged contractors are permitted to participate.

To be concrete, consider an auction that occurs in year c ("cohort") and let t denote the number of years since the auction occurred. For notational convenience, we omit the firm subscript on all variables. Denote a firm's observed outcome by $Y_{c,t}$. Let $D_c = 1$ if the firm wins its first auction at c, and $D_c = 0$ if it bids in an auction at c but loses.[3] Let $Y_{c,t}(1)$ and $Y_{c,t}(0)$ represent the realization of $Y_{c,t}$ that would have been experienced by the firm had its win status been exogenously set to 1 or 0. The relationship between observed and potential outcomes is given by

$$Y_{c,t} = D_c Y_{c,t}(1) + (1 - D_c)\, Y_{c,t}(0)$$

The parameter of interest is $\mathbb{E}\left[Y_{c,t}(1) - Y_{c,t}(0) \middle| D_c = 1\right]$, which is the cohort-specific average treatment effect on the treated (ATT), t years after the auction among firms winning their first auction in cohort c. In the empirical analysis, we average the estimates of these cohort-specific ATTs across cohorts.

The key identification challenge is to recover the average potential outcome among auction winners if they had lost, $\mathbb{E}\left[Y_{c,t}(0) \middle| D_c = 1\right]$. A natural control group for the treated firms that won their first auction at c is the set of firms that also had never won an auction before c and placed a bid at c but lost. Let $X_c = 1$ indicate firms that belong to this control group. Given this control group, one possibility is to use a matching estimator that infers the average potential outcome of the winners if they had lost the auction, $\mathbb{E}\left[Y_{c,t}(0) \middle| D_c = 1\right]$, from the observed outcome of the control group that did lose the auction, $\mathbb{E}\left[Y_{c,t}(0) \middle| D_c = 0, X_c = 1\right]$. However, matching is unlikely to perform well in our setting. In the standard model of optimal bidding in first-price auctions with private information on costs, firms bid monotonically in costs, so low-cost firms are more likely to win. This implies there are likely to be differences in the composition of auction winners and losers, even after conditioning on $X_c = 1$.

To account for such differences, we consider a cohort-specific DiD estimator of the form,

$$\mathbb{E}\left[Y_{c,t} - Y_{c,s} \middle| D_j = 1\right] - \mathbb{E}\left[Y_{c,t} - Y_{c,s} \middle| D_j = 0, X_c = 1\right] \tag{1}$$

for a given pre-period $s < 0$. The data must satisfy two conditions in order for the DiD to recover the ATT. The first condition is *parallel trends*, $\mathbb{E}\left[Y_{c,t}(0) - Y_{c,s}(0) \middle| D_j = 1\right] = \mathbb{E}\left[Y_{t,s}(0) - Y_{c,s}(0) \middle| D_c = 0, X_c = 1\right]$. The second condition is *no anticipation*, $\mathbb{E}\left[Y_{c,s}(1) \middle| D_j = 1\right] = \mathbb{E}\left[Y_{c,s}(0) \middle| D_j = 1\right]$. Under parallel trends and no anticipation, it is straightforward to show that the DiD estimator in equation (1) recovers the ATT. Intuitively, compositional differences between the winners and losers (such as differences in costs) are differenced out by comparing changes over time for the treated and control groups.

Although the DiD estimator adjusts for differences across firms in levels, one could still be concerned that firms in the treated group may have different underlying trends as compared to firms in the control group. For example, firms with different cost functions may experience different changes in market conditions. To make the firms even more comparable and, thus, more likely to satisfy the parallel trends condition, it may be useful to place stronger restrictions on X_c. In our empirical analysis, we make several such restrictions. One of these is to restrict the control group to firms that are close to winning the auction in a cardinal sense. These are firms that placed bids within a certain threshold of the winning bid in dollar value. Another is to restrict the control group to firms that are close to winning the auction in an ordinal sense. These are firms

[3] To account for left-censoring, we do not define a win as a "first win" unless there were at least two observed years of data during which the firm could have won and did not win an auction. For example, if a state provided auction records for 2001-2015, and a firm is first observed winning in 2001 or 2002, we do not consider this firm a first-time winner, but if the firm is first observed winning in 2003 or later, we consider it a first time winner. The results do not materially change if we use all the years of the data.

that placed bids lower than other bidders in the auction, but still higher than the winner's bid. Though stronger sample restrictions reduce sample size and thus lead to less precise estimates, it is reassuring to find that the main estimates do not materially change across the various specifications.

Another potential concern is that the winners might anticipate that they are relatively likely to win an upcoming auction and change their behavior even prior to the outcome of the procurement auction. To investigate this concern, we directly assess the pre-trends. If such anticipation occurs, we can change the before and after contrast to avoid the periods at which anticipation is likely. Our data indicate that winning firms may be adjusting their behavior in the year just before the outcome of the auction. If we were to use $s = -1$ as the omitted relative time, this anticipatory behavior may create bias in the DiD estimates. However, there is no evidence of anticipatory behavior in earlier time periods. Thus, our baseline DiD specification contrasts the outcomes in the post-treatment periods to those in the pre-treatment periods $s < -1$.

A final possible concern is that firm composition may change over time due to differential firm survival between winning and losing firms. We investigate this explicitly by defining a firm death indicator and estimating survival probabilities for the treated and control group, finding a relatively precisely estimated zero effect on differential survival.

3.3 Graphical Evidence

Before presenting our main results, we provide graphical evidence on the effects of winning an auction for the first time. We consider as the treated group the firms that win their first procurement auction in year c. The control group is the set of firms that had not won an auction before c and placed a bid at c but lost. Letting j denote a firm and considering each relative time $t = -4, ..., 4$, we consider the regression,

$$Y_{j,c,t} = \underbrace{\sum_{t' \neq s} 1\{t' = t\}\, \mu_{t'}}_{\textit{relative time fixed effect}} + \underbrace{\sum_{j'} 1\{j' = j\}\, \psi_{j'}}_{\textit{firm fixed effect}} + \underbrace{\sum_{t' \neq s} 1\{t' = t\}\, D_{j,c}\tau_{t,c}}_{\textit{treatment status by relative time}} + \underbrace{\epsilon_{j,t}}_{\textit{residual}} \tag{2}$$

where the empirical counterpart to $\tau_{t,c}$ is the DiD estimand defined in equation (1). We use the regression implementation to make it easier to include additional covariates and calculate the standard errors (which are clustered at the firm level j to account for serial correlation).[4]

Appendix Figure A.1 presents estimates from equation (2) for two outcomes from the procurement auctions: Subfigure (a) plots the share of firms that are first-time winners of a procurement auction, and subfigure (b) plots the share of firms that win a procurement auction in the relative year. Mechanically, both treated and control units have no wins prior to $t = 0$, so the effect is zero on $t < 0$ for both subfigures. At $t = 0$, the treated group wins a contract and the control group bids for a contract but loses, so the treatment effect is mechanically one for both subfigures. The mean winnings for first-time winners at $t = 0$ is \$2.7 million. On $t > 0$ in subfigure (a), we see that some control units win auctions, with around 15% of control units winning their first auction at $t = 1$ and around 5% at $t = 4$. This means the losers continue to bid and partially catch up to the winners. However, as shown in subfigure (b), treated units are more likely than control units to win any procurement auction on $t > 0$. Treated firms are around 21% more likely to win at least one auction at $t = 1$ and 14% more likely at $t = 4$.

[4] To estimate $\tau_{t,c}$ for all c and t, we stack all cohorts. We then average $\tau_{t,c}$ across c to get a mean impact per relative time period, $\bar{\tau}_t$. We use the delta method to compute the standard error of the $\bar{\tau}_t$ for each t.

In Appendix Figure A.2, we plot pre-trends and post-trends at annual frequency from the estimates of the treatment effect for 12 outcomes of interest. Four patterns stand out. First, in the pre-treatment relative times -4 to -2, there is no evidence of differential trends between the winners and losers. This is consistent with the contract winners at relative time 0 being similar before they begin differentially winning procurement contracts. Second, at relative time -1, there is suggestive evidence of winners changing behavior as compared to losers, though the estimates are only a small fraction of the effects at relative times 0 and onward.[5] This is reassuring given our difference-in-differences strategy. If we were to use $s = -1$ as the omitted relative time, this anticipatory behavior may create bias in the DiD estimates. However, there is no evidence of anticipatory behavior in earlier time periods. Thus, our baseline DiD specification contrasts the outcomes in the post-treatment periods to those in the pre-treatment periods $s < -1$. Third, at relative times 0 to 2, the contract has been awarded and the outcomes of treatment firms jump in an economically and statistically significant manner relative to control firms. At relative times 3 to 4, the differences show some evidence of fading out, as the control group begins to catch up to the treatment group, though the difference is economically and statistically significant.

Based on the patterns observed at annual frequency, our main estimates of the effects of winning a procurement contract on winners relative to losers will classify relative times $\{-4, -3, -2\}$ as the pre-treatment period ("Before"), and relative times $\{0, 1, 2\}$ as the post-treatment period ("After"). One sometimes sees empirical studies that restrict the behavior of the control group in the post-period by requiring that the control group remains untreated long after the event. Though such a restriction helps to clarify the counterfactual, it risks biasing the estimate by conditioning on an endogenous outcome. As a robustness check, we nevertheless restrict the control group to firms that do not win an auction during the "After" interval. As evidenced by Appendix Figure A.6, the main patterns are unchanged, though the point estimates become slightly larger.

Furthermore, one may worry that our results depend strongly on our choice of "Before" and "After" time intervals. As a robustness check, Appendix Figure A.7 considers including all of the pre-periods $\{-4, -3, -2, -1\}$ in the "Before" interval and all of the post-periods $\{0, 1, 2, 3, 4\}$ in the "After" interval, finding that the estimates are very similar.

4 Effects of Winning Procurement Auctions

4.1 Main Estimates

Balance Sheet Outcomes

Figures 1(a)-1(e) present the main estimates of the effects of auction winnings on five outcomes constructed from the firm balance sheet data, where "Before" refers to relative times $\{-4, -3, -2\}$ and "After" refers to relative time $\{0, 1, 2\}$. The "Before" effects are small in magnitude and are not statistically different from zero (all p-values above 0.10). This is consistent with the graphical evidence presented in the previous section indicating no differential pre-trends between auction winners and losers on the balance sheet outcomes, among firms bidding at the same time.

The "After" effects in Figure 1 provide our main results on balance sheet outcomes. We find the largest effect on sales, which increase by 17% for winners relative to losers (p-value below 0.01). Removing the

[5] The evidence of a partial effect at relative time -1 may be driven in part by the fact that fiscal years of firms can differ from calendar years, and relative time is defined based on calendar years. For example, a firm may report its activity from the second half of relative time -1 and the first half of event year 0 as the calendar year corresponding to relative time -1.

procurement contract value from the measure of sales allows us to study private market activity. A negative estimate on sales net of procurements would indicate that public expenditure "crowds-out" revenues from the private market, but we find no evidence of crowd-out. Instead, the effect on sales net of procurements remains positive at 6% (p-value below 0.10), suggesting a small crowd-in effect. In Section 6, our economic model shows how the evidence of crowd-in can be informative about firm technology and market power.

Next, we estimate the effects on two distinct measures of profits that accrue to firms: sales net of expenses (on intermediate inputs and labor), and EBITD (earnings before interest, taxes, and depreciation). Both measures show that profits increase by over 14% (p-values around 0.01). In Section 6, our economic model will use this estimate to assess the economic incidence of winning a procurement contract on firms.

Lastly, we estimate the effect of winning a procurement contract on intermediate expenditures, finding an increase of 15% (p-value below 0.01). In Section 6, our economic model will clarify that the ratio between the sales response and the intermediate expenditures response is proportional to the degree of market power the firm commands in the local product market. Market power diminishes as these two effects become closer, and their difference (17% versus 15%) suggests that market power in the local product market is modest but non-zero for the construction sector.

Employment and Earnings Outcomes

Figures 1(f)-1(l) present the main estimates of the effects of auction winnings on seven variables constructed from the worker earnings data. The "Before" effects are small in magnitude and are not statistically different from zero (all p-values above 0.10). This is consistent with the graphical evidence presented in the previous section indicating no differential pre-trends between auction winners and losers on the employment and earnings outcomes, among firms bidding at the same time.

The "After" effects in Figure 1 provide our main results on employment and earnings outcomes. We find that the effect of winning a procurement auction is a 10 percent increase in the wage bill, an 8 percent increase in the number of employees, and a 2 percent increase in earnings per employee (all p-values below 0.01). We also estimate effects on the broader measure of the firm's workforce that includes part-time and contracted labor. Like Kline et al. (2019), we find larger effects for this broader measure of the firm's workforce, with a 14% increase in wage bill, 10% increase in number of workers, and 4% increase in mean earnings (all p-values below 0.01).

The evidence that wining a procurement auction causes the firm to bid up wages and hire more workers is at odds with the textbook model in which the labor supply curve facing the firm is perfectly elastic. Instead, it suggests that firms face upward sloping labor supply curves and, therfore, have wage setting power in the labor market. In Section 5, we recover the slope of the firm-specific labor supply curve, and thus the degree of imperfect competition in the labor market, from the employment and earnings impacts of the procurement win. The estimated 2% increase in earnings per worker relative to an 8% increase in employment is consistent with firms having non-negligible monopsony power.

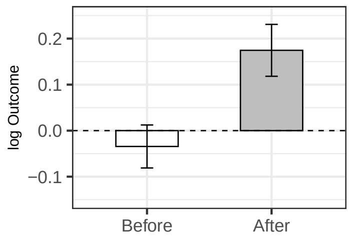

(a) Sales

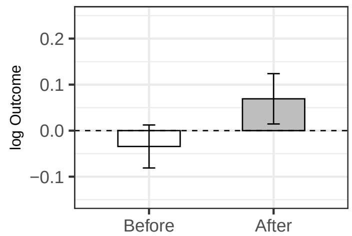

(b) Sales net of procurements

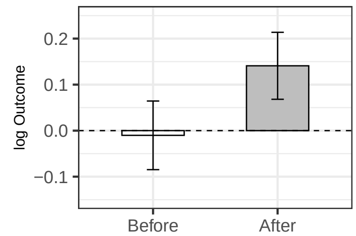

(c) Sales net of expenditures

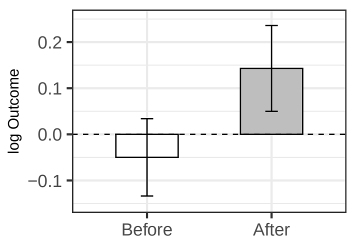

(d) EBITD

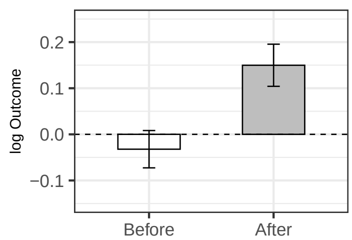

(e) Intermediate Costs

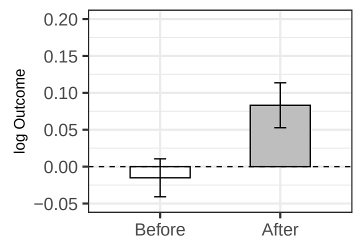

(f) Employees

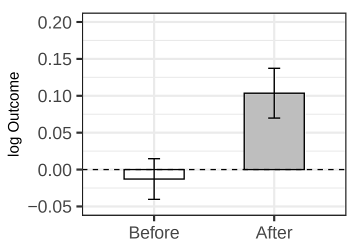

(g) Wage bill

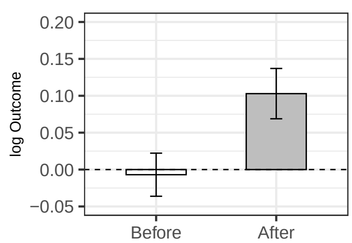

(h) Employees, broader measure

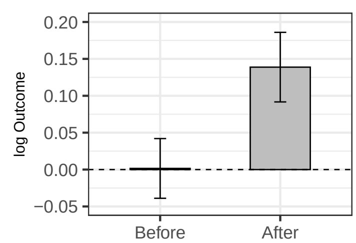

(i) Wage bill, broader measure

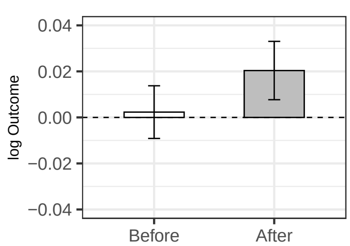

(j) Mean earnings, all employees

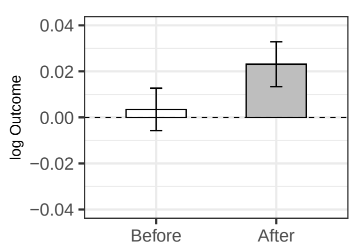

(k) Mean earnings, stayers

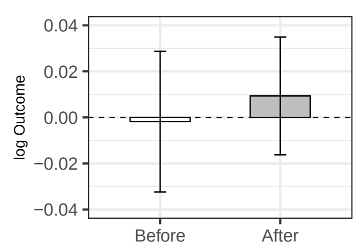

(l) Mean past earnings, new hires

Figure 1: Effects of Winning Procurement Auctions

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. "Before" refers to relative times {-4,-3,-2} and "After" refers to relative times {0,1,2}. Control firms are restricted to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. The omitted relative time is −2. 90% confidence intervals are displayed, clustering on firm.

One potential concern with this interpretation of the estimated effects on labor outcomes is that firms might engage in skill-upgrading in response to winning the auction. In other words, the increase in earnings per worker may arise from composition changes from low-skill to high-skill labor, not movement along the labor supply curve. To investigate this, we use we perform two checks. First, we consider earnings at previous firms as a proxy for worker quality of new hires. We estimate the difference-in-differences regression using this proxy as the outcome. As evidenced from Figure 1(l), we find that the average previous earnings of new hires does not experience a statistically significant change in response to winning a procurement auction (p-value about 0.5). This suggests firms do not engage in skill-upgrading, as the new hires are not significantly different before and after winning the auction.

The second way we check whether or not the observed increase in earnings per worker is due to skill-upgrading is to condition on the sample of workers that do not change firms. Mechanically, conditioning on stayers ensures we look at earnings for the same set of workers before and after the auction. In the baseline stayers estimates, we consider workers employed by the same firm during relative times (-2,...,2). This point estimate is virtually identical to the result for the full sample, finding in Figure 1(k) that earnings per incumbent worker increases by just over 2%. In Appendix Figure A.3(a), we vary the definition of a stayer by expanding and contracting the stayer window, always finding an increase of around 2% across definitions. In the analysis of the stayers sample, one may be concerned with conditioning on the potentially endogenous outcome of staying in the same firm after the auction outcome. This motivates the exercise in Appendix Figure A.3(b), where we condition on workers who have been employed at the firm for a certain number of years prior to the auction ("tenure"), but do not condition on the workers remaining in the firm after the auction. Thus, if the worker moves to a new firm, we use the earnings at the new firm as their outcome. We find a consistent 2% increase in earnings per worker, regardless of tenure.

4.2 Robustness Checks

As discussed in the previous section, the key concern with a comparison of outcomes for those who win and lose is that control firms may differ from treated firms in time-invariant characteristics. Our baseline specification addressed this concern by including firm fixed effects to capture time-invariant differences, and by verifying that there are no differential pre-trends in event years -4 to -2 (see the results in the "Before" bars). We do not use relative time -1 in the before and after contrast to ensure it cannot create bias in the DiD estimates. For completeness, in Appendix Figure A.7, we nevertheless examine results when including -1 and later time periods, finding that the estimates are very similar. In Appendix Figure A.6, we consider restricting the control group to firms that do not win any auctions throughout the "After" time interval, again finding the same patterns with slightly stronger effects. In Appendix Figure A.9, we provide main estimates when restricting to the subsample of firms from the five states that provided the EIN. For this subsample of firms, we link auction records to tax records exactly instead of relying on a fuzzy matching algorithm. The results are nearly identical, suggesting that the fuzzy matching algorithm recovers the same main estimates as exact matching.

To make the firms even more comparable and thus more likely to satisfy the parallel trends assumption, it may be useful to place stronger restrictions on the control group. In Appendix Figure A.4, we restrict to firms that are close to winning the auction in a cardinal sense. These are firms that placed bids within a certain threshold of the winning bid in dollar value. In Appendix Figure A.5, we restrict to firms that are close to winning the auction in an ordinal sense.[6] These are firms that placed bids lower than other bidders in

[6]The estimates when only considering second-place firms in the control group are similar but even less precise.

the auction, but still higher than the winner's bid. Though stronger sample restrictions weaken the precision of the estimator by reducing sample size, it is reassuring to find that the estimates are materially unchanged across the various alternatives.

Our main specification controls for differences in the composition of treated and control firms through firm fixed effects. This removes any time-invariant characteristics of firms, subsuming the identity of the auctions in which they participate at time c. However, these fixed effects are included in an additive fashion. As a robustness check, we re-estimate the difference-in-differences estimator separately for each auction, then average the treatment effect estimates across auctions. Standard errors are calculated using the block bootstrap, where a block is taken to be an auction. Results are displayed in Appendix Figure A.8. It is reassuring to find that the estimates do not depend on whether we include additive fixed effects or estimate the model separately for each auction.

5 Model

In this section, we develop a model where multiple construction firms compete with one another for projects in the product market and for workers in the labor market. Workers have heterogeneous preferences over non-wage job characteristics or amenities. This heterogeneity gives rise to imperfect competition in the labor market. There are two product markets in which the construction firms may choose to participate: projects in the private market and government projects which are procured through auctions. Firms are allowed to have price-setting power in both product and labor markets.

At the outset, it is useful to make clear the purposes of the model. First, it helps us to economically interpret the treatment effects of winning a procurement auction that we presented in Section 4. Without a model, it is hard to gauge the size of these treatment effects and difficult to understand the responses to the procurement-induced changes in product demand. Second, the model makes explicit the assumptions needed to recover the parameters that govern the behavior of firms and workers in the construction industry. This includes the firm-specific labor supply curve, the technology, and the market structure. Third, the model lets us quantify the incidence of the surplus in the construction industry in general and of the government spending on infrastructure projects in particular. Fourth, the model makes it possible to perform counterfactuals to explore how imperfect competition in the product market or the labor market affects the outcomes and behavior of both workers and firms in the construction industry.

5.1 Preferences and Labor Supply

Worker i has the following preferences over being employed at a given firm j,

$$u_i(j, W_j) = \log W_j + g_j + \eta_{ij} \tag{3}$$

where W_j represents earnings, g_j represents the average value of firm-specific amenities, and η_{ij} captures worker i's idiosyncratic tastes for the amenities of firm j. Since we allow amenities to be unobserved to the analyst, they can include a wide range of characteristics such as distance of the firm to the worker's home, flexibility in the work schedules, the type of work being performed, and so on.

Our specification of preferences allows for the possibility that workers view firms as imperfect substitutes. The term g_j gives rise to *vertical* employer differentiation: some employers offer good amenities while other employers have bad amenities. The term η_{ij} gives rise to *horizontal* employer differentiation: workers are heterogeneous in their preferences over the same firm. The importance of horizontal differentiation is governed by the variability across workers in the idiosyncratic taste for a given firm. We parameterize the distribution of η_{ij} as i.i.d. Type-1 Extreme Value (T1EV) with dispersion θ. When θ is larger, horizontal employer differentiation becomes relatively more important, as η_{ij} has greater variability.

We consider an environment where labor is hired in a spot market and assume that firms are endowed with a fixed set of amenities g_j (or, more precisely, we restrict amenities to be fixed over the estimation window). It is important to note that this restriction neither imposes nor precludes that employers *initially* choose amenities to maximize profits. Indeed, it is straightforward to show that permitting firms to initially choose amenities would not affect any of our estimates.

We consider two additional assumptions on the supply of labor. First, firms do not observe the idiosyncratic taste for amenities of any given worker η_{ij}. This information asymmetry implies that employers cannot price discriminate with respect to workers' reservation wages. Instead, if a firm wants to hire more labor, it needs to offer higher wages to both marginal and inframarginal workers. Second, since we find no evidence of changes in worker quality in response to winning a procurement auction, we assume homogenous labor. It is straightforward to extend the model and the empirical analysis to allow for differences in worker quality, provided that the firm's output depends on the total efficiency units of labor (see Lamadon et al., 2019).

Given these assumptions, the number of workers who accept a job at firm j for a posted wage offer W_j can be expressed as,

$$L_j = W_j^{1/\theta} g_j \bar{\theta} = W_j^{1/\theta} \mu_j^{-1/\theta},$$

where $\mu_j^{-1/\theta} \equiv g_j \bar{\theta}$ captures the vertical differentiation g_j and the size of the labor market $\bar{\theta}$.[7] Equivalently, we can write the inverse labor supply curve as,

$$W_j = L_j^{\theta} \mu_j. \tag{4}$$

The labor supply curve facing the firm, and thereby the wages it needs to pay to hire more workers, is increasing in the variability of the idiosyncratic tastes θ. The elasticity of the labor supply curve facing the firm is constant and given by $1/\theta$. In what follows, we assume firms are strategically small in the sense that $\frac{\partial \mu_j}{\partial W_j} \approx 0$. This means that the size of the labor market does not depend on marginal changes to the wage posted by any one firm, holding fixed all other firms' posted wages, own amenities, and all other firms' amenities.

5.2 Technology, Product Market, and Firm Behavior

We model firm behavior as a two-stage problem which we solve backwards. In the first stage, firms submit a bid for a government project that is procured through a first-price sealed-bid auction.[8] The project specifies the amount of output that must be produced within a given time frame. At the end of the first stage, the firm learns the auction outcome. If the firm wins the auction, it receives as revenue the winning bid amount and commences production. In the second stage, the firm chooses inputs to maximize profit from production in the private market, taken as given the outcome in the procurement auction. Private and government production occur simultaneously at the end of the second stage. Production in both private and

[7] Formally, $\bar{\theta} \equiv \bar{L}/\bar{W}$, where $\bar{L}$ is the population size and $\bar{W} \equiv \sum_{j'} W_{j'}^{1/\theta} g_{j'}$ is the price index of labor.

[8] For simplicity, we assume exogenous participation. Indeed, it is straightforward to extend the model to allow for a common, fixed cost of entry. The entry cost becomes sunk upon entry, and thus irrelevant for the firm's subsequent choices. As a result, the size of this fixed cost would not affect any of our estimates.

procurement projects occur simultaneously at the end of the second stage.

We begin by specifying the technology and the structure of the product market, before describing the firm's problem through the two stages.

Technology

Following Ackerberg et al. (2015), the production function (in physical units) is,

$$Q_j = \min\{\phi_j L_j^{\beta_L} K_j^{\beta_K}, \beta_M M_j\} \times \exp(e_j), \tag{5}$$

where ϕ_j denotes total factor productivity (TFP), K_j denotes capital, M_j denotes intermediate inputs (in physical units), and e_j represents measurement error.

We assume that capital markets are perfect, so firms can rent capital at constant rate r. The first-order condition for capital implies a composite production function,

$$Q_j = \min\{\Omega_j L_j^{\rho}, \beta_M M_j\} \times \exp(e_j), \tag{6}$$

where $\Omega_j \equiv \phi_j[\frac{\beta_K}{\beta_L}\frac{(1+\theta)\mu_j}{p_K}]^{\beta_K}$ and $\rho \equiv (1+\theta)\beta_K + \beta_L$. We refer to Online Appendix E for the derivation of the composite production function in equation (6). The Leontief production function in intermediate inputs (Ackerberg et al., 2015) implies that the optimal M_j is given by the function,

$$M_j = \Omega_j L_j^{\rho} / \beta_M. \tag{7}$$

This expression for optimal intermediates will prove useful in our identification strategy by providing an invertible relationship between labor and intermediates.

Product Market

We assume there are two product markets in which firms may choose to participate. First, they may participate in the market for private projects, which we denote H. Quantity produced by firm j in the private market is denoted by Q_j^H, which is endogenously chosen by the firm. Private projects are priced at $p\left(Q_j^H\right)^{-\epsilon}$ which implies revenues $R_j^H = p\left(Q_j^H\right)^{1-\epsilon}$. The parameter $1/\epsilon \geq 0$ is the price elasticity of demand in the private market. When $\epsilon > 0$, the demand curve facing the firm is downward-sloping and firms have price-setting power in the private market. Our derivations in the text focus on $\epsilon > 0$. We refer to Online Appendix F for derivations with perfect competition, $\epsilon = 0$. As we discuss below, $\epsilon = 0$ is at odds with our empirical results.

Second, firms may participate in the market for government projects, denoted by G. Government projects are allocated through procurement auctions, and the government sets the size of a project, $\bar{Q}^G$. If firm j loses the auction ($D_j = 0$), it does not produce in the government market ($Q_j^G = 0$). If firm j wins the auction ($D_j = 1$), it must produce exactly $\bar{Q}^G$ in the goverment market ($Q_j^G = \bar{Q}^G$). The quantity produced by firm j in the government market can then be expressed $Q_j^G = \bar{Q}^G D_j$. Revenues from winning a project of size $\bar{Q}^G$ are determined by equilibrium auction bidding, which is discussed below. At the end of the second stage,

At the end of the second stage, firm j produces total output $Q_j = Q_j^H + Q_j^G$ simultaneously across both markets using the production function in equation (6).

Second stage: Optimal firm choice in private market, given government project

We first solve for the optimal private market behavior of firm j if it wins an auction and if it loses an auction. Denote profits before auction revenues by $\pi^H_{1,j}$ if $D_j = 1$ and $\pi^H_{0,j}$ if $D_j = 0$. Denote the bid by b_j. Total profits are then $\pi_{1,j} = b_j + \pi^H_{1,j}$ for winners and $\pi_{0,j} = \pi^H_{0,j}$ for losers. Observed profits are $\pi_j = \pi_{1,j} D_j + \pi_{0,j}(1 - D_j)$. Given $\bar{Q}^G$ and D_j, the firm's second stage problem is to choose inputs $L_{d,j}$ and $M_{d,j}$ to maximize private market profits,

$$\pi^H_{d,j} = p\left(Q_{d,j} - \bar{Q}^G d\right)^{1-\epsilon} - W_j L_{d,j} - p_M M_{d,j}, \tag{8}$$

for $d = 0, 1$, subject to the labor supply curve (equation 4), the production function (equation 6), the condition of the optimal choice of intermediates (equation 7), and $Q_{1,j} \geq \bar{Q}^G$.

We define the opportunity cost of winning the auction as the difference in private market profits between losing and winning, $\Delta(\Omega_j) = \pi^H_{0,j} - \pi^H_{1,j}$, which emphasizes that firm productivity Ω_j is the only source of heterogeneity in the opportunity cost. Since the winning firm must allocate sufficient resources to the government project to produce $\bar{Q}^G$, increasing marginal cost implies that the marginal cost of projects in the private market is greater for winners. As a result, the opportunity cost of winning a contract is strictly positive, $\Delta(\Omega_j) > 0$. Note that the profit function for auction winners depends on $\bar{Q}^G$, so the opportunity cost also depends on $\bar{Q}^G$. For notational convenience, we suppress this dependence.

In Appendix C, we characterize the profit-maximizing choices for winners and losers. Here we emphasize several properties of the optimal solution. First, both winners and losers always produce strictly positive output in the private market. This follows from the fact that contractors have market power which implies that the marginal revenue in the private market is strictly greater than marginal cost as private market output approaches zero. For the same reason, total production is strictly greater if the firm wins the auction than if it loses.

Second, we note that the government project crowds-out private projects if $1 + \theta > \rho$, and conversely, crowds-in private projects if $1 + \theta < \rho$. To see why this is the case, note that winning a government project increases the total output level. This requires more employment to achieve a greater level of production. Due to the upward-sloping labor supply curve, greater employment leads to higher costs of labor, determined by $1+\theta$. On the other hand, greater scale induces private production under economies of scale, $\rho > 1$. Thus, the magnitude of $1 + \theta$ relative to ρ determine how the winning of a procurement auction affects the firm's production in the private market.

First stage: Auction model and optimal bidding for government project

In the first stage, bidders observe common information about the size of the project, $\bar{Q}^G$, and about the distribution of TFP, which we now describe. Each bidder begins by receiving an i.i.d. random draw of TFP $\Omega_j \sim \tilde{F}(\cdot)$, where $\tilde{F}(\cdot)$ is the TFP distribution. The TFP distribution and second stage optimization imply a distribution of the opportunity cost, i.e., $\Delta(\Omega_j) \sim F(\cdot)$, where the distribution function $F(\cdot)$ is known by all contractors. The benefit of winning the auction is the winning bid amount, b_j. Thus, the difference between the benefit and opportunity cost of winning the auction with bid b_j is $b_j - \Delta(\Omega_j)$.

Given the realized draw of Ω_j, firm j chooses the optimal bid b_j that solves the problem,

$$\max_{b_j} \underbrace{\left(b_j - \Delta(\Omega_j)\right)}_{payoff} \times \underbrace{\Pr(D_j = 1 | b_j)}_{probability\ of\ winning}.$$

The first term is the payoff to winning an auction, which is increasing in b_j, while the second term is the probability of winning an auction, which is decreasing in b_j. Thus, the contractor faces the usual trade-off between profits if one wins and the chances of winning. The profit-maximizing bidding strategy is,

$$b_j^* = \Delta\left(\Omega_j\right) + \frac{\int_{\Delta(\Omega_j)}^{\bar{\Delta}} [1 - F(u)]^{I-1} du}{\left[1 - F\left(\Delta\left(\Omega_j\right)\right)\right]^{I-1}}, \tag{9}$$

where I is the number of bidders.

This bidding strategy defines the unique symmetric equilibrium. To understand why, note that the existence of the private market provides a “walkaway” value for the bidder, to produce only in the private market. This gives an implicit participation constraint: the firm’s optimal bid must yield an expected payoff at least as high as the profit when losing, $\pi_{0,j}^H$, net of the private market profits received when winning, $\pi_{1,j}^H$ (the total payoff upon winning must be at least as great as the total payoff upon losing). Therefore, the bid must satisfy $b_j^* > \pi_{0,j}^H - \pi_{1,j}^H$, which is the opportunity cost, $\Delta\left(\Omega_j\right)$. Since $b_j^* > \Delta\left(\Omega_j\right)$, the bidding strategy is said to be “above the 45 degree line”. Furthermore, the bidding strategy is strictly increasing in $\Delta\left(\Omega_j\right)$. If a bidding strategy is “above the 45 degree line” and strictly increasing, it defines the unique symmetric equilibrium (Milgrom and Weber, 1982; Maskin and Riley, 1984).

5.3 Rents and the Economic Incidence of Procurements

Given the specification of the labor and product markets above, we can now define the surplus or rents that firms and their workers accrue. We focus both on the total rents from production and the additional rents from procurement projects. We define incidence as the additional rents generated by winning a procurement contract, and use the terms “additional rents” and “incidence” interchangeably.

Rents and Incidence for Workers

In our model, the employer may face an upward-sloping supply curve for labor, implying that the wage a firm pays can be an increasing function of its size. Since employers do not observe the idiosyncratic taste for amenities of any given worker, they cannot price discriminate with respect to workers’ reservation values. Instead, if a firm becomes more productive and thus wants to increase its size, the employer must offer higher wages to all workers. As a result, the equilibrium allocation of workers to firms creates surpluses or rents for inframarginal workers, defined as the excess return over that required to change a decision, as in Rosen (1986).

To define the additional rents to worker i from an exogenous wage increase at firm j from W_j to $\widetilde{W}_j$, we consider an equivalent variation (EV) representation. Denote worker i’s preferred firm excluding j as j^*. The EV of worker i for the wage increase at firm j, V_{ij}, is defined by the equation,

$$\underbrace{\max\left\{\log \widetilde{W}_j + g_j + \eta_{ij}, \log W_{j^*} + g_{j^*} + \eta_{ij^*}\right\}}_{\textit{utility with wage increase at firm } j} = \underbrace{\max\left\{\log\left(W_j + V_{ij}\right) + g_j + \eta_{ij}, \log\left(W_{j^*} + V_{ij}\right) + g_{j^*} + \eta_{ij^*}\right\}}_{\textit{equivalent utility at the initial choice of firm}}$$

The EV is the amount of compensation required at the initial choice of firm (right-hand side) to provide the same utility as the worker receives after the wage increase at firm j (left-hand side). There are two cases. If j is the initial choice of firm, then $V_{ij} = \widetilde{W}_j - W_j$. This is because the worker is an incumbent at firm j, so the wage gain at j is the amount of compensation required to achieve the same utility. If j^* is the initial choice of firm, V_{ij} is more complicated, as it must account for the relative differences in wages and preferences at firms j and j^*.

Letting $V_j \equiv \sum_i V_{ij}$ denote the total EV at firm j, Theorem 2 of Bhattacharya (2015) implies that,

$$V_j = \int_{W_j}^{\widetilde{W}_j} l_j(W) dW, \tag{10}$$

where $l_j(\cdot)$ is firm j's labor supply curve, which depends only on the wage at firm j under the assumption that each firm is strategically small. From our labor supply curve in equation (4), the solution to this integral is,

$$V_j = \frac{\widetilde{W}_j \widetilde{L}_j - W_j L_j}{1 + 1/\theta} = \frac{\widetilde{B}_j - B_j}{1 + 1/\theta} \tag{11}$$

where $L_j = l_j(W_j)$ is the initial labor, $\widetilde{L}_j = l_j\left(\widetilde{W}_j\right)$ is labor after the wage increase, $B_j = W_j L_j$ is the initial wage bill, and $\widetilde{B}_j = \widetilde{W}_j \widetilde{L}_j$ is the wage bill after the wage increase.

It is useful to observe that V_j answers two types of questions. The first is the incidence of procurements, i.e., the additional rents that workers gain from firm j winning a procurement contract. This can be computed from equation (11) by setting $W_j = W_{0,j}$ (the wage at firm j if it loses the auction) and $\widetilde{W}_j = W_{1,j}$ (the wage at firm j if it wins the auction), denoted $V_{\Delta,j}$, which is equal to $\frac{B_{1,j} - B_{0,j}}{1+1/\theta}$. The second is the total rents that workers earn from being inframarginal in the current choice of firm. This can be computed from equation (11) by setting $W_j = 0$ (the wage at which firm j shuts down production) and $\widetilde{W}_j = W_{0,j}$ (the wage at firm j if it loses the auction), denoted $V_{0,j}$, which is equal to $\frac{B_{0,j}}{1+1/\theta}$. This is identical to the total rents expression of Lamadon et al. (2019). Intuitively, it can be interpreted as the willingness-to-pay to stay at the current firm, which is greater when horizontal employer differentiation is more important (i.e., when θ is greater). In our context, it can be thought of as the "baseline" rents that the worker would receive if the firm lost the procurement auction. For convenience, we also define the total rents to workers if the firm wins the auction, which we denote $V_{1,j}$, so that

$$\underbrace{V_{1,j}}_{\textit{Total worker rents}} = \underbrace{V_{0,j}}_{\textit{Baseline worker rents}} + \underbrace{V_{\Delta,j}}_{\textit{Incidence on workers}} = \underbrace{\frac{B_{0,j}}{1 + 1/\theta}}_{\textit{Baseline worker rents}} + \underbrace{\frac{B_{1,j} - B_{0,j}}{1 + 1/\theta}}_{\textit{Incidence on workers}} \tag{12}$$

When using equation (11) to analyze the incidence of procurements, it is useful to decompose $V_{\Delta,j}$ into additional rents captured by incumbent workers and additional rents captured by new hires drawn into firm j by the wage increase. Expanding equation (10), we can write,

$$\begin{aligned} V_{\Delta,j} &= \underbrace{L_{0,j}\left(W_{1,j} - W_{0,j}\right)}_{\textit{gain for incumbents}} + \underbrace{\int_{W_{0,j}}^{W_{1,j}} \left(W_{1,j} - W\right) \frac{dl_j}{dW_j} dW}_{\textit{gain for new hires}} \\ &= \underbrace{L_{0,j}\left(W_{1,j} - W_{0,j}\right)}_{\textit{gain for incumbents}} + \underbrace{W_{1,j}\left(\frac{1}{1 + 1/\theta} L_{1,j} - L_{0,j}\right) + \frac{1}{1 + 1/\theta} B_{0,j}}_{\textit{gain for new hires}} \end{aligned} \tag{13}$$

This expression allows us to directly evaluate the share of welfare gains from procurements that are captured by incumbent workers relative to new hires.

Rents and Incidence for Firms

As our measure of firm rents, we use profits. There are three relevant measures of profits. First, $\pi_{0,j}$ is the profits that the firm captures from production in the private market if it loses the auction. Second, $\pi_{1,j}$ is the profits that the firm captures from joint production in the government and private markets if it wins the auction. Third, $\pi_{\Delta,j} \equiv \pi_{1,j} - \pi_{0,j}$ is the additional rents earned by the firm due to winning a procurement contract. We will make use of the decomposition,

$$\underbrace{\pi_{1,j}}_{\textit{Total firm rents}} = \underbrace{\pi_{0,j}}_{\textit{Baseline firm rents}} + \underbrace{\pi_{\Delta,j}}_{\textit{Incidence on firms}} \tag{14}$$

It is important to observe that profits do not necessarily represent ex-ante rents for the employer. Suppose, for example, that each employer initially chooses the amenities offered to the workers by deciding on the firm's location, the working conditions, or both. Next, the employers compete with one another for the workers who have heterogeneous preferences over the chosen amenities. These heterogeneous preferences give rise to wage-setting power which employers can use to extract additional profits or rents. Of course, the existence of such ex-post rents could simply be returns to costly ex-ante choices of amenities. On top of this, profits from the procurement auctions may, in part, reflect a fixed cost of entry to the auction. For example, in order to bid on procurement contracts, firms must hold licenses which are costly. While the presence of a fixed entry cost will affect the interpretation of profits, it will not affect identification of model parameters.

6 Model Parameters: Identification, Estimation and Interpretation

The purpose of this section is to identify, estimate and economically interpret the model parameters of interest $\left(\theta, \epsilon, \rho, \beta_L, \beta_K, \bar{\theta}, p, \beta_M/p_M, F(\cdot), \mathbb{E}\left[V_{\Delta,j}\right], \mathbb{E}\left[\pi_{\Delta,j}\right]\right)$ given the data $(L_j, W_j, \pi_j, R_j, p_M M_j, D_j, b_j)$. We provide a formal identification argument in the text while summarizing, in Table 2, the moments used to identify each parameter of interest. The identification argument forges a direct link between our model in Section 5 and the procurement effects analysis in Section 4. For notational simplicity and without loss of generality, the formal argument keeps the conditioning on $X = 1$ and the fixed effects implicit, or, equivalently, D_j is treated as if it is randomly assigned.

At the outset, we emphasize three assumptions that are key for our identification results:

Assumption 1: Workers' idiosyncratic taste over non-wage attributes of firms is distributed T1EV.
This specification of preferences is standard in the empirical literature on monopsony. It is useful for identification becuse it gives a parsimonious measure of wage-setting power through a constant elasticity of labor supply.

Assumption 2: The production function is Leontief in intermediate inputs.
This common specification of technology implies that intermediate input is proportional to output. It makes it possible to derive an invertible relationship between intermediates and labor, which can be used to control for unobserved productivity differences across firms.

Assumption 3: Demand elasticity in private market is constant and positive.
A constant elasticity offers tractability, while positivity ensures that firms that win auctions also participate in the private market (as we observe in the data).

As our identification argument makes clear, many of the parameters of interest do not require all three assumptions. Thus, some of our findings may be considered more reliable than others.

6.1 Labor Supply Elasticity and Incidence

We begin by identifying and estimating the labor supply elasticity, $1/\theta$, and the incidence terms for workers and firms, $\mathbb{E}[V_{\Delta,j}]$ and $\mathbb{E}[\pi_{\Delta,j}]$, defined in Subsection 5.3.

Identification and Estimate of the Labor Supply Elasticity

Under Assumption 1, total wage bill B_j can be expressed,

$$\log B_j = \log(W_j L_j) = (1+\theta)\log L_j + \log \mu_j.$$

We can write θ in terms of changes induced by winning a procurement auction as,

$$1+\theta = \frac{\mathbb{E}[\log B_j|D_j=1] - \mathbb{E}[\log B_j|D_j=0]}{\mathbb{E}[\log L_j|D_j=1] - \mathbb{E}[\log L_j|D_j=0]}. \tag{15}$$

Intuitively, θ determines the cost of hiring more workers, which can be recovered from the change in payments to labor compared to the change in the number of workers due to winning the procurement auction.

The main estimate of θ is displayed in Panel A of Table 2. The point estimate of the firm-specific labor supply elasticity is 4.1 in our sample of firms. This indicates that, if an American construction firm aims to increase the number of employees by 10%, it can accomplish this by increasing wages by around 2.4%. Suárez Serrato and Zidar (2016) and Lamadon et al. (2019) estimate local labor supply elasticities of 4.2 and 4.6, respectively, while Card et al. (2018) pick 4.0 as the preferred value in their calibration exercise. A related literature using experimentally manipulated piece-rate wages for small tasks typically finds labor supply elasticities ranging from 3.0 to 5.0 (Caldwell and Oehlsen, 2018; Dube et al., 2020; Sokolova and Sorensen, 2018).

Figure 2(a) provides the labor supply elasticity estimates, $1/\theta$, corresponding to each of the robustness exercises discussed in Section 4. The robustness checks include restricting to those firms that made bids close to the winning bid, in a cardinal sense ("Close Bidders Cardinal") and in an ordinal sense ("Close Bidders Ordinal"), including all time periods in the estimation ("All Time Periods"), restricting the control group to those that did not win subsequent auctions in the near future ("Non-winners"), separately estimating the effect within each auction and then averaging ("Within Auction"), and estimating the labor supply elasticity as the log change in employees divided by the log change in mean earnings of stayers ("Stayers") rather than relying on the log change in wage bill. Across all the robustness checks, the range of labor supply elasticity estimates is 3.5 to 5.1, which are close to the baseline estimate and align with the range of estimates in the literature.

As shown above, our identification of θ relies on the argument that winning an auction shifts the firm's demand for labor along the labor supply curve. One potential reason this argument may fail is adjustment costs: If labor enters the firm slowly over time rather than immediately when the new wage is posted, the short-run relation between wages and quantity of labor may understate the longer-run elasticity of labor supply. The evidence in Section 3 is at odds with such adjustment costs. Both wages and labor quantity appear to respond relatively quickly to winning the auction (see Appendix Figure A.2f,g). Indeed, the implied labor supply elasticity varies relatively little over time.

Panel A. Labor Supply and Baseline Incidence Analysis		
	Moments used in Identification	
	Definition	**Estimate**
Estimand for winning an auction:		
Estimand log wage bill	$\tau_b \equiv E[\log B_j \mid D_j = 1] - E[\log B_j \mid D_j = 0]$	0.104
Estimand log employment	$\tau_l \equiv E[\log L_j \mid D_j = 1] - E[\log L_j \mid D_j = 0]$	0.083
Estimand log profits	$\tau_\pi \equiv E[\log \pi_j \mid D_j = 1] - E[\log \pi_j \mid D_j = 0]$	0.143
	Labor Supply and Incidence Parameters	
	Identifying Equation	**Estimate**
Labor supply elasticity	$1/\theta = \tau_l/\tau_b - 1$	4.084
Incidence on Workers ($1,000/employee)	$V_\Delta = \frac{B}{L}\frac{\tau_b}{1+1/\theta}$	0.959
Incidence on Firms ($1,000/employee)	$\pi_\Delta = \frac{\pi}{L}\tau_\pi$	2.873
Workers' Share of Incidence	$V_\Delta/(V_\Delta + \pi_\Delta)$	0.250
Panel B. Firm's Problem and Private Market		
	Moments used in Identification	
	Definition	**Estimate**
Estimand for winning an auction:		
Estimand log inputs relative to labor	$\tau_{m/l} \equiv \frac{E[\log p_M M_j \mid D_j=1] - E[\log p_M M_j \mid D_j=0]}{E[\log L_j \mid D_j=1] - E[\log L_j \mid D_j=0]}$	1.600
Estimand log private revenues	$\tau_{R-b} \equiv E[\log(R_j - b_j) \mid D_j = 1] - E[\log R_j \mid D_j = 0]$	0.069
Auction loser log covariances:		
Var. of log inputs	$\sigma^2_{m\mid 0} \equiv Var[\log(p_M M_j) \mid D_j = 0]$	3.062
Cov. of log inputs and log revenues	$\sigma_{m,r\mid 0} \equiv Cov[\log(p_M M_j), \log(R_j) \mid D_j = 0]$	2.559
	System of Equations	
	Identifying Equation	**Evaluated Right-hand Side**
Optimal intermediate inputs to employees (eq 18)	$\rho = \tau_{m/l}$	1.600
Optimal labor to value added ratio (eq 19)	$\beta_L = \mathbb{E}\left[\frac{(1+\theta)B_j}{(1-\epsilon)R_j - M_j}\right]$	0.587
Diminishing returns to private output (eq 20)	$1 - \epsilon = \sigma_{m,r\mid 0}/\sigma^2_{m\mid 0}$	0.836
First-order condition with crowd-in (eq 21)	$\rho = \frac{\tau_{R-b}}{\tau_l} + \frac{\mathbb{E}\left[\log\left(\frac{1+\theta}{\beta_L}B_j + p_M M_j\right) - \log((1-\epsilon)R_j)\right]}{\tau_l}$	1.401
	GMM Results	
	Parameters	**Estimate**
Private demand elasticity	$1/\epsilon$	7.343
Returns to labor	β_L	0.597
Composite scale parameter	ρ	1.419
Panel C. Remaining Parameters for Price, Scale, and TFP		
	Moments used in Identification	
	Definition	**Estimate**
Mean log wage bill	$\mu_b \equiv E[\log(B_j)]$	13.682
Mean log employment	$\mu_l \equiv E[\log(L_j)]$	2.914
Mean log intermediate expenditure	$\mu_m \equiv E[\log(p_M M_j)]$	14.958
Mean log revenues	$\mu_r \equiv E[\log(R_j)]$	15.291
	Remaining Parameters	
	Identifying Equation	**Estimate**
Scale of optimal log wage	$\log \bar{\theta} = \mu_b - (1+\theta)\mu_l$	10.055
Scale term for intermediates	$\log \frac{\beta_M}{p_M} = \rho\mu_l - \mu_m$	-10.824
Scale of log output price	$\log p = \mu_r - (1-\epsilon)(\log \frac{\beta_M}{p_M} + \mu_m)$	11.719
Interquartile range of log TFP	$\text{IQR}(\log \Omega_j) = \text{IQR}(\log p_M M_j - \rho \log L_j)$	0.973

Table 2: Model Identification and Parameter Estimates

Notes: This table summarizes all results on identification and estimation of the model. In Panel A, it provides the parameters needed to estimate the incidence of procurements on firms and workers. Incidence is evaluated at the median firm in our sample. In Panels B and C, it provides the parameters needed to simulate counterfactual results from the model.

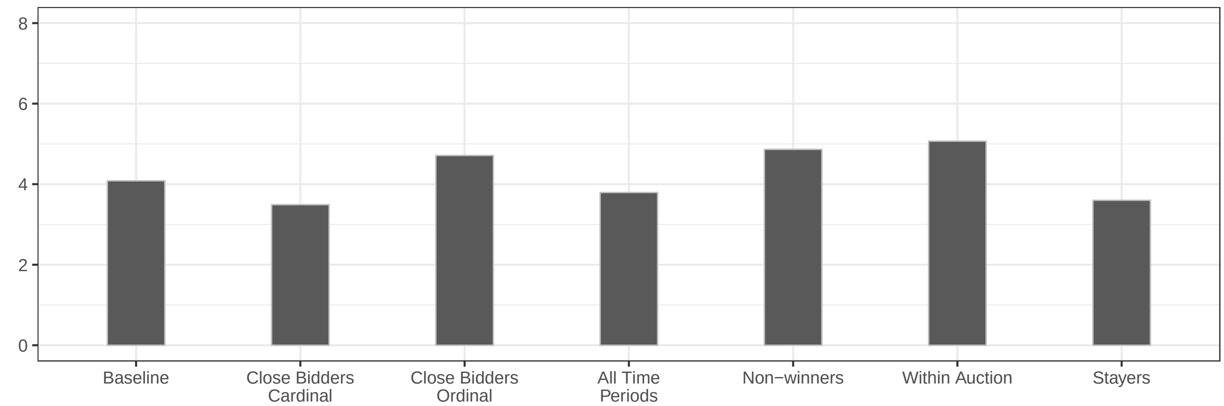

(a) Labor Supply Elasticity

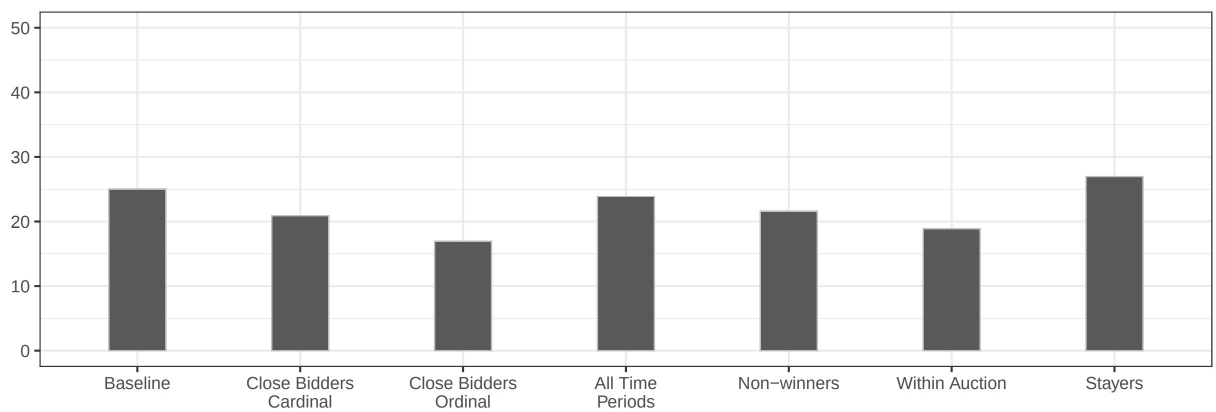

(b) Worker Share of Incidence (%)

Figure 2: Labor Supply Elasticity and Incidence Share Estimates

Notes: This figure presents the baseline estimate and robustness checks for the labor supply elasticity, $1/\theta$, and the worker share of incidence. For the sample of Close Bidders (Ordinal), we consider control group firms that finished in the three lowest bidders in an auction. For the sample of Close Bidders (Cardinal), we consider control group firms that bid no more than 10% higher than the winning bid. Specification details on the robustness checks and alternative sample definitions are provided in Section 4 and Appendix A.

Another possible threat to our identification of the labor supply curve is skill upgrading: If the wage bill increases for the winning firm both because more workers are hired and because the new workers are more efficient, then the estimator will include a bias term related to the change in worker composition. In Section 4, we provided evidence that the composition of new hires does not appear to change in response to winning an auction. When conditioning on the incumbent workers in the firm so that composition mechanically does not affect the wage change, we estimate approximately the same labor supply elasticity (see the “stayers” bar in Figure 2). Thus, we argue our data is at odds with significant changes in the skill composition of workers.

Identification and Estimates of the Incidence on Workers and Firms

In Subsection 5.3, we provided expressions for the total rents earned by workers and firms, as well as the additional rents (incidence) due to procurements. We now show how these rents and incidence expressions can be quantified in the data.

Given that D_j is (conditionally) randomly assigned, we can express the incidence on workers as the observed difference in wage bill between firms that win and lose a procurement auction divided by one plus the labor supply elasticity,

$$\mathbb{E}\left[V_{\Delta,j}\right] = \frac{\mathbb{E}[B_j|D_j = 1] - \mathbb{E}[B_j|D_j = 0]}{1 + 1/\theta}, \tag{16}$$

By comparison, the incidence on firms is the observed difference in profits between firms that win and lose a procurement auction,

$$\mathbb{E}\left[\pi_{\Delta,j}\right] = \mathbb{E}[\pi_j|D_j = 1] - \mathbb{E}[\pi_j|D_j = 0]. \tag{17}$$

Given the incidence expressions (equations 12 and 14), we can express total rents earned if the firm wins the auction by $\mathbb{E}\left[V_{1,j}\right] = \frac{1}{1+1/\theta}\mathbb{E}[B_j|D_j = 1]$ for workers and $\mathbb{E}\left[\pi_{1,j}\right] = \mathbb{E}[\pi_j|D_j = 1]$ for firms, and total rents in the economy are $\mathbb{E}\left[V_{1,j}\right] + \mathbb{E}\left[\pi_{1,j}\right]$. We can then use the relationships $\mathbb{E}\left[V_{1,j}\right] = \mathbb{E}\left[V_{0,j}\right] + \mathbb{E}\left[V_{\Delta,j}\right]$ and $\mathbb{E}\left[\pi_{1,j}\right] = \mathbb{E}\left[\pi_{0,j}\right] + \mathbb{E}\left[\pi_{\Delta,j}\right]$ to decompose worker and firm rents into the baseline components earned by auction losers and the additional rents due to auction winnings. Furthermore, we can decompose the additional rents for workers, $\mathbb{E}\left[V_{\Delta,j}\right]$, into the components earned by incumbent workers and new hires using equation (13).

Panel A of Table 2 provides the baseline point estimate of the incidence of procurements on workers using the expressions above. We report dollar values in per-worker units, where the number of workers corresponds to the actual number of employees observed for auction winners. The results in Table 2 suggest that workers gain, on average, about \$1,000 at the typical firm from winning a procurement auction. This calculation includes both incumbent workers and new hires. The results in Table 2 suggest an incidence of procurements on firms of about \$2,900 per worker at the typical firm. Together, these estimates imply that workers receive around 25% of the additional rents from winning a procurement auction.

Figure 2(b) provides the worker share of incidence estimates corresponding to each of the robustness exercises discussed in Section 4. Across all the robustness checks, the range of incidence share estimates is 17% to 27%, which are relatively close to the baseline estimate. By comparison, Suárez Serrato and Zidar (2016) find that the incidence of corporate tax changes on workers in the U.S. is 28%.

In Figure 3, we compare the size of additional rents, $V_{\Delta,j}$ and $\pi_{\Delta,j}$, to the baseline rents workers and firms would have earned if they had lost the procurement auction, $V_{0,j}$ and $\pi_{0,j}$. This uses the decompositions in equations (12) and (14). In total, auction winners generate rents of \$30,328 per worker, of which \$20,102 (66%) is captured by firms while the remaining \$10,226 (34%) is captured by workers. Out of the \$20,102 captured by firms, \$17,229 (86%) are due to private market activity if the firm loses the auction, while the additional rents of \$2,873 (14%) are due to winning the auction. Out of the \$10,226 captured by workers, \$8,307 (90%) are due to private market activity if the firm loses the auction, while the additional rents of \$960 (10%) are due to winning the auction. The additional rents to workers are split between incumbent workers and new hires, with \$880 (92%) captured by incumbent workers.

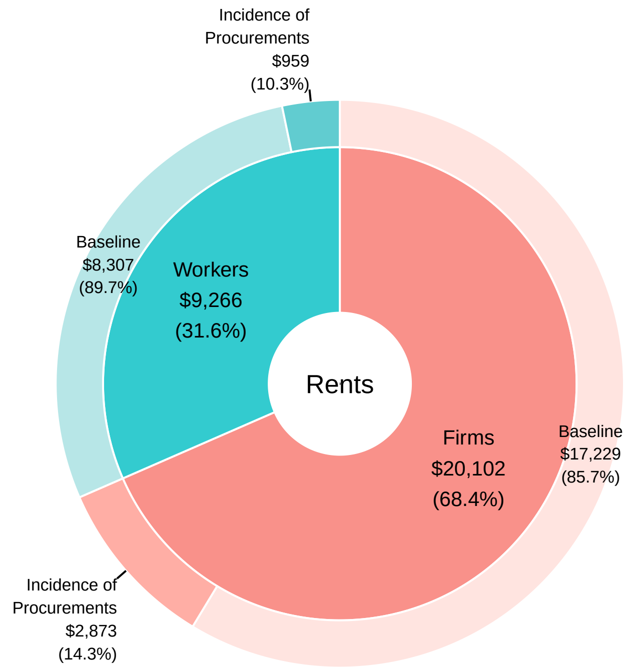

Figure 3: Aggregate Rents

Notes: This figure presents the total rents among auction winners in the economy. It divides the rents into those captured by firms (red) and those captured by workers (blue). It further decomposes rents into those earned if the firm loses the auction ("Baseline") and the additional rents due to winning the procurement auction ("Incidence of Procurements").

6.2 Firm Technology and Product Market

The previous subsection demonstrated identification of the labor supply elasticity and incidence under Assumption 1. In this subsection, we provide identification results for the three main parameters that govern the product market and firm technology. These parameters are the product demand elasticity in the private market, $1/\epsilon$, the returns to labor in production, β_L, and the composite production parameter, ρ. Identification of these parameters requires Assumptions 1, 2, and 3, as well as the estimate of $1/\theta$ from the previous subsection. Knowledge of these parameters is essential to simulate optimal firm behavior in counterfactual environments, which will be important for studying the relationship between incidence and imperfect competition in the next section. We provide main results in the text and refer the reader to Appendix C for derivations.

Identification of the Main Parameters Governing Firm Technology and Product Market

Consider first how to recover ρ from the optimal materials condition (equation 7) under Assumptions 2. This condition implies that the ratio of differences in intermediate input expenditures and the quantity of labor hired must equal the composite returns ρ,

$$\rho = \frac{\mathbb{E}\left[\log\left(p_M M_j\right)|D_j = 1\right] - \mathbb{E}\left[\log\left(p_M M_j\right)|D_j = 0\right]}{\mathbb{E}\left[\log L_j | D_j = 1\right] - \mathbb{E}\left[\log L_j | D_j = 0\right]} \tag{18}$$

Identification relies on the optimal proportionality between labor and materials as the firm increases output. Intuitively, if ρ is high, this means that for a given change in labor, materials must adjust by more to maintain proportionality.

Second, under Assumptions 1, 2, and 3, the first-order condition with respect to labor for firms that only operate in the private market ($D_j = 0$) provides a relationship between the returns to labor, β_L, and the ratio between labor expenses and revenues. Using that $\rho \equiv (1+\theta)\beta_K + \beta_L$ and taking the expectation across losing firms, the first-order condition with respect to labor (equation 8) implies,

$$\beta_L = \mathbb{E}\left[\frac{(1+\theta)B_j}{(1-\epsilon)R_j - p_M M_j} | D_j = 0\right]. \tag{19}$$

The firm equates the marginal cost of labor $W_j(1+\theta)$ with the marginal benefit of labor $\frac{1}{L_j}\left((1-\epsilon)R_j\beta_L - p_M M_j \beta_L\right)$. This expression for the marginal benefit of labor includes an adjustment for materials because, in order for the firm to produce more output, it must increase both labor and materials under Assumption 2. In the limiting case as $\theta \to 0$ and $\epsilon \to 0$, β_L is equal to the labor share (wage bill divided by value added). The terms related to θ and ϵ account for the fact that both the output and labor markets are imperfectly competitive. If labor is very productive at the margin, the firm will optimally choose a higher wage bill. Conversely, the extent to which the wage bill is high relative to value-added reveals that the marginal product of labor must also be high.

Third, under Assumptions 2 and 3, notice that revenues if firm j loses the auction are $R_{0,j} = p(\beta_M M_{0,j} \exp(e_j))^{1-\epsilon}$, leading to,

$$\log R_{0,j} = \log p + (1-\epsilon)[\log(\beta_M/p_M) + \log(p_M M_{0,j}) + e_j],$$

Using the assumption of Ackerberg et al. (2015) that e_j is unobserved to the firm at the time it chooses materials (or, equivalently, it is measurement error in log revenues), it follows that,

$$1 - \epsilon = \frac{\text{Cov}[\log(R_j), \log(p_M M_j)|D_j = 0]}{\text{Var}[\log(p_M M_j)|D_j = 0]} \tag{20}$$

which uses that $\text{Cov}\left[e_j, \log(p_M M_j)\right] = 0$. This is an application of the "control function" approach of Ackerberg et al. (2015): when the firm makes its optimal input choices, Ω_j is known. Thus, the firm's choice of inputs is conditional on Ω_j, as can be seen in equations (4) and (7). On the other hand, since output is equal to $log\Omega_j + \rho log L_j$, revenues depend directly on Ω_j and indirectly through L_j. Since Ω_j is unobservable to the econometrician, there is an endogeneity problem. The Leontief production function implies that $p_M M_j$ is a control function for Ω_j. Thus, we can rewrite the revenue function in terms of materials and what remains is e_j which is independent of $p_M M_j$.

Fourth, from equation (8) and Assumptions 1, 2, and 3, we can write output in the private market as,

$$\log Q^H_{d,j} = [\log \Omega_j + \rho \log L_{d,j}] - d\left[\log\left(\frac{1+\theta}{\beta_L} B_{1,j} + p_M M_{1,j}\right) - \log\left((1-\epsilon)R^H_{1,j}\right)\right].$$

Combining the cases $d = 0, 1$ and using the random assignment of D_j, we arrive at the moment condition,

$$\underbrace{\mathbb{E}\left[\log\left(R_j - b_j\right)\Big|D_j = 1\right] - \mathbb{E}\left[\log R_j\Big|D_j = 0\right]}_{\text{Crowd−in}} = \rho\underbrace{\mathbb{E}\left[\log L_j\Big|D_j = 1\right] - \mathbb{E}\left[\log L_j\Big|D_j = 0\right]}_{\text{Total labor response}} \tag{21}$$

$$- \underbrace{\mathbb{E}\left[\log\left(\frac{1+\theta}{\beta_L}B_j + p_M M_j\right) - \log\left((1-\epsilon)R^H_{1,j}\right)\big|D_j = 1\right]}_{\text{Adjustment for labor used to complete procurement projects}}$$

On the left-hand side of this condition is the definition of crowd-in and on the right-hand side are two terms: the revenue that would be generated by the observed labor difference, and an adjustment for the partial labor adjustment required to complete the procurement contract.

Estimates of the of the Main Parameters Governing Firm Technology and Product Market

We use the general method of moments (GMM) to jointly estimate $(1/\epsilon, \rho, \beta_L)$ based on equations (18-21). To simplify the search space in the numerical solver, we impose the natural constraints $\epsilon \geq 0$, $\rho \geq 0$, $\beta_L \in [0, 1]$, and $\rho(1 - \epsilon) < (1 + \theta)$, where the latter constraint ensures that firms do not optimally choose to be infinitely large. However, none of the constraints bind at the numerical solution.

We estimate $1/\epsilon$ to be 7.3. This implies that, in order for a firm to increase output by 7.3%, it must reduce price by 1%. Though we do not find directly comparable estimates of the price elasticity of demand from the construction industry, some estimates from the literature suggest our estimate is within a reasonable range. Goldberg and Knetter (1999) estimated residual demand elasticities for German beer ranging from 2.3 to 15.4. Goldberg and Verboven (2001) estimated average price elasticities of demand for foreign cars ranging from 4.5 to 6.5.

We estimate β_L to be 0.60 and ρ to be 1.4. The value of β_L implies that a 100% increase in a firm's employment results in 60% more output, all else equal. This 60% share is broadly similar to the aggregate labor share of income in the U.S. over the same time period. The value of ρ implies that, if a firm has 100% more labor than another firm, we expect it to produce 140% more output, *not* holding all else equal. The larger firm will optimally have greater utilization of capital from the rental market and materials from the intermediates market, which are accounted for in the ρ parameter. From the definition of ρ, we can back out the implied returns to scale as $\beta_L + \beta_K = \beta_L + \left(\frac{\rho-\beta_L}{1+\theta}\right)$, which equals about 1.2. Reassuringly, our returns to scale estimate is comparable to the range of estimates from 1.0 to 1.2 by Levinsohn and Petrin (2003). Note that imperfect competition in the product and labor markets attenuates incentives for firms to grow infinitely large even with increasing returns to scale.

Model Fit

Because of the tight link between the data and the parameters of interest, our model fits perfectly many moments. However, some moments could potentially fit poorly. For example, it is possible for the model to imply a private market crowd-in of procurements that does not match our estimate from the data. It is reassuring to find that the crowd-in rate implied by the estimated model is 0.07, which is nearly identical to our estimate in Section 4. Similarly, the returns to scale and diminishing returns to private output equations are closely fit by the estimated parameters.

6.3 Identification and Estimation of the Remaining Parameters

For the few remaining model parameters, the identifying equations and estimates are provided in Panel C of Table 2. These include the scale of the labor market, $\bar{\theta}$, the relative price in the private market, p, the relative material returns versus cost, β_M/p_M, and the interquartile range of distribution of estimated TFP, to characterize $F(\cdot)$. Identification requires Assumptions 1, 2, and 3, as well as the estimates of $1/\theta$, ρ, and $1/\epsilon$ from the previous subsections. Although the magnitudes of these parameters are perhaps not of interest on their own, they are needed to perform counterfactual simulations in the next section. Appendix C provides derivations of the identifying equations.

7 Imperfect Competition and Incidence

In this section, we use the model to understand how imperfect competition in the labor and product markets affects the outcomes and behavior of workers and firms in the American construction industry.

Before presenting the results from this analysis, it is important to observe that simulating counterfactuals is computationally challenging. In particular, since $1/\theta$ and $1/\epsilon$ both appear in the firm's opportunity cost $\Delta(\Omega_j)$ (equation 8), it follows that changing these parameters also changes the optimal bid b_j^* (equation 9). In turn, the bid affects the additional rents captured by firms from winning a procurement contract. To perform the counterfactuals, we first solve the second stage problem for each Ω_j to find the counterfactual distribution of opportunity costs. Next, we solve the first stage problem to obtain the distribution of optimal bids given the counterfactual opportunity costs. Finally, we combine the optimal bid distribution from the first stage with the optimal private market profits from the second stage. From this, we recover the counterfactual outcomes, such as profits. To ease the computational burden in solving for these distributions in the two-stage problem, we implement the quantile representation method of Luo (2019). We focus on counterfactual results for the typical firm (as defined by the firm with the median value of Ω_j), which further reduces the computational burden. The algorithm and computational details are provided in Online Appendix G.

7.1 The Importance of Imperfect Competition in the Labor Market

Defining and Interpreting the Counterfactual

To study the importance of imperfect competition in the labor market, we consider a compensated rotation of the labor supply curve of a given firm, holding all other firms fixed, so that the labor supply elasticity, $1/\theta$, decreases while the initial equilibrium labor and wage choices remain feasible for the firm. In practice, this means that we first solve for the initial monopsonistic equilibrium (ME) in the labor market, (L_j, W_j), shift $1/\theta$ to $1/\theta'$, then compensate firm j for this increase in the average cost of labor by shifting the firm-specific labor supply curve intercept μ_j to μ_j' so that (L_j, W_j) is still on the labor supply curve.

Figure 4 provides an illustration of rotating the labor supply curve. It considers a fictional firm j, with labor on the x-axis and the wage on the y-axis. The initial equilibrium is in black, while the equilibrium after rotating the labor supply curve is in red. The initial average cost of labor curve (ACL, solid line) and its associated marginal cost of labor curve (MCL, dashed line) are in black. The marginal revenue product of labor curve (MRPL) is also in black. To determine the equilibrium (ME), the monopsonistic firm chooses labor to equate MCL and MRPL, then chooses the lowest feasible wage at this quantity of labor, which is on the ACL curve directly below the intersection of MCL and MRPL.

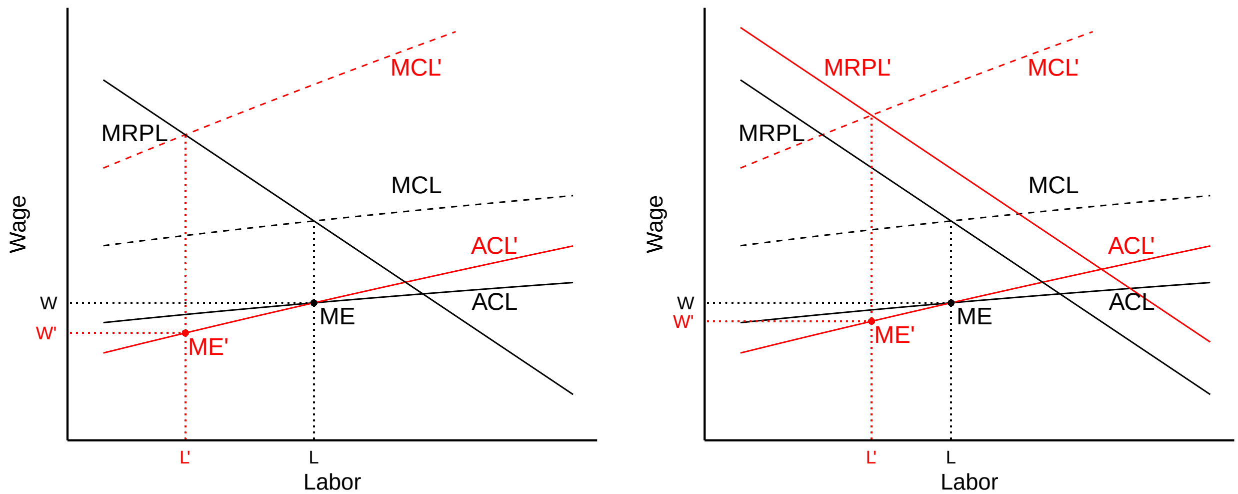

(a) Without MRPL Shift

(b) With MRPL Shift

Figure 4: Illustration: Rotation of the Labor Supply Curve

Notes: This diagram visualizes the counterfactual exercise of making the labor market less competitive. ACL denotes the average cost of labor, MCL denotes the marginal cost of labor, ME denotes the monopsonistic equilibrium, and MRPL denotes the marginal revenue product of labor. Black colors denote the initial economy, and red colors denote the new economy after rotating the labor supply curve through a compensated decrease in the labor supply elasticity, $1/\theta$. In subfigure (a), MRPL is not allowed to adjust when $1/\theta$ decreases, while in subfigure (b), the MRPL is allowed to adjust when $1/\theta$ decreases.

The red lines in Figure 4(a) demonstrate how the equilibrium adjusts when the labor supply curve is made "steeper" by lowering $1/\theta$ to $1/\theta$'. Lowing $1/\theta$ raises the average cost of labor, so we compensate the firm by decreasing the intercept of each firm's labor supply curve until the initial ME is on the new labor supply curve (that is, we ensure ME is on ACL'). The new marginal cost of labor curve (MCL') is higher, so the point at which MRPL equals MC' is at a lower level of labor (that is, L' is less than L on the x-axis). Since ACL' is lower than ACL, it follows that ME' must also have a lower wage than ME (W' is less than W on the y-axis). Furthermore, since ME is a feasible choice of the firm, and the firm chooses a different point ME' to maximize profits, it must have higher profits with the counterfactual labor supply curve. Thus, the counterfactual exercise in Figure 4(a) always results in firm j employing fewer workers, paying a lower wage to each employee, producing less private market output, becoming more capital-intensive, and earning higher profits.

Figure 4(b) presents the same exercise, but allowing the MRPL to also shift in response to the decrease in $1/\theta$. The reason decreasing $1/\theta$ results in higher MRPL is because firms produce using capital, which is rented in a perfect capital market. When the marginal cost of labor is higher, firms will choose to use more capital per worker. Thus, a one worker increase in labor corresponds to a greater marginal increase in capital if $1/\theta$ is lower, which implies that MRPL increases when $1/\theta$ decreases. As a result, the point at which MC' equals MRPL' is to the right of the point at which MC' equals MRPL, which implies that L' is greater when MRPL shifts than when MRPL does not shift. Thus, the counterfactual exercise in Figure 4(b) results in a less extreme decrease in labor and wages than the exercise in Figure 4(a).

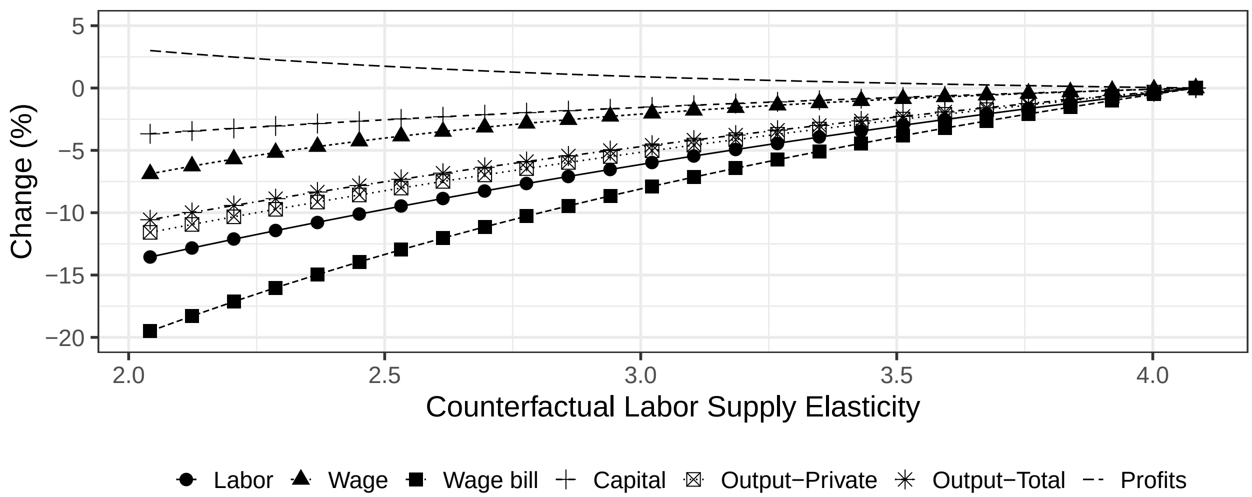

Figure 5: Counterfactual Rotation of the Labor Supply Curve

Notes: This figure presents the counterfactual median values of labor, wages, the wage bill, capital, private market output, total output, and profits when the labor market becomes less competitive. It expresses these values as percentage changes relative to the actual economy.

Characteristics of the Actual and Counterfactual Labor Markets

In Figure 5, we present the counterfactual analysis where we increase the degree of imperfect competition in the labor market by rotating the labor supply curve. We find that, as the firm gains more market power, it employs fewer workers and pays a lower wage to each employee. By taking advantage of its market power to increasingly markdown wages, the firm earns higher profits. These results are consistent with the predictions from the diagram in Figure 4. Because of the monotonic relationships across values of the labor supply elasticity, we focus on comparing the actual value ($1/\theta = 4.08$) and half of this amount ($1/\theta' = 2.04$).

The empirical estimates from our counterfactuals show that the firm employs 15% fewer workers when the labor supply elasticity is reduced by half. Wages decrease by 7% and the wage bill decreases by nearly 20%. Capital only decreases by 4%, indicating a shift toward capital-intensive production as the marginal cost of labor rises. As a result of the reductions in labor and capital, the output quantity (in the private market) falls by about 12%. Since government projects (which are of the same size in the actual and counterfactual economies) make up a small share of all output, total output falls by less than private market output. Because firms increasingly markdown wages, profits increase by more than 3%.

Comparing Incidence of Procurements in the Actual and Counterfactual Labor Markets

In Table 3, we quantify the impact of auctions by estimating the percentage change in various outcomes of interest if the firm wins the auction versus if it loses the auction. It compares these percentage changes induced by auctions in the actual economy ($1/\theta = 4.08$, first column) and the counterfactual economy in which firms have greater labor market power ($1/\theta' = 2.04$, second column).

In both economies, winning the auction induces the firm to increase total output, both to complete the procurement project and to produce additional private market output (due to crowd-in). The key difference in the economies is that the marginal cost of increasing output is greater when the labor supply curve is steeper. The firm chooses to increase employment by less and thus increase output less in the private

Changes Induced by Winning an Auction		
Labor Supply Elasticity $(1/\theta)$	**Actual** (4.08)	**Counterfactual** (2.04)
Employment	10.5%	9.2%
Wage	2.6%	4.5%
Wage bill	13.0%	13.8%
Output - Private Market	3.5%	1.7%
Output - Total	14.8%	14.6%
Firm Rents (Profits)	7.6%	7.6%
Worker Rents	15.0%	16.0%

Table 3: Changes Induced by Winning an Auction: Actual and Counterfactual Labor Markets

Notes: This table presents the percentage change in each outcome if the firm wins the auction versus if it loses the auction. It shows these percentage changes both in the actual economy (first column) and in the counterfactual economy in which the labor supply elasticity is half as great (second column).

market. When the labor supply curve is steeper, the firm must bid up wages more to achieve the increase in employment. Given that the firm hires less additional labor but at a greater wage, the wage bill may in theory be higher or lower when the labor supply curve is steeper. Empirically, we find that when the labor supply curve is steeper, total output increases by 14.6% (versus 14.8% in the actual economy), private market output increases by 1.7% (versus 3.5% in the actual economy), employment increases by 9.2% (versus 10.5% in the actual economy), the wage rate increases by 4.5% (versus 2.6% in the actual economy), and the wage bill increases by 13.8% (versus 13.0% in the actual economy).

We now examine the additional rents (incidence) captured by workers if the firm wins versus loses the procurement auction. If the firm wins, it wishes to increase output and must hire more labor. In the counterfactual economy, the firm must bid up wages more to achieve an increase in employment because the labor supply curve is steeper. This results in each incumbent worker capturing greater rents in the counterfactual economy from being inframarginal in the current firm. However, because labor is more expensive at the margin, the firm chooses to hire less additional labor in the counterfactual economy, which lowers the additional rents captured by workers in the counterfactual economy. Accounting for these opposing forces in the model, we find that if the firm wins the auction, workers earn 16% additional rents in the counterfactual economy, which is slightly greater than the 15% additional rents in the actual economy.

We also examine the additional rents (incidence) captured by the firm if it wins versus loses the procurement auction, which is equal to the gain in profits (equation 17). If it wins, the firm wishes to increase output. To do so, it hires additional labor. The marginal cost of labor is greater when the labor supply curve is steeper, which tends to decrease marginal profits. However, the firm also has greater market power and suppresses the number of new hires in order to markdown wages, which tends to increase marginal profits. Accounting for these opposing forces in the model, we find that the change in profits is nearly identical in the actual and counterfactual economies, with the firm increasing profits by about 8% if it wins the auction versus if it loses the auction.

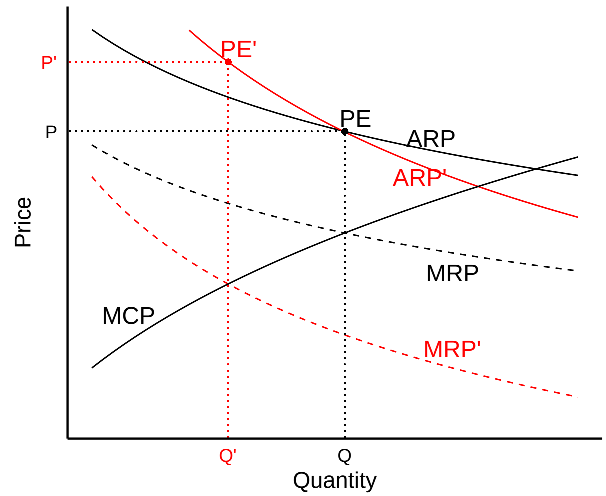

Figure 6: Illustration: Rotation of the Private Demand Curve

Notes: This diagram visualizes the counterfactual exercise of making the private product market less competitive. MCP denotes the marginal cost of production, MRP denotes marginal revenue product, ARP denotes average revenue product, and PE denotes the monopolistic product market equilibrium. Black colors denote the initial economy, and red colors denote the new economy after rotating the private demand curve through a compensated decrease in the demand elasticity, $1/\epsilon$.

7.2 The Importance of Imperfect Competition in the Private Product Market

Defining and Interpreting the Counterfactual

To study market power in the private product market, we consider a compensated rotation of the private product market demand curve (hereafter, demand curve) so that the demand elasticity, $1/\epsilon$, decreases while the initial equilibrium output and price combination remains feasible. In practice, this means that we first solve for the initial monopolistic product market equilibrium (PE), (Q_j^H, R_j^H), shift $1/\epsilon$ to $1/\epsilon'$, then compensate firm j for this decrease in the average revenue per unit of output by shifting the demand curve scale term p to p' so that (Q_j^H, R_j^H) is still on the demand curve.

Figure 6 provides an illustration of rotating the demand curve. It considers a fictional firm j, with private market output quantity on the x-axis and private market price on the y-axis. The initial equilibrium is in black, while the equilibrium after rotating the private demand curve is in red. The initial average revenue product (ARP, solid line) and its associated marginal revenue product (MRP, dashed line) are in black. The marginal cost of production curve (MCP) is also in black. To determine the equilibrium (PE), the monopolistic firm optimally chooses output to equate MCP and MRP, then chooses the highest feasible price at this quantity of output, which is on the ARP curve directly above the intersection of MCP and MRP.

The red lines in Figure 6 demonstrate how the equilibrium adjusts when the private demand curve is made "steeper" by lowering $1/\epsilon$ to $1/\epsilon$'. Lowing $1/\epsilon$ lowers the average revenue product, so we compensate the firm by increasing the multiplicate term p until the initial PE is on the new demand curve (that is, we ensure PE is on ARP'). The new marginal revenue product curve (MRP') is lower, so the point at which MCP equals MRP' is at a lower level of output (that is, Q' is less than Q on the x-axis). Since ARP' is above ARP at Q', it follows that PE' must also have a higher price than PE (P' is greater than P on the y-axis). Since the firm wishes to produce less, it needs to decrease inputs. To decrease labor, the firm lowers the wages it pays each employee, which also results in a lower wage bill. The relative price of labor to capital

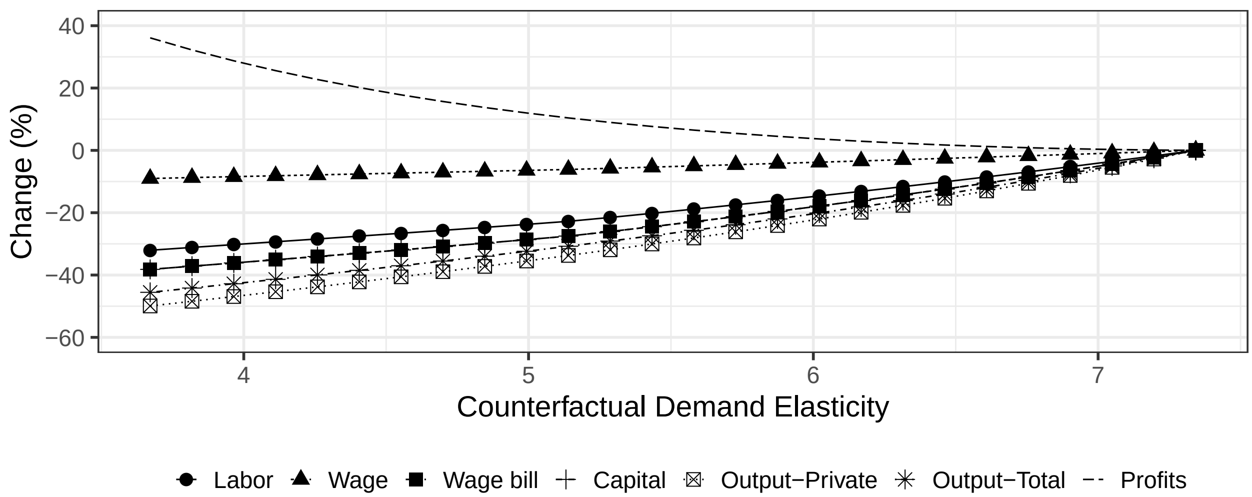

Figure 7: Counterfactual Rotation of the Private Demand Curve

Notes: This figure presents the counterfactual median values of labor, wages, the wage bill, capital, private market output, total output, and profits when the private product market becomes less competitive. It expresses these values as percentage changes relative to the actual economy.

is unchanged, so the firm reduces labor and capital in fixed proportion. Furthermore, since PE is a feasible choice of the firm, and the firm chooses a different point PE' to maximize profits, it must have higher profits with the counterfactual demand curve. Thus, the counterfactual exercise in Figure 6 always results in firm j producing less private market output, employing fewer workers, paying a lower wage to each employee, using less capital, and earning higher profits. Note that increasing returns to scale may partially reverse some of these effects, and we account for these returns to scale in the empirical application.

Characteristics of the Actual and Counterfactual Private Product Markets

In Figure 7, we present the counterfactual analysis where we increase the degree of imperfect competition by rotating the demand curve. We find that, as the firm gains more market power, it cuts production in the private market and marks-up the price. This allows it to increase profits. To reduce output, it employs fewer workers and pays a lower wage to each employee. These results are consistent with the predictions from the diagram in Figure 6. Because of the monotonic relationships across values of the demand elasticity, we focus on comparing the actual value ($1/\theta = 7.34$) and half of this amount ($1/\theta' = 3.67$).

The empirical estimates from our counterfactuals show that, when the demand elasticity is reduced by half, the firm produces 50% less output in the private market and 45% less total output. To do so, it employs 30% fewer workers. Wages decrease by 10% and the wage bill decreases by nearly 40%. Capital also decreases by 40%. Despite the large reduction in output, profits increase by nearly 40% in the counterfactual, as firms increasingly markup prices.

Incidence in Actual and Counterfactual Private Product Markets

In Table 4, we quantify the impact of auctions by estimating the percentage change in various outcomes of interest if the firm wins the auction versus if it loses the auction. It compares these percentage changes induced by auctions in the actual economy ($1/\epsilon = 7.34$, first column) and the counterfactual economy in

Changes Induced by Winning an Auction		
Private Demand Elasticity $(1/\epsilon)$	**Actual** (7.34)	**Counterfactual** (3.67)
Employment	10.5%	17.1%
Wage	2.6%	4.2%
Wage bill	13.0%	21.2%
Output - Private Market	3.5%	3.0%
Output - Total	14.8%	24.2%
Firm Rents (Profits)	7.6%	5.6%
Worker Rents	15.0%	27.0%

Table 4: Changes Induced by Winning an Auction: Actual and Counterfactual Private Product Demand

Notes: This table presents the percentage change in each outcome if the firm wins the auction versus if it loses the auction. It shows these percentage changes both in the actual economy (first column) and in the counterfactual economy in which the private demand elasticity is half as great (second column).

which firms have greater private product market power ($1/\epsilon' = 3.67$, second column). When the private market demand curve is steeper, the firm chooses to produce less in total so that it can markup prices. Thus, a smaller share of output is produced in the private market when the demand curve is steeper. Due to increasing returns to scale, winning the auction makes the firm wish to increase private market activity, both in the actual and counterfactual economies. But when the firm has more market power, this increase is mitigated by the firm choosing to keep output low so that it can markup prices. Because government output accounts for a greater share of total output when the demand curve is steeper, winning the contract gives a greater percent change in total output but a smaller percent change in private market output.

We now examine the additional rents (incidence) captured by workers and firms if the firm wins versus loses the procurement auction. Because the firm increases total output by a greater percent if it wins the auction in the counterfactual economy, it must also increase employment by a greater percent. To increase employment by a greater percent, it must bid up wages by a greater percent in the counterfactual economy. The greater percent increase in both wages and employment translates into a greater percent increase in rents captured by workers in the counterfactual economy. However, in the counterfactual economy, marginal revenue in the private market is lower, making it less profitable to increase private market output, which decreases the additional rents captured by firms in the counterfactual economy. Empirically, we find that if the firm wins the auction, workers earn 27% additional rents and firms earn 6% additional rents in the counterfactual economy, which is substantially greater than the 15% additional rents in the actual economy for workers but less than the 8% gain in rents for firms in the actual economy.

8 Conclusions

TBD.

References

Ackerberg, D. A., K. Caves, and G. Frazer (2015): "Identification properties of recent production function estimators," *Econometrica*, 83, 2411–2451.

Anagol, S. and T. Fujiwara (2016): "The runner-up effect," *Journal of Political Economy*, 124, 927–991.

Athey, S. and P. A. Haile (2007): "Nonparametric approaches to auctions," *Handbook of econometrics*, 6, 3847–3965.

Berger, D. W., K. F. Herkenhoff, and S. Mongey (2019): "Labor Market Power," Tech. rep., National Bureau of Economic Research.

Bhaskar, V., A. Manning, and T. To (2002): "Oligopsony and Monopsonistic Competition in Labor Markets," *The Journal of Economic Perspectives*, 16, 155–174.

Bhattacharya, D. (2015): "Nonparametric Welfare Analysis for Discrete Choice," *Econometrica*, 83, 617–649.

Boal, W. M. and M. R. Ransom (1997): "Monopsony in the labor market," *Journal of Economic Literature*, 35, 86–112.

Caldwell, S. and E. Oehlsen (2018): "Monopsony and the Gender Wage Gap: Experimental Evidence from the Gig Economy," *Massachusetts Institute of Technology Working Paper*.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.

Dube, A., A. Manning, and S. Naidu (2020): "Monopsony and employer mis-optimization explain why wages bunch at round numbers," *Unpublished manuscript*.

Ferraz, C., F. Finan, and D. Szerman (2015): "Procuring firm growth: the effects of government purchases on firm dynamics," Tech. rep., National Bureau of Economic Research.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, firm turnover, and efficiency: selection on productivity or profitability?" *American Economic Review*, 98, 394–425.

Garin, A. and F. Silverio (2019): "How Does Firm Performance Affect Wages? Evidence from Idiosyncratic Export Shocks," Tech. rep., Working Paper.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal nonparametric estimation of first-price auctions," *Econometrica*, 68, 525–574.

Hamermesh, D. S. (1999): "Changing inequality in markets for workplace amenities," *The Quarterly Journal of Economics*, 114, 1085–1123.

Howell, S. T. and J. D. Brown (2020): "Do Cash Windfalls Affect Wages? Evidence from R&D Grants to Small Firms," Working Papers 20-06, Center for Economic Studies, U.S. Census Bureau.

Jofre-Bonet, M. and M. Pesendorfer (2003): "Estimation of a dynamic auction game," *Econometrica*, 71, 1443–1489.

Kline, P., N. Petkova, H. Williams, and O. Zidar (2019): "Who profits from patents? rent-sharing at innovative firms," *The Quarterly Journal of Economics*, 134, 1343–1404.

Lamadon, T., M. Mogstad, and B. Setzler (2019): "Imperfect Competition, Compensating Differentials and Rent Sharing in the US Labor Market," Tech. rep., National Bureau of Economic Research.

Levinsohn, J. and A. Petrin (2003): "Estimating production functions using inputs to control for unobservables," *The review of economic studies*, 70, 317–341.

Luo, Y. (2019): "Unobserved heterogeneity in auctions under restricted stochastic dominance," *Journal of Econometrics*.

Maestas, N., K. J. Mullen, D. Powell, T. Von Wachter, and J. B. Wenger (2018): "The Value of Working Conditions in the United States and Implications for the Structure of Wages," Working Paper 25204, National Bureau of Economic Research.

Manning, A. (2003): *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press.

Mas, A. and A. Pallais (2017): "Valuing alternative work arrangements," *American Economic Review*, 107, 3722–59.

Maskin, E. and J. Riley (1984): "Optimal auctions with risk averse buyers," *Econometrica: Journal of the Econometric Society*, 1473–1518.

Milgrom, P. R. and R. J. Weber (1982): "A theory of auctions and competitive bidding," *Econometrica: Journal of the Econometric Society*, 1089–1122.

Pierce, B. (2001): "Compensation Inequality," *The Quarterly Journal of Economics*, 116, 1493–1525.

Rosen, S. (1986): "The theory of equalizing differences," *Handbook of labor economics*, 1, 641–692.

Sokolova, A. and T. Sorensen (2018): "Monopsony in labor markets: A meta-analysis," .

Suárez Serrato, J. C. and O. Zidar (2016): "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms," *American Economic Review*, 106, 2582–2624.

Wiswall, M. and B. Zafar (2017): "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 133, 457–507.

A Sample and Variable Definitions

All firm-level variables are constructed from annual business tax returns over the years 2001-2015: C-Corporations (Form 1120), S-Corporations (Form 1120-S), and Partnerships (Form 1065). Worker-level variables are constructed from annual tax returns over the years 2001-2015: Direct employees (Form W-2) and independent contractors (Form 1099).

Tax Return Variable Definitions:

- **Earnings:** Reported on W-2 box 1 for each Taxpayer Identification Number (TIN). Each TIN is de-identified in our data.
- **Employer:** The Employer Identification Number (EIN) reported on W-2 for a given TIN. Each EIN is de-identified in our data.
- **Employees:** Number of workers matched to an EIN in year t from Form W-2 with annual earnings above the annualize full-time minimum wage and where the EIN is this worker's highest-paying employer.
- **Wage bill:** Total earnings among employees in year t.
- **Contracted worker:** Number of 1099-MISC individuals matched to an EIN in year t.
- **Employees, broader measure:** Total number of W-2 workers plus 1099 contracted workers matched to an EIN in year t.
- **Wage bill, broader measure:** Sum of earnings for the broader sample of employees matched to an EIN in year t.
- **NAICS Code:** The NAICS code is reported on line 21 on Schedule K of Form 1120 for C-corporations, line 2a Schedule B of Form 1120S for S-corporations, and Box A of form 1065 for partnerships. We consider the first three digits to be the industry. We code invalid industries as missing.
- **Sales:** Line 1 of Form 1120 for C-Corporations, Form 1120S for S-Corporations, and Form 1065 for partnerships. Also referred to as gross revenues.
- **Intermediate Costs:** Line 2 of Form 1120 for C-Corporations, Form 1120S for S-Corporations, and Form 1065 for partnerships. Also referred to as cost of goods sold.
- **EBITD:** We follow Kline et al. (2019) in defining Earnings Before Interest, Taxes, and Depreciation (EBITD) as the difference between total income and total deductions other than interest and depreciation. Total income is reported on Line 11 on Form 1120 for C-corporations, Line 1c on Form 1120S for S-corporations, and Line 1c on Form 1065 for Partnerships. Total deductions other than interest and depreciation are computed as Line 27 minus Lines 18 and 20 on Form 1120 for C-corporations, Line 20 minus Lines 13 and 14 on Firm 1120S for S-corporations, and Line 21 minus Lines 15 and 16c on Form 1065 for partnerships.
- **Sales net of expenditures:** This is sales minus intermediate costs minus wage bill.

Procurement Auction Variable Definitions:

- **Bid:** The dollar value submitted by the firm as a price at which it would be willing to complete the procurement project.
- **Auction winner:** A firm is an auction winner if it placed the lowest bid in a procurement auction.
- **Amount of winnings:** Bid placed by the winner in each auction.
- **Year of first win:** First year in which the firm is an auction winner. To account for left-censoring, we do not define a win as a "first win" unless there were at least two observed years of data during which the firm could have won and did not win an auction. For example, if a state provided auction records for 2001-2015, and a firm is first observed winning in 2001 or 2002, we do not consider this firm a first-time winner, but if the firm is first observed winning in 2003 or later, we consider it a first time winner.

Firm sample definitions:

- **Baseline sample:** A firm that files tax form 1120, 1120-S, or 1065 is considered part of the baseline sample centered around auction cohort c if it is observed bidding in an auction in year c.
- **Sample of close bidders (cardinal):** A firm in the baseline sample at c is also in the sample of close bidders (cardinal) if, in at least one auction in year c, its bid was less than $\mathcal{T} \times 100$ percent greater than the bid of the winner, where $\mathcal{T}$ is referred to as the loss margin. For example, if $\mathcal{T} = 0.10$, the firm must have placed a bid within 10% of the winner's bid in at least one auction at c.
- **Sample of close bidders (ordinal):** A firm in the baseline sample at c is also in the sample of close bidders (ordinal) if, in at least one auction in year c, it had at least the R lowest bid in the auction, where R is referred to as the bid rank. For example, if $R = 3$, the firm must have been among the 3 lowest bidders (inclusive of the winner) in at least one auction at c.
- **Sample of non-winners:** A firm in the baseline sample at c that does not win an auction before or during c is called a non-winner if it continues to not win any auctions until at least relative time $t \geq 4$. For example, if $c = 2005$, then a non-winner must not win its first auction until at least 2009.

Worker sample definitions:

- **Main sample:** A worker is considered part of the main sample at c if the worker's highest-paying firm at c on Form W-2 is in the baseline sample of firms and the W-2 wage payments from that firm are greater than $15,000 in 2015 USD. We also restrict to workers aged 25-60.
- **Broader sample:** A worker is considered part of the broader sample at c if the worker receives a W-2 or 1099 form associated with a firm in the baseline sample of firms at c.
- **Stayers:** A worker is a stayer for $2k + 1$ years at firm j in the baseline sample of firms at c if the worker's highest-paying W-2 firm is the same firm during each time period in $(c - k, ..., c + k)$ and the W-2 wage payments from that firm in each year are greater than $15,000 in 2015 USD. We also restrict to workers aged 25-60.
- **New Hires:** A worker is a new hire at firm j in year t if the worker's highest-paying W-2 employer in year t was firm j and highest-paying W-2 employer in year $t - 1$ was firm $j' \neq j$, where the worker recieved W-2 wage payments greater than $15,000 in 2015 USD from j' in $t - 1$ as well as from j in t. We also restrict to workers aged 25-60.

B Additional Tables and Figures

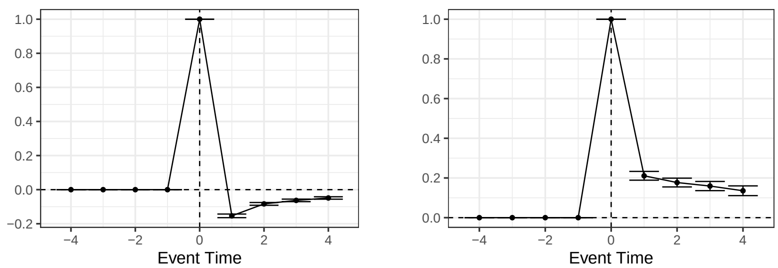

(a) Indicator for First-time Winning (b) Indicator for Any Wins

Figure A.1: Visualizing the Difference-in-differences Research Design

Notes: We only consider treatment and control units that place a bid in the cohort year that have never won an auction before. All results include firm fixed effects. The omitted relative time is −2. 90% confidence intervals are displayed.

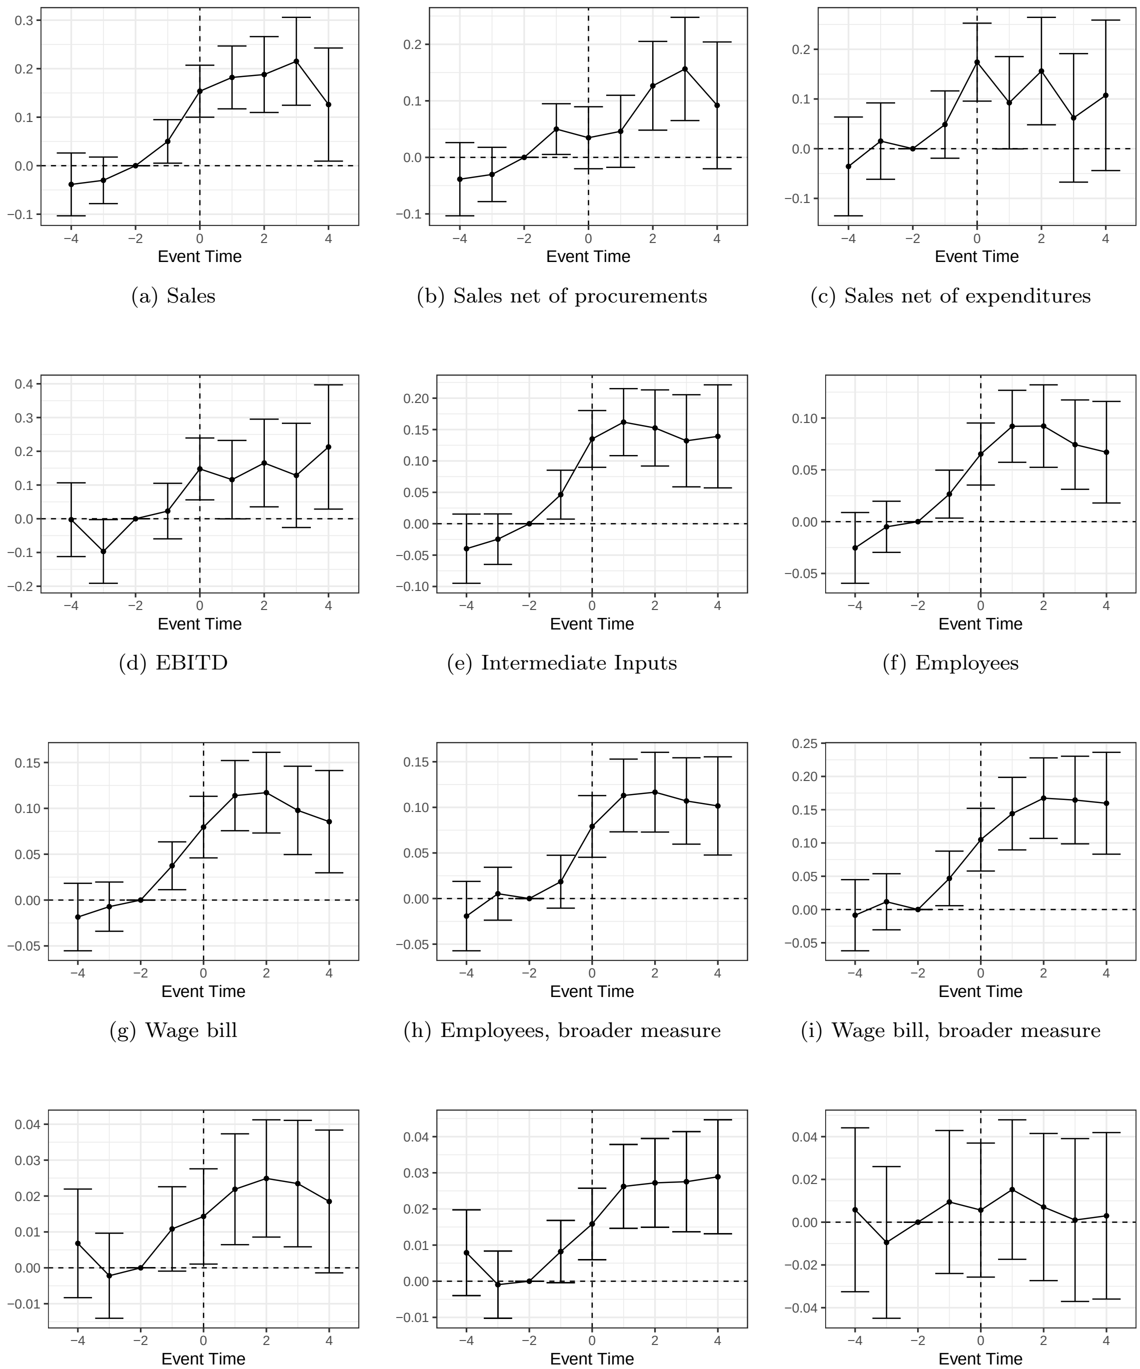

(a) Sales (b) Sales net of procurements (c) Sales net of expenditures

(d) EBITD (e) Intermediate Inputs (f) Employees

(g) Wage bill (h) Employees, broader measure (i) Wage bill, broader measure

(j) Mean Wage, All Workers (k) Mean Wage, Stayers (l) Mean Previous Earnings, New Hires

Figure A.2: Difference-in-differences Estimates at Annual Frequency

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. The control units are those firms that place a bid in a procurement auction in the same year that the reference treatment cohort wins. The omitted relative time is -2. 90% confidence intervals are displayed, clustering on firm.

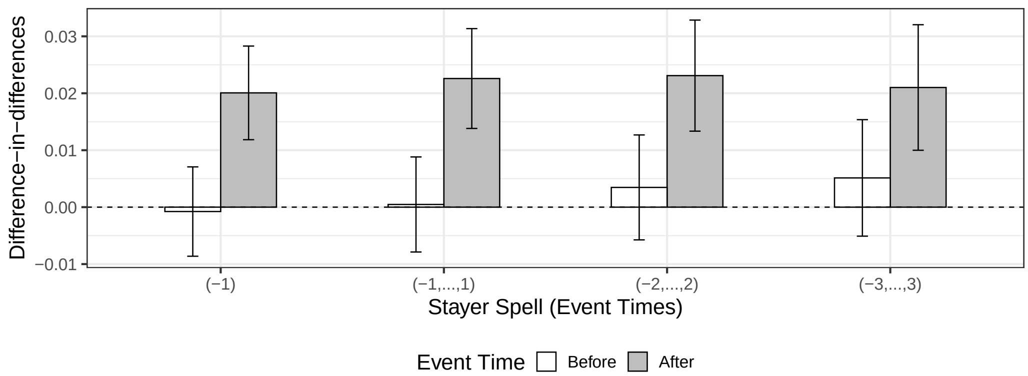

(a) Robustness by Stayer Spell

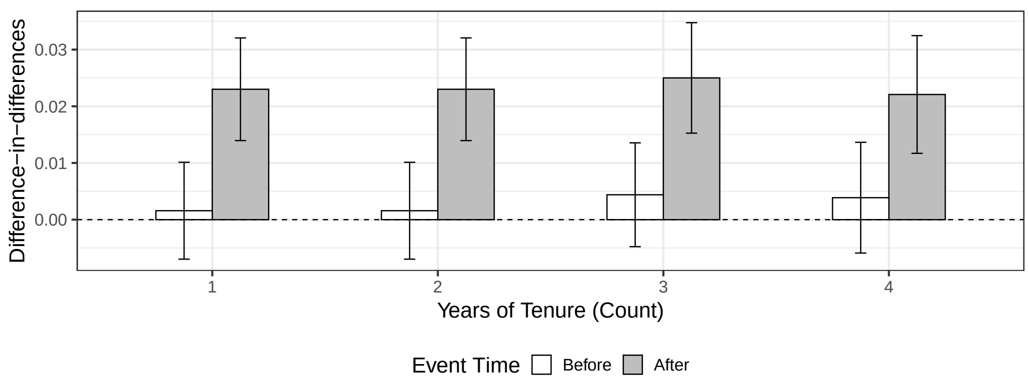

(b) Robustness by Tenure Length

Figure A.3: Robustness: Stayer and Tenure Samples

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. "Before" refers to relative times {-4,-3,-2} and "After" refers to relative times {0,1,2}. The Baseline sample restricts the control units to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. The omitted relative time is -2. 90% confidence intervals are displayed, clustering on firm. Subfigure (a) varies the window in which the worker must have been employed by the bidding firm, where (-2,...,2) is treated as the baseline definition of stayers. Subfigure (b) varies the window over which the worker must have been employed prior to the auction bid, e.g., tenure of -4 means that the worker was employed from relative time -4 until at least relative time 0.

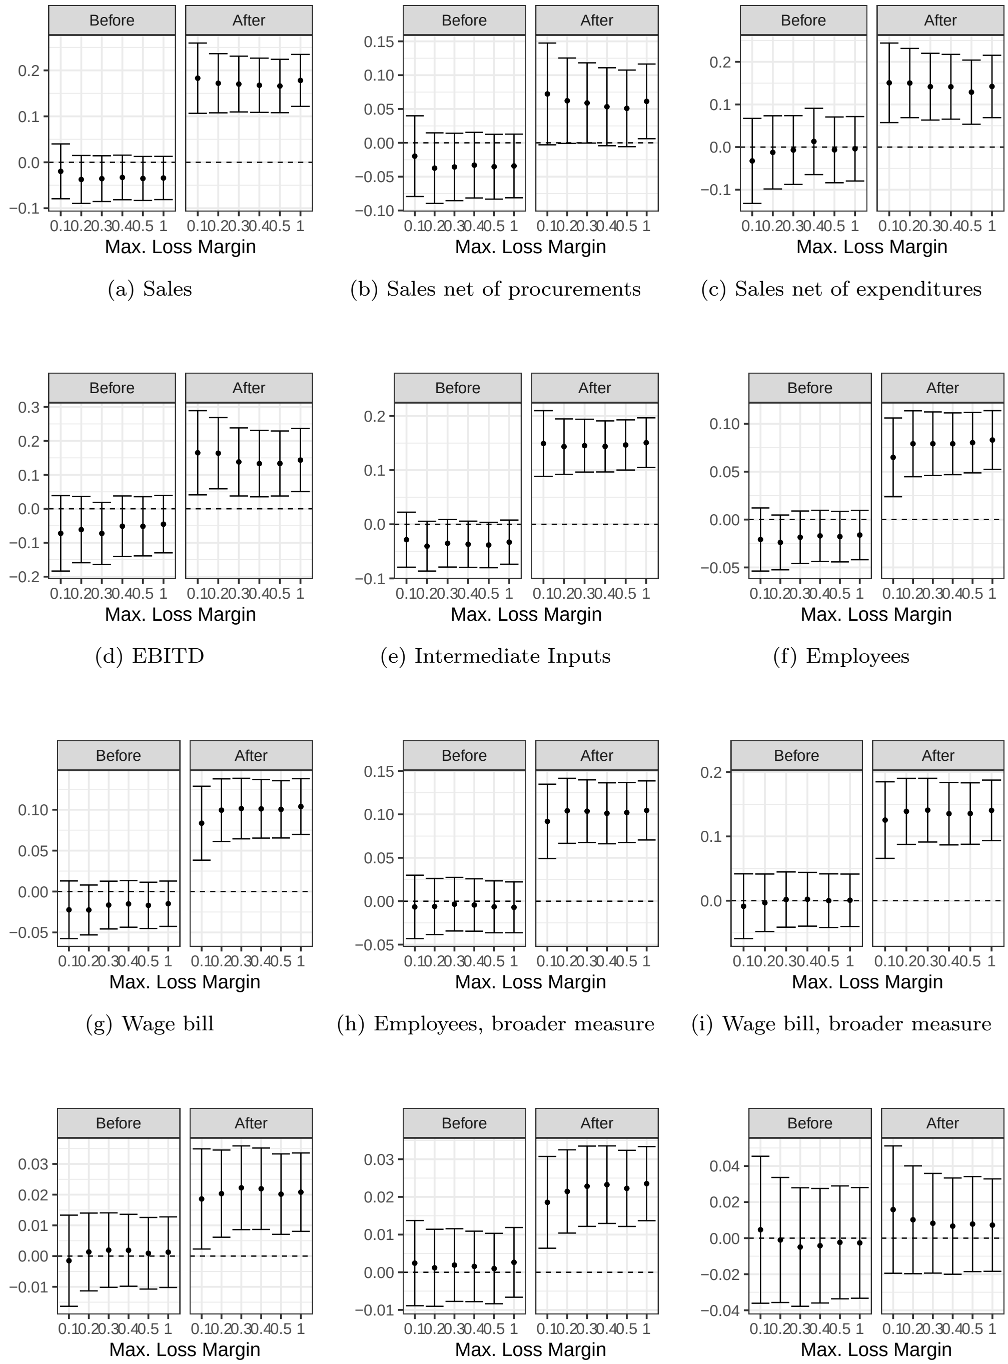

(a) Sales (b) Sales net of procurements (c) Sales net of expenditures

(d) EBITD (e) Intermediate Inputs (f) Employees

(g) Wage bill (h) Employees, broader measure (i) Wage bill, broader measure

(j) Mean Wage, All Workers (k) Mean Wage, Stayers (l) Mean Previous Wage, New Hires

Figure A.4: Restricting the Control Group's Bid Loss Margin

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. "Before" refers to relative times {-4,-3,-2} and "After" refers to relative times {0,1,2}. The Baseline sample restricts the control units to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. We then restrict the control group to firms whose bid loss margin was lower than the number displayed on the x-axis.

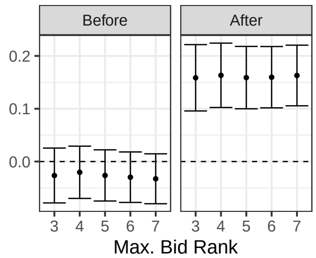

(a) Sales

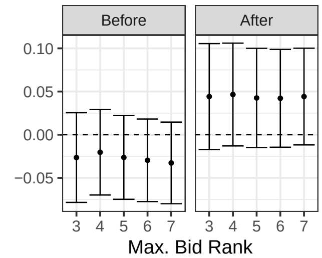

(b) Sales net of procurements

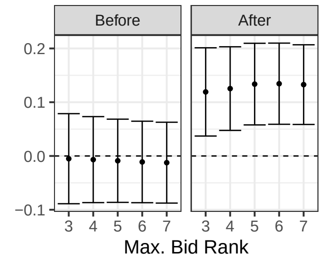

(c) Sales net of expenditures

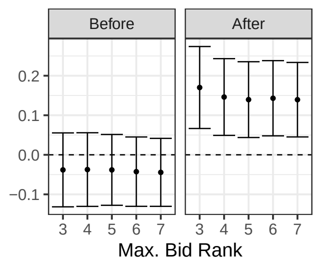

(d) EBITD

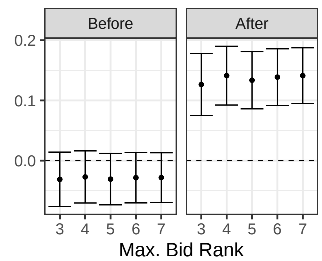

(e) Intermediate Inputs

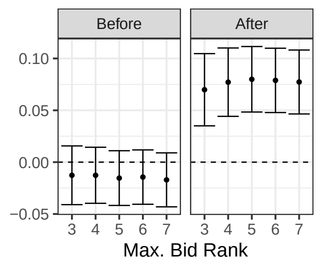

(f) Employees

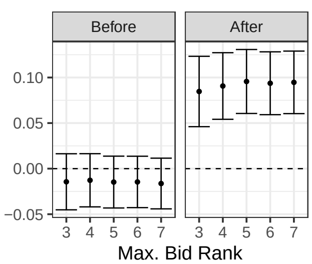

(g) Wage bill

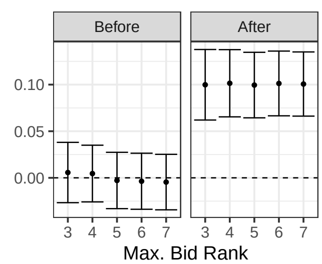

(h) Employees, broader measure

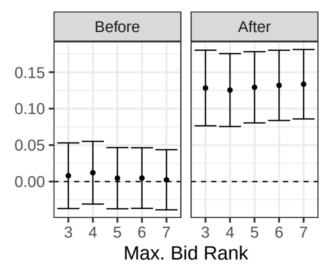

(i) Wage bill, broader measure

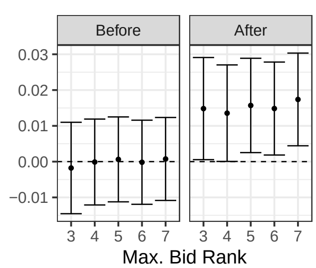

(j) Mean Wage, All Workers

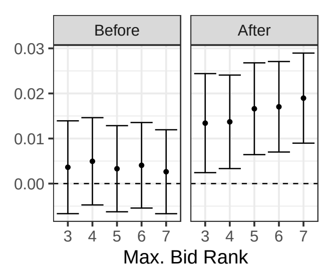

(k) Mean Wage, Stayers

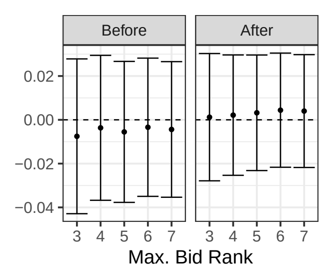

(l) Mean Previous Wage, New Hires

Figure A.5: Restricting the Control Group's Bid Ranks

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. "Before" refers to relative times {-4,-3,-2} and "After" refers to relative times {0,1,2}. The Baseline sample restricts the control units to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. We then restrict the control group to firms whose bid rank was less than or equal to the number displayed on the x-axis.

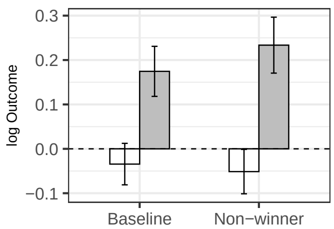

(a) Sales

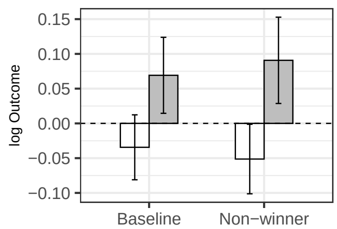

(b) Sales net of procurements

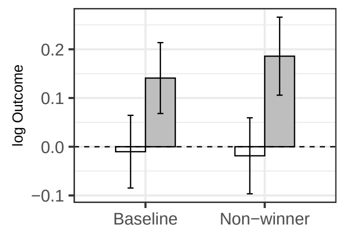

(c) Sales net of expenditures

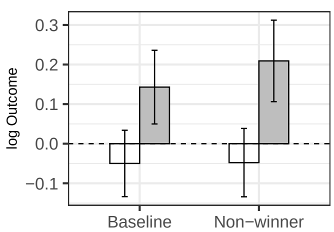

(d) EBITD

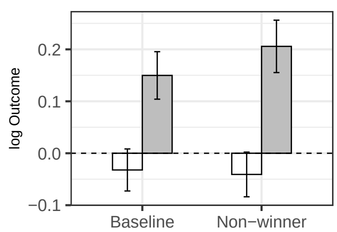

(e) Intermediate Costs

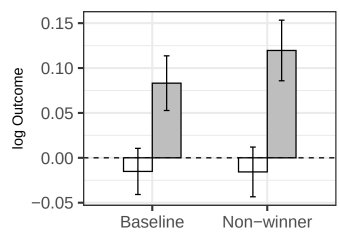

(f) Employees

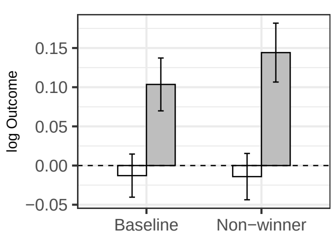

(g) Wage bill

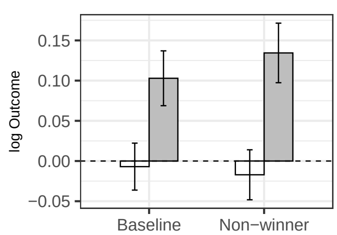

(h) Employees, broader measure

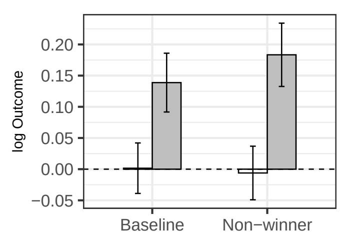

(i) Wage bill, broader measure

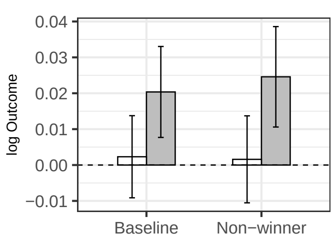

(j) Mean earnings, all employees

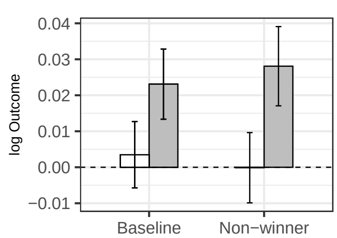

(k) Mean earnings, stayers

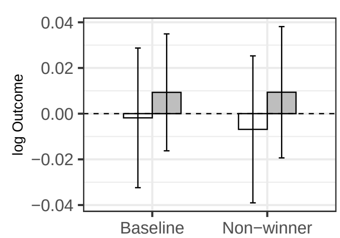

(l) Mean past earnings, new hires

Figure A.6: Robustness: Non-winner Control Firms

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. "Before" refers to relative times {-4,-3,-2} and "After" refers to relative times {0,1,2}. "Baseline" restricts the control firms to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. "Non-winner" further restricts the control firms to those that do not win an auction until at least relative time 4. The omitted relative time is -2. 90% confidence intervals are displayed, clustering on firm.

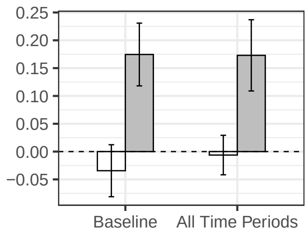

(a) Sales

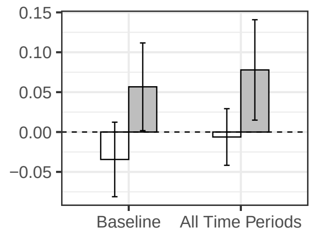

(b) Sales net of procurements

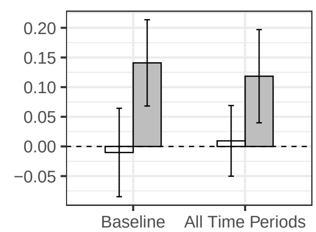

(c) Sales net of expenditures

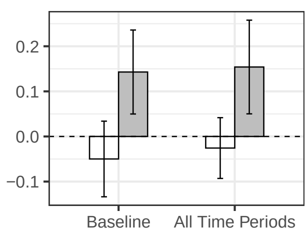

(d) EBITD

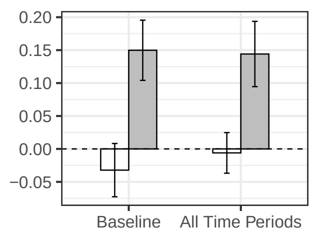

(e) Intermediate Costs

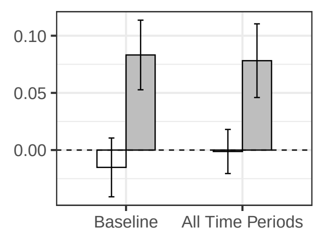

(f) Employees

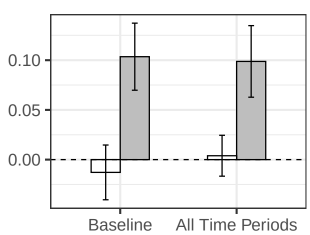

(g) Wage bill

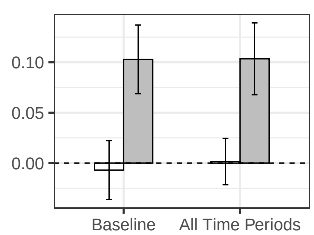

(h) Employees, broader measure

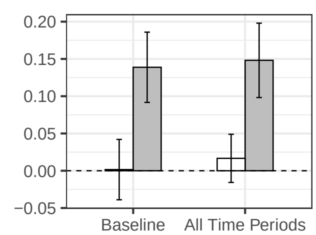

(i) Wage bill, broader measure

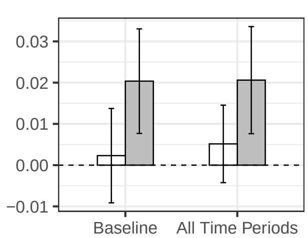

(j) Mean earnings, all employees

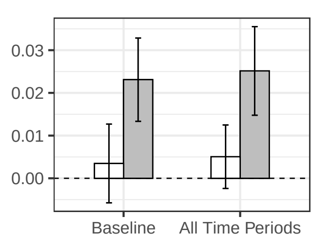

(k) Mean earnings, stayers

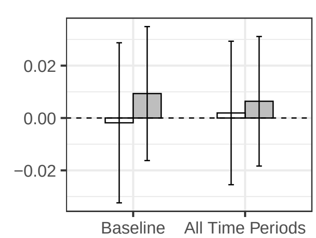

(l) Mean past earnings, new hires

Figure A.7: Robustness: All Time Periods

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. “Before” refers to relative times {-4,-3,-2} and “After” refers to relative times {0,1,2}. In “Baseline”, “Before” refers to relative times {-4,-3,-2} and “After” refers to relative time {0,1,2}. In “All Time Periods”, “Before” refers to relative times {-4,-3,-2,-1} and “After” refers to relative time {0,1,2,3,4}.The omitted relative time is -2. 90% confidence intervals are displayed, clustering on firm.

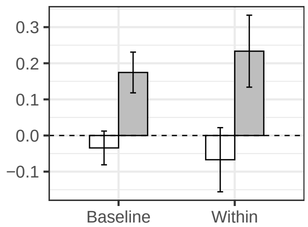

(a) Sales

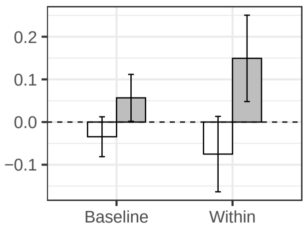

(b) Sales net of procurements

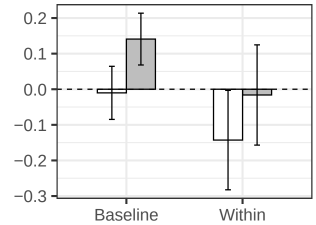

(c) Sales net of expenditures

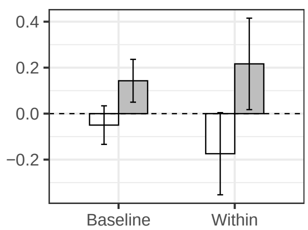

(d) EBITD

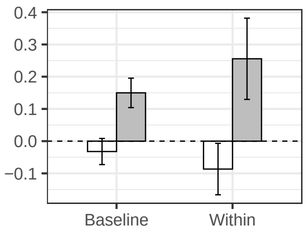

(e) Intermediate Costs

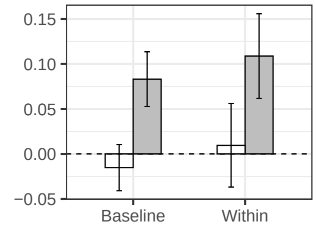

(f) Employees

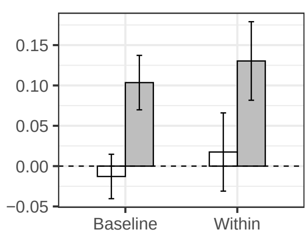

(g) Wage bill

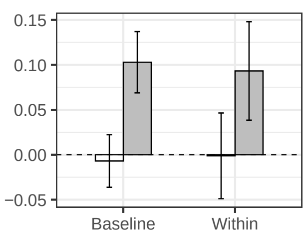

(h) Employees, broader measure

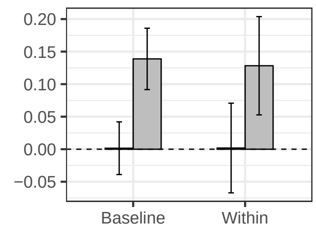

(i) Wage bill, broader measure

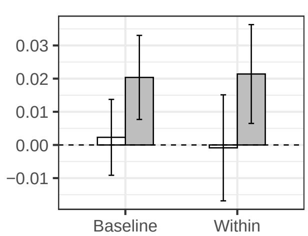

(j) Mean earnings, all employees

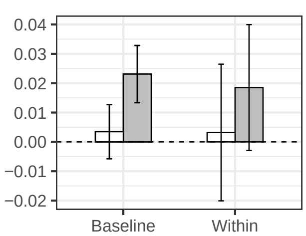

(k) Mean earnings, stayers

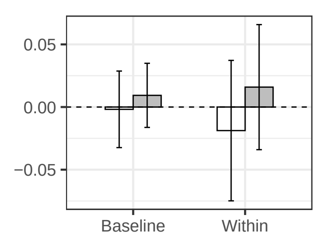

(l) Mean past earnings, new hires

Figure A.8: Robustness: Within Auction Specification

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. "Before" refers to relative times {-4,-3,-2} and "After" refers to relative times {0,1,2}. "Baseline" restricts the control firms to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. "Within Auction" further restricts the control firms to those that placed bids within the same auction as the winner, then estimates the difference-in-differences specification separately by auction and averages the estimates across auctions. In the "Within Auction" specification, standard errors are calculated using the block bootstrap, where a block is taken to be an auction. The omitted relative time is -2. 90% confidence intervals are displayed, clustering on firm.

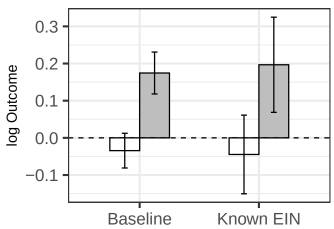

(a) Sales

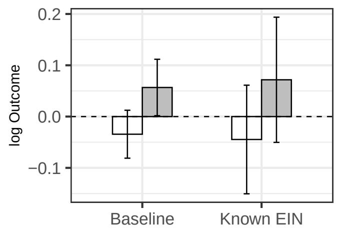

(b) Sales net of procurements

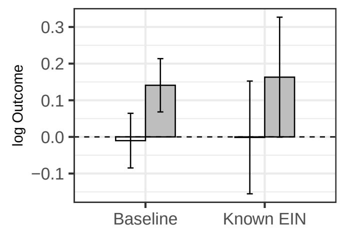

(c) Sales net of expenditures

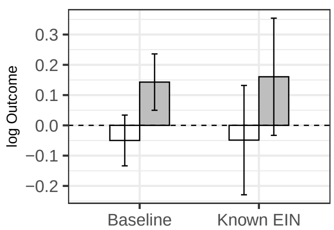

(d) EBITD

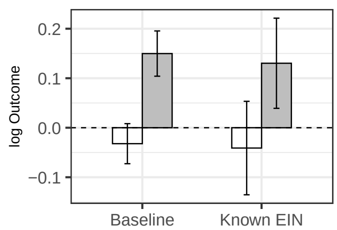

(e) Intermediate Costs

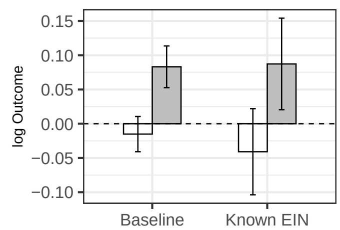

(f) Employees

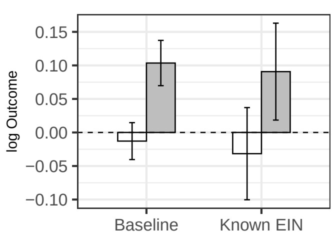

(g) Wage bill

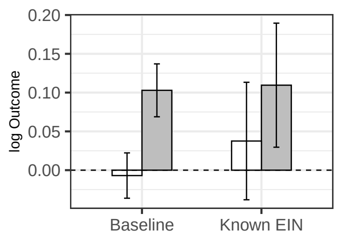

(h) Employees, broader measure

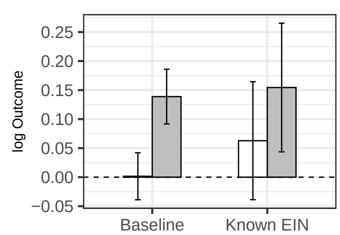

(i) Wage bill, broader measure

Figure A.9: Robustness: 5 States with Known EIN

Notes: This figure presents average treatment effect on the treated estimates using the difference-in-differences specification defined in the text. All results include firm fixed effects. “Before” refers to relative times {-4,-3,-2} and “After” refers to relative times {0,1,2}. “Baseline” restricts the control firms to those that place a bid in a procurement auction in the same year that the reference treatment cohort wins. “Known EIN” further restricts the control firms to those that firms for which we were provided the EIN by the state DOT, thus allowing us to link records exactly rather than using a fuzzy matching algorithm. Earnings per worker outcomes are omitted due to the very large standard errors when considering only 5 states. The omitted relative time is -2. 90% confidence intervals are displayed, clustering on firm.

C Additional Model Derivations

C.1 Second stage: the optimal input choices

This subsection solves the second stage assuming a monopolistic competitive product market ($\epsilon > 0$). We consider the winning firm and the losing firm separately. Substituting the labor supply curve, revenue curve, and optimal materials choices from above, the losing firm's problem can be represented as,

$$\max_{L_j} \quad p[\Omega_j L_j^{\rho}]^{1-\epsilon} - \Theta_0 L_j^{1+\theta} - p_M m(L_j, \Omega_j)$$

where $m(L_j, \Omega_j) = \Omega_j L_j^{\rho}/\beta_M$. This leads to an interior solution, $L_{0,j}$, that satisfies the FOC,

$$\frac{\partial \pi_{0,j}}{\partial L_j} = p\Omega_j^{1-\epsilon}(1-\epsilon)\rho L_{0,j}^{(1-\rho)\epsilon-(1-\rho)-\epsilon} - \Theta_0(1+\theta)L_{0,j}^{\theta} - \frac{p_M}{\beta_M}\Omega_j \rho L_{0,j}^{-(1-\rho)} = 0$$

Multiplying all terms by $L_{0,j}^{1-\rho}$, the first term in the FOC converges to infinity when labor input approaches zero. On the other hand, the other two terms in the FOC reduces to zero. Thus, a corner solution cannot satisfy the equation, so $L_{0,j} > 0$.

Similarly, the winning firm's problem can be represented as:

$$\max_{L_j : \Omega_j L_j^{\rho} \geq Q_{1,j}} \quad p[\Omega_j L_j^{\rho} - \bar{Q}^G]^{1-\epsilon} - \Theta_j L_j^{1+\theta} - p_M m(L_j, \Omega_j)$$

Taking the derivative with respect to labor gives the FOC:

$$\frac{\partial \pi_{1,j}}{\partial L_j} = p\Omega_j(1-\epsilon)\rho[\Omega_j L_j^{\rho} - \bar{Q}^G]^{-\epsilon} L_j^{-(1-\rho)} - \Theta_j(1+\theta)L_j^{\theta} - \frac{p_M}{\beta_M}\Omega_j \rho L_j^{-(1-\rho)} = 0$$

As evident from this expression, if the firm wins an auction, it will still choose a positive amount of private market output. To see this, note that $[\Omega_j L_j^{\rho} - \bar{Q}^G]^{-\epsilon} = [Q_j - \bar{Q}^G]^{-\epsilon}$ tends to infinity when Q_j approaches $\bar{Q}^G$. Meanwhile, the second term converges to a constant and the third term is a constant. That is, the marginal benefit must be larger than the marginal cost for sufficiently small production in the private sector Q_j^H. This ensures an interior solution also if the firm wins the auction. Furthermore, it is always true that $\frac{\partial \pi_{1,j}}{\partial L_j}|_{L_j = L_{0,j}} > 0$. Thus, total production will be larger if the firm wins the auction than if it loses.

Lastly, it is interesting to consider if winning a procurement project will lead a firm to produce more in the private market (crowd in) or less (crowd out). First, the losing firm's FOC implies that $\frac{\partial \pi_{0,j}}{\partial L_j} L_{0,j}^{(1-\rho)} = p_0 \Omega^{1-\epsilon}(1-\epsilon)\rho L_{0j}^{-\rho\epsilon} - \theta_0(1+\theta)L_{0j}^{\theta+(1-\rho)} - \frac{p_M}{\beta_M}\Omega\rho = 0$. Second, we check the marginal profits of the winner when the total output amount is at $\bar{Q}^G + Q_0^{H*}$, where $\bar{Q}^G$ is the size of the DOT project and Q_0^{H*} is a losing firm's output in the private market. The winner wants to produce more than this hypothetically amount $\bar{Q}^G + Q_0^{H*}$ if the derivative (i.e., the marginal profit) is positive, and less otherwise. Let the corresponding labor choice be $L_{1,j}$ such that $\Omega L_{1,j}^{\rho} - \bar{Q}^G = Q_0^{H*} = \Omega L_{0j}^{\rho}$. Marginal profits for the firm if it wins are,

$$\frac{\partial \pi_{1,j}}{\partial L_j}|_{L_j = L_{1,j}} = p_0 \Omega(1-\epsilon)\rho(Q_0^{H*})^{-\epsilon} L_{1,j}^{\rho-1} - \theta_0(1+\theta)L_{1,j}^{\theta} - \frac{p_M}{\beta_M}\Omega\rho L_{1,j}^{\rho-1}$$

which has the same sign as,

$$p_0\Omega(1-\epsilon)\rho(Q_0^{H*})^{-\epsilon} - \theta_0(1+\theta)L_{1,j}^{\theta+1-\rho} - \frac{p_M}{\beta_M}\Omega\rho = p_0\Omega^{1-\epsilon}(1-\epsilon)\rho L_{0j}^{-\rho\epsilon} - \theta_0(1+\theta)L_{1,j}^{\theta+1-\rho} - \frac{p_M}{\beta_M}\Omega\rho = \theta_0(1+\theta)(L_{0j}^{\theta+1-\rho} - L_{1,j}^{\theta+1-\rho})$$

The second equation applies the above equation $\frac{\partial \pi_{0,j}}{\partial L_j} L_{0,j}^{(1-\rho)} = 0$. The definition of $L_{1,j}$ implies that $L_{1,j} > L_{0j}$. Therefore, $\frac{\partial \pi_{1,j}}{\partial L_j}|_{L_j = L_{1,j}} < 0$ if $\theta + 1 - \rho > 0$, and $\frac{\partial \pi_{1,j}}{\partial L_j}|_{L_j = L_{1,j}} > 0$ otherwise. Therefore, winning a DOT project crowd out private projects when $1 + \theta > \rho$ and crowd in them otherwise. Intuitively, winning a DOT project increases the total output level. This leads to a higher wage rate due to an increasing labor supply curve, as well as a larger scale of production. When the benefits, i.e., the return to scale, is smaller than the cost, i.e., more expensive labor, there is crowd out effects. Otherwise, there is a crowd in effect.

C.2 Proof of Worker Rents Expression

Defining $u = l, dv = dW$, we can calculate this integral using integration by parts:

$$
\begin{aligned}
\mathbb{E}[EV] &= \int_{W_1^0}^{W_1^1} l(W)dW \\
&= [l(W)W]_{W_1^0}^{W_1^1} - \int_{W_1^0}^{W_1^1} \frac{dl}{dW} W dW \\
&= l(W_1^1)W_1^1 - l(W_1^0)W_1^0 - \int_{W_1^0}^{W_1^1} \frac{dl}{dW} W dW \\
&= \underbrace{l(W_1^0)\Delta W}_{\text{inframarginals}} + \underbrace{\int_{W_1^0}^{W_1^1} \left(W_1^1 - W\right) \frac{dl}{dW} dW}_{\text{marginals}}
\end{aligned}
$$

where $\Delta W \equiv W_1^1 - W_1^0$.

Following Lamadon et al. (2019), define $\omega \equiv \frac{W}{W_1^1}$ so that $\frac{d\omega}{dW} = \frac{1}{W_1^1}$. Next, $l(\omega_j W_1^1) = \frac{\left(\omega_j W_1^1\right)^{1/\theta} g_j}{\sum_j \left(\omega_j W_1^1\right)^{/\theta} g_j} = \omega_j^{1/\theta} l(W_1^1)$. Thus, $\frac{dl}{dW} dW = \frac{dl}{d\omega} d\omega$. Moreover, $\frac{dl}{d\omega} = \frac{\partial \omega^{1/\theta}}{\partial \omega} l(W_1^1)$. Then

$$
\begin{aligned}
\mathbb{E}[EV] &= l(W_1^0)\Delta W + W_1^1 \int_{\frac{W_1^0}{W_1^1}}^{1} (1-\omega) \frac{dl}{d\omega} d\omega \\
&= l(W_1^0)\Delta W + W_1^1 l(W_1^1) \int_{\frac{W_1^0}{W_1^1}}^{1} (1-\omega) \frac{\partial \omega^{1/\theta}}{\partial \omega} d\omega \\
&= l(W_1^0)\Delta W + W_1^1 l(W_1^1) \left[\left[(1-\omega)\, \omega^{1/\theta} \right]_{\frac{W_1^0}{W_1^1}}^{1} + \int_{\frac{W_1^0}{W_1^1}}^{1} \omega^{1/\theta} d\omega \right] \\
&= l(W_1^0)\Delta W + W_1^1 l(W_1^1) \left[-\left(1 - \frac{W_1^0}{W_1^1}\right) \left(\frac{W_1^0}{W_1^1}\right)^{1/\theta} + \left[\frac{1}{1+1/\theta} - \frac{1}{1+1/\theta} \left(\frac{W_1^0}{W_1^1}\right)^{1+1/\theta} \right] \right] \\
&= l(W_1^0)\Delta W + \left[\frac{W_1^1 l(W_1^1)}{1+1/\theta} - l(W_1^0)W_1^1 + \frac{1/\theta}{1+1/\theta} W_1^0 l(W_1^0) \right] \\
&= \frac{W_1^1 l(W_1^1)}{1+1/\theta} - \frac{W_1^0 l(W_1^0)}{1+1/\theta}
\end{aligned}
$$

C.3 Identification of the rest of the model parameters

We obtain the size of the labor market $\bar{\theta}$ by taking the expectation of the log wage equation:

$$\log \bar{\theta} = \mathbb{E}[\log B_j - (1+\theta) \log L_j]$$

using that the residual $\log g_j$ is normalized to mean zero.

Given ρ, we can easily estimate the return to materials from the Leontief FOC: $\Omega_j L_j^{\rho} = \beta_M M_j$, which implies,

$$\log(\beta_M/p_M) = \mathbb{E}[\rho \log L_j - \log(p_M M_j)],$$

where we have normalized $\log \Omega_j$ to have zero mean. Combining these parameters, we can then construct TFP for every firm in our data:

$$\log \Omega_j = \log(\beta_M/p_M) + \log(p_M M_j) - \rho \log L_j.$$

Moreover, since we obtained $\log p + (1-\epsilon)\log(\beta_M/p_M)$ above, we can then estimate $\log p$ as,

$$\log p = E[\log R_j] - (1-\epsilon)(\log \beta_M/p_M + \mathbb{E}[\log(p_M M_j)]).$$

For Online Publication

D Online Appendix: Data Sources and Sample Selection

This Appendix describes the procedure adopted to match the bidders in our BidX sample to the tax data. For a subset of bidders the Employer Identification Number (EIN) is available in the BidX data, providing a unique identifier for the matching. For those observations an exact matching can be performed. We refer to this subset of perfect matches as the *training data*. In any other case, we rely on the fuzzy matching algorithm described below.

D.1 Algorithm

The procedure takes advantage of some regularities in the denomination of firms and common abbreviations to improve the quality of matching. Furthermore, in order to properly distinguish different branches of the same company, additional information on value added or state will be used.

Overview of denominations

Generally, a business name consists of three parts: a distinctive part, a descriptive part, and a legal part.[1] The distinctive part is named by the business owner and is usually required by governments to be "*substantially different*"[2] from any other existing name. The descriptive part describes what the business does, or its sector. Finally, the legal part refers to the business structure of a corporation. For example, for the name "Rogers Communications Inc.", "Rogers" is the distinctive part, "Communication" is the descriptive part, and "Inc." is the legal part. Most of the discrepancies of company names between different sources arise from the descriptive and the legal parts, since they are more subject to be abbreviations or common synonyms.

The legal part of corporations names takes a fairly small number of denominations, therefore can be identified using a properly constructed dictionary and treated separately. Conversely, disentangling the distinctive and the descriptive parts is not as straightforward. However, conventionally, the descriptive

[1] Although there is no specific regulations that demands such composition, it is in alignment with naming convention and government guidelines. `https://www.ic.gc.ca/eic/site/cd-dgc.nsf/eng/cs01070.html`

[2] An example would be California Code of Regulations for business entities. `https://www.sos.ca.gov/administration/regulations/current-regulations/business/business-entity-names/#section-21000`

part follows the distinctive one within the string. This observation motivates a procedure that gives more weight to the first words within a company name, since they are more likely to be part of the distinctive part.

Legal-Parts Dictionay

In order to uniform abbreviation in the legal part to a unique term, we constructed a many-to-one dictionary using a subsample of our training data. In particular, we manually selected the abbreviations and typing errors in the legal part looking at the disagreements between our sample and tax names. For example, "Incorporated" appears as "Inc." "INC", "Incorp" and so on in our data. Therefore all this abbreviations, when found, are mapped into "Incorporated" as described below. The dictionary is available online with the code.

Matching Procedure

For each company name in the BidX data, the algorithm searches the best match in the tax database. Although the algorithm is meant for the comparison of corporate names, it can be augmented with additional information if available. For example, in our main application the BidX data contains information about the name and the state of origin of the bidding firms. The latter can be used to improve the quality of the matching introducing a blocking for firms that do not match the origin state, as explained below. Let A be the firm we want to match, S^a be the string name and $State^a$ its state of origin. The state of origin is only used if the *state* option is enabled in the code provided. The algorithm proceeds as follow.

1. Name Normalization
   All non-alphanumeric characters with the exception of spaces are removed from S^a and all letter characters are capitalized. Consecutive white spaces are replaced with one white space. Any substring separated by one space is defined as a "word". Every word so identified that belongs to the legal-part dictionary is removed. For example, "Amnio Brothers Inc." is composed by the three words "Amnio" "Brothers" and "Inc.". After the first step, it would be normalized to "AMNIO BROS", since the word "INC" is recognized in our dictionary as a legal part and therefore removed. We refer to the normalized string as S^a_{norm}. The same normalization is applied to every company name in the tax data. If the normalized name is not unique in the tax data, we restrict to the ones that ever filed at least once one of the three firm tax returns (1120,1120-S or 1065). If more than one firm did, we select the one with highest value added.

2. Shortlisting
   Let S^a_{norm} be composed by $n \geq 2$ words. Starting from the first word, we search in the list of normalized tax data company names the subset of names that contains that word. If the subset is empty, no matching occurs and the matching for A ends. If the subset is a singleton, A is paired with the unique element of the set and the shortlisting step ends for A. If the subset has more than one element, we proceed with the second word in S^a_{norm} and consider only the candidate matches that also contains the second word. If the set still contains more than one element, we proceed with the third word and so on, until all the n words are used or we obtain either a singleton or an empty set. If this iteration leads to a singleton, A is paired with the unique element of the set. If it leads to an empty set, then A is paired with the smallest non-empty subset from the previous iterations. In short, this

Table A.1: Example Search

Steps	Output
String Normalization	Normalized Name: HANNAFORD BROS DISTRIBUTION
Shortlisting	The names(in bracket) and normalized names in the shortlist are shown below. The shared word is in bold. KELLY **HANNAFORD BROS DISTRIBUTION** (Kelly Hannaford Brothers Distribution Company) **HANNAFORD BROS DISTRIBUTION**(Hannaford Brothers Distri. C.) HASTING **HANNAFORD BROS DISTRIBUTION** (Hasting Hannaford Bros. Distribution Inc.)
Scoring	Normalized names in the shortlist are shown below. The scores are shown on the right of the names. KELLY HANNAFORD BROS DISTRIBUTION **(LR = 0.9)** HANNAFORD BROS DISTRIBUTION **(LR =1)** HASTING HANNAFORD BROS DISTRIBUTION **(LR =0.87)**
Unique match	HANNAFORD BROS DISTRIBUTION(Hannaford Brothers Distri. C.)

step selects a shortlist of candidate matches that share, after normalization, the highest number of initial words with A.

If the *state* option is enabled, only firms with that match exactly the $State^a$ are considered for shortlisting.

3. Scoring

This step employs the Levenshtein ratio (LR), a widespread measure of distance between strings, to select the best match from the shortlist. For each element of the set paired to A we compute its LR with respect to S^a. The company whose name has the highest score is selected as the match. If multiple companies tie for the top score, the one with the highest value added is selected.

If the option *strict* is enabled, all the company names that do not reach a minimum threshold $T \in (0, 1)$ in their LR are dropped. If all candidate matches are dropped, then A is considered unmatched. Hence the higher the T, the more stringent is the matching process. In our application, we considered $T = 0.6$.

A summary of the algorithm in pseudocode is presented below. Table A.1, instead, illustrates how the algorithm works with an example search. In our example, *strict* and *state* are disabled. Let "Hannaford Bros. Distribution Co." be our search query.

Matching Performance

We run the program on tax data to search bidders from our training data, in order to evaluate the accuracy of the procedure. Overall, the algorithm outperformed a simple string matching in both accuracy and number

Algorithm 1 Matching Algorithm Pseudocode

Input: S^a, $State^a$, $Dict$, $IRS\_Firmlist\_normalized$
Output: $Match^a$, $Score^a$
1. Name Normalization;
S^a_{norm} = Remove non-alphanumeric characters and double spaces from S^a and set uppercase letters;
$W^a \equiv \{w^a_1, w^a_2, ..., w^a_n\}$ = Split substrings at spaces (S^a_{norm});
for $i = 1, \dots, n$ **do**
    if $w_i \in Dict$ **then** $W = W -' \{w_i\}$;
S^a_{norm} = Merge words in W ;
2. Shortlisting; $Shortlist = IRS\_Firmlist\_normalized$
Candidate="Unmatched"
Out=0
i=0
repeat
    i=i+1
    C= {FirmName ∈ IRS_Firmlist_normalized | $w^a_i \in FirmName$}
    Shortlist= Shortlist ∩ C
    if *Shortlist is singleton* **then**
        Candidate= Shortlist
        Out=1
    if *Shortlist is empty* **then**
        $Out = 1$
    else
        Candidate= Shortlist
until *Out=1*;
3. Scoring
for $c \in Candidate$ **do**
    $Score^c$= Levenshtein distance(c,S^a_{norm})
Best= $argmax\{Score^c\}$
if *Levenshtein distance(Best,S^a_{norm})<0.6* **then**
    $Match^a$ ="Unmatched"
else
    $Match^a$=Best;
    $Score^a$= Levenshtein distance(Best,S^a_{norm})

Table A.2: Performance in BidX to tax data matching

	Simple Search	Fuzzy Match using Sieve			
	(1)	(2)	(3)	(4)	(5)
% Bidders Matched to Any Tax Record	80.2	99.9	97.6	99.9	95.8
% Bidders Matched to the True Tax Record	65.3	63.0	62.5	71.0	70.3
% Potential Matches Correctly Matched to Tax Records	78.6	75.8	75.1	85.4	84.5
Algorithm Parameters:					
Match must be perfect (string score = 1.0)	✓	✗	✗	✗	✗
Match must be high-quality (string score ≥ 0.6)	✗	✗	✓	✗	✓
Prefer matches in same state as auction	✓	✗	✗	✓	✓

of matches achieved. In our main application, the algorithm correctly matched 85.4 % of the bidders whose EIN are known and could be found in tax database. The use of the *State* option proved effective in increasing the number of true matches, while the *Strict* option with $T = 0.6$ improved accuracy by reducing the false matches. Table 2 provides further details on the quality of matching.

However, in order to assess the external validity of the algorithm outside our specific application, we constructed two test data sets using data from the Employee Benefits Security Administration (ESBA).

Using publicly available pension filings as a test data set

Our test data sets, *PensionData* and *PensionTest*, are constructed using Form 5500 Data sets that are published by Employee Benefits Security Administration (ESBA)[1]. Form 5500 data sets contain information, including company names and EINs, about the operations, funding and investments of approximately 800,000 business entities. We consider both retirement and Health and Welfare data sets, drop every variable except the Company Name and the Employer and remove every duplicate observation. For every unique EIN in this subset, we find all names that are associated with it, then we discard any name duplicate. Most of EIN are associated with multiple company names, reproducing a condition similar to the one in our main application. For each EIN, if multiple names are associated with it, we select the first name and put in into the *PensionData* data set and all the others into the *PensionTest* data set. If there is only one name associated with the EIN, we still add that name into *PensionData*. This gives us 709,850 companies in *PensionTest* and 1,270,079 companies in *PensionData*. We then proceeded to test our program using *PensionData* as a main data set and *PensionTest* as a query set.

Performance Tests

We tested the program by searching in *PensionData* all the $709,850$ *PensionTest* firms. Since we have the EIN for all the names in the two data set, we can evaluate the matching performance. The program achieved an average speed of 152 queries per second and an average accuracy of 73.39% among matched queries for a $T = 0.6$ using the *strict* option. Table 3 presents the percentage of correctly matched firms and false matches for different values of T. Should be noted that the percentage of correct matches is not monotone in T when T is close to 1. In fact, requiring extreme level of string similarity lead to a loss of correct matches

[1] https://www.dol.gov/agencies/ebsa/researchers/data

that outweighs the gains in precision. Therefore we do not recommend setting T above 0.9. In Table 4, instead, we provide a closer look to the effectiveness of the shortlisting step. Looking at the distribution of the shortlists' length, we see that over 50% of the sample is matched in the shortlisting step and 70% of the candidate matches requires the scoring of at most 2 candidates. Furthermore, the top 99% percentile of the longest shortlists amount to 2733 candidates, that is just the 0.2% of the potential matches that a standard matching algorithm would have to consider for each query.

Table A.3: Matching Performance With Different T

T	0	0.2	0.4	0.6	0.8	0.9	1
Matches	99.05%	99.04%	98.46%	91.68 %	74.52%	64.44%	49.01%
Correct Matches	70.36%	70.37%	70.57%	73.39 %	80.69%	84.12%	82.58%

Table A.4: The Quantiles of Shortlist Lengths

Quantile	1%	10%	30%	50 %	70%	90%	99%
Length	1	1	1	1	2	37	2733

D.2 Building the Auction data

This appendix describes our data source for auction bids and how we build the data set for our main application.

BidX Website

The BidX website collects information on bids and bidders for procurement auctions held by the transportation departments of over 43 US states. It can be freely accessed at `www.bidx.com`, although the access to information on the bidders requires a payed account registration.

Scraping procedure

We performed the scraping using the Python library *Selenium* to automate browser actions. We registered a Bidx.com account, which is required to access bidder information.

We collect the auction information for a given state using the following procedure:

1. We go to the web page of that state on Bidx.com and select the latest letting.

   Browser actions: visit `www.bidx.com`, select the desired agency from "Select a U.S. Agency" drop down menu and click the button "go". Then click the "Letting" tab on the top left corner of the new refreshed web page and click the first letting date hyperlink in "List of Letting" table.

2. There are two different sources of information - "Apparent Bids" and "Bid Summary" - on a letting page. More specifically, "Apparent Bids" and "Bid Summary" contains auction information but in different formats, and both of them have links to additional bidder information, which requires a paid account to access (see Figure 3 and 4). Starting from the latest letting page, our scraper clicks the

Figure A.1: Bidx.com

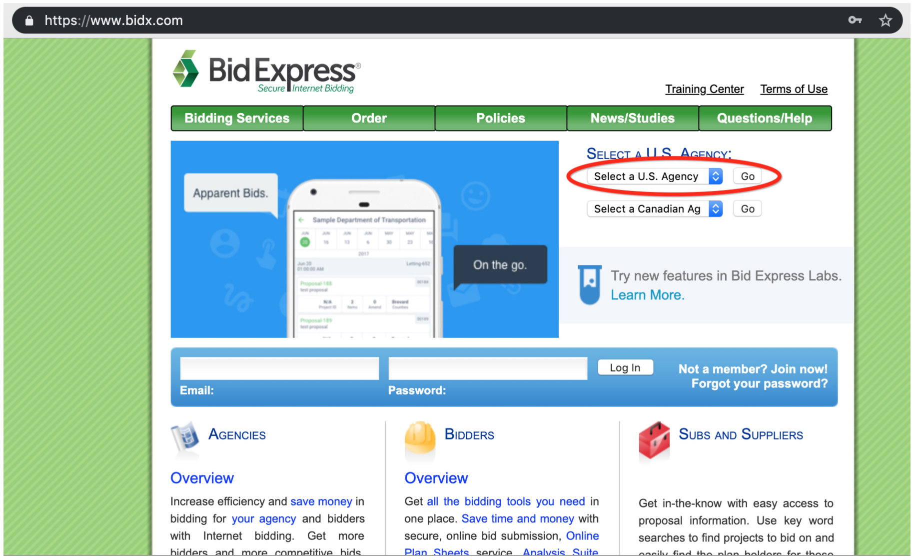

hyperlink "Apparent Bids" (Fig. 2) then downloads a csv file for every bidder by clicking on the bidder hyperlink (Fig. 3) and "Export(csv)" on the refreshed page.

Figure A.2: A Letting Page

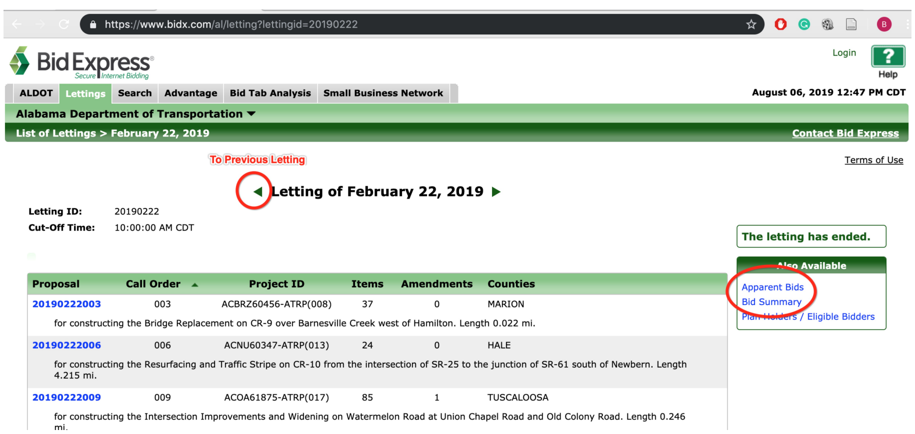

If there is no information on the refreshed page, it moves to a new letting by clicking the arrow with html class "prev_arrow". The procedure is iterated until the arrow is not clickable. We repeat the same procedure for the "Bid Summary" hyperlink.

Through this procedure, we obtain three tables for each letting:

a. auction information from "Apparent Bids", which contains: bidder names, bidder ID, bidder ranks, bid amounts, bidder call orders, project description, counties[2], letting ID and letting date. We do note that a few states record two extra variables: DBE Percentage and DBE Manual.

b. auction information from "Bid Summary", which contains: bidder names, bidder ID, bidder ranks, bid amounts, bidder call orders, counties, proposal ID and letting date.

[2]the location of the project

Figure A.3: An "Apparent Bids" Page

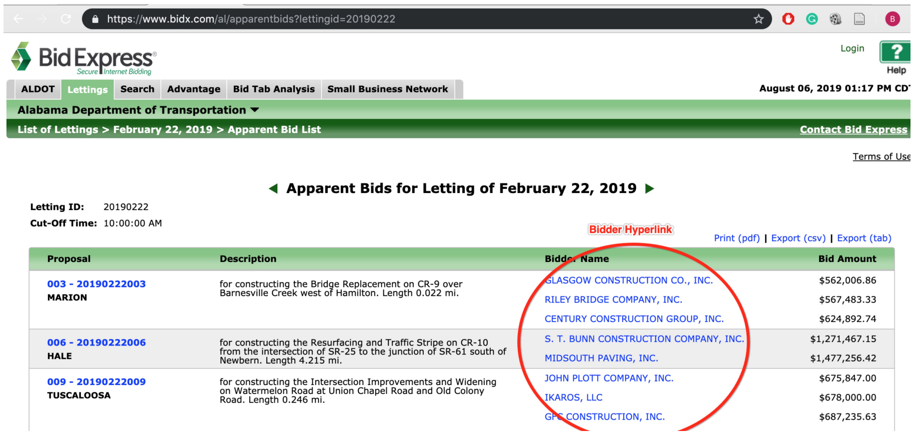

Figure A.4: A "Bid Summary" Page

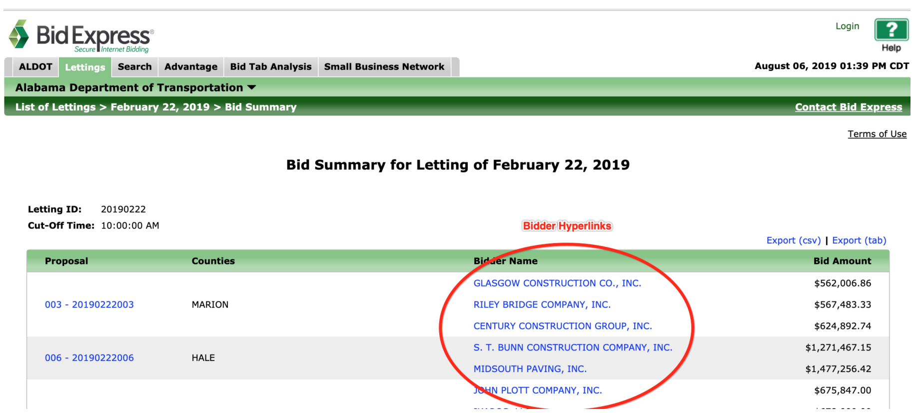

c. additional bidder information from bidder links, which contains: company name, company address, company phone Number, company fax number.

We then merge the table c into a and b. Therefore, 2 files are created for every letting, one for "Apparent Bids" and one for "Bid Summary" with both auction and firm level information.

The information at the letting level are then further aggregated for each state as follow:

1. For a state X, we merge its "Apparent Bids" files into one single file *X_ apparentbid* and "Bid Summary" files into one single file *X_ bidsummary*. Then we add a new variable *State*, which is two-letter abbreviation of states, in *X_ apparentbid* and *X_ bidsummary*.

2. Then we find lettings that are in *X_ bidsummary* but not in *X_ apparentbid*, and augment them so that they have the same variables as lettings in *X_ apparentbid*[3]. The variables added are filled with "N/A". Then we merged these lettings with *X_ apparentbid* into one file *X_ all*

3. We merge all **_ all* files into one final file.

As a result, we obtain a comprehensive file that has the following variables: Bidder Rank, Bidder Call Order, Project Description, Counties, Letting ID, Bidder ID, Bidder Name, Bid Amount, DBE Percentage,

[3]Proposal in *X_ bidsummary* is treated as Letting ID

DBE-Manual, Company Name, Company Address, Company Phone Number, Company Fax Number, State, Letting Date.

State	Bidding Firms	Share of Value Added	Share of FTE Workers
AL	196	15.7%	17.4%
AR	149	7.9%	12.8%
AZ	*	*	*
CA	1,041	8.3%	11.2%
CO	241	12.6%	14.7%
CT	126	9.4%	15.5%
FL	344	30.7%	10.6%
GA	137	4.3%	7.0%
IA	256	15.4%	20.7%
ID	112	17.2%	13.6%
IL	*	*	*
IN	213	10.6%	16.6%
KS	130	13.7%	21.6%
KY	*	*	*
LA	167	11.5%	10.8%
MA	*	*	*
MD	*	*	*
ME	141	13.7%	16.9%
MI	391	9.5%	16.3%
MN	262	13.5%	19.8%
MO	179	14.9%	13.3%
MS	*	*	*
MT	122	15.0%	23.6%
NC	135	5.2%	9.8%
ND	*	*	*
NE	*	*	*
NH	*	*	*
NJ	*	*	*
NM	*	*	*
NV	*	*	*
NY	*	*	*
OH	320	43.7%	17.5%
OK	*	*	*
OR	*	*	*
PA	*	*	*
SC	*	*	*
SD	*	*	*
TN	140	5.3%	11.5%
TX	551	4.9%	9.6%
UT	*	*	*
VA	241	14.2%	12.0%
VT	*	*	*
WA	200	7.5%	14.0%
WI	194	12.1%	14.6%
WV	103	13.7%	19.0%
National	6,792	10.7%	9.9%

Table A.5: 2010 Shares of Construction Sector due to Bidding Firms

Notes: Among firms in the construction sector in 2010, this table considers the share of value added and FTE workers due to the firms that participated in auctions in our sample. We drop from this table firms that have missing values on the variables displayed, so the total sample size must be smaller than in Table 1 in the main text.

E Online Appendix: Accounting for Capital

E.1 Set up

We consider the problem of including capital in the production function of firms. We build on the approach of Lamadon et al. (2019) [hereafter, LMS] by allowing for firms to rent capital in a competitive market with rental rate p_K. This approach in this paper [hereafter, KLMS] differs in an important way:

- LMS consider profit maximization that accounts for revenue and cost jointly, but do it sequentially: first, they find the optimal K for each L to maximize conditional profit; second, they find the optimal L.
- KLMS first writes revenue and costs as functions of Q so as to separate the joint maximization in two steps: first find optimal combination (K,L) for each Q; second find the optimal Q. Since there is a one to one mapping between Q and L, we can rewrite everything in terms of L.

KLMS needs to consider the optimal inputs for each Q because Q may be set externally by the auction, whereas LMS assume firms choose their own Q.[4]

Formally, denote the underlying production function in capital and labor as $\Omega K^{\beta_K} L^{\beta_L}$ and the inverse labor supply curve as $W = \theta_0 L^\theta$. Both papers start by rewriting the firm's problem into the following one with only labor:

$$\widetilde{\Omega} L^\rho - \widetilde{\theta}_0 L^{1+\theta}$$

where:

- LMS: $\widetilde{\Omega} = \frac{p_K}{\beta_K}(1-\beta_K)[\frac{p_0 \Omega \beta_K}{p_K}]^{\frac{1}{1-\beta_K}}$, $\rho = \beta_L/(1-\beta_K)$ and $\widetilde{\theta}_0 = \theta_0$ when it is perfect competition. The new $\widetilde{\Omega} L^\rho$ is revenue minus capital cost (value added) for any labor choice (at optimal capital choice in profit maximization conditioning on L). The new $\widetilde{\theta}_0 L^{1+\theta}$ is (only) the labor cost. Note: Under monopolistic competition parameterized by $\epsilon > 0$, we should replace Ω, β_K, β_L with $\Omega^{1-\epsilon}, \beta_K(1-\epsilon), \beta_L(1-\epsilon)$, respectively.
- KLMS: $\widetilde{\Omega} = \Omega[\frac{\beta_K}{\beta_L}\frac{(1+\theta)\theta_0}{p_K}]^{\beta_K}$, $\rho = (1+\theta)\beta_K + \beta_L$ and $\widetilde{\theta}_0 = [\frac{\beta_K}{\beta_L}(1+\theta)+1]\theta_0$. the new $\widetilde{\Omega} L^\rho$ is (gross) revenue for any labor choice. The new $\widetilde{\theta}_0 L^{1+\theta}$ is capital cost + labor cost (total cost) for any labor choice (at optimal capital choice in cost minimization conditioning on Q). Note: these values are the same regardless of monopolistic competition $\epsilon \geq 0$.

The difference is that LMS substracts the capital costs from revenues to form value added, while KLMS adds the capital costs to the labor costs to form the total costs. This difference results in very different ρ definitions: for the same underlying parameters $\beta_L = 1/3$ and $\beta_K = .58$, we have that $\rho = 0.79$ in LMS and $\rho = 1.08$ in KLMS (which recovers the actual estimates of ρ from both papers).

E.2 Details: KLMS

Now, we think about the problem when there are both capital and labor in the production function. Contractors can always rent capital at price p_K and hire labor at price $p_L = \theta_0 L^\theta$. Consider a Cobb–Douglas

[4] Intuition for the optimization problems: Think about climbing a mountain where x is L, y is K and profit is height. On a contour map, LMS do a grid search of the best y for each x; and then find the best y to maximize profit. KLMS finds all the cost minimizing combinations (K,L) and then follow this different route but reach the same peak point as in LMS. Both approaches reach the same peak point simply because they are dual problems, but get there through different strategies.

production function (in physical units)

$$Q = \Omega K^{\beta_K} L^{\beta_L}.$$

Given any production level Q, the firm can find the most cost efficient combination (K, L) by solving the following problem

$$\min_{(K,L):Q=\Omega K^{\beta_K} L^{\beta_L}} p_K K + \theta_0 L^{1+\theta},$$

where $p_K K + \theta_0 L^{1+\theta}$ represents the total cost. This leads to the Lagrange function $p_K K + \theta_0 L^{1+\theta} + \lambda(Q - \Omega K^{\beta_K} L^{\beta_L})$ and FOCs

$$p_K = \lambda\Omega\beta_K K^{\beta_K - 1} L^{\beta_L}$$
$$(1+\theta)\theta_0 L^{\theta} = \lambda\Omega\beta_L K^{\beta_K} L^{\beta_L - 1},$$

which leads to the optimal choice of capital as a function of labor $K^* = \frac{\beta_K}{\beta_L}\frac{(1+\theta)\theta_0}{p_K} L^{1+\theta}$ or simply $K^* = x L^{1+\theta}$, where $x = \frac{\beta_K}{\beta_L}\frac{(1+\theta)\theta_0}{p_K}$. .

Recall that $Q = \Omega[xL^{1+\theta}]^{\beta_K} L^{\beta_L}$. We have the optimal labor choice $L^* = [\frac{Q}{\Omega x_K^{\beta}}]^{\frac{1}{(1+\theta)\beta_K+\beta_L}}$ and cost becomes $c^*(Q) = p_K x L^{1+\theta} + \theta_0 L^{1+\theta} = (p_K x + \theta_0)[\frac{Q}{\Omega x^{\beta_K}}]^{\frac{1+\theta}{(1+\theta)\beta_K+\beta_L}}$. Or $Q = L^{*\rho}\Omega x^{\beta_K}$ and $c^* = (p_K x + \theta_0)L^{*\rho}$. In other words, we can write everything in terms of labor again. Hereafter, we omit the star.

Using this relationship in the production function gives

$$Q = \Omega K^{\beta_K} L^{\beta_L} = \Omega[\frac{\beta_K}{\beta_L}\frac{(1+\theta)\theta_0}{p_K}]^{\beta_K} L^{(1+\theta)\beta_K+\beta_L} = \widetilde{\Omega} L^{\rho},$$

where $\widetilde{\Omega} = \Omega[\frac{\beta_K}{\beta_L}\frac{(1+\theta)\theta_0}{p_K}]^{\beta_K}$ and $\rho = (1+\theta)\beta_K + \beta_L$.

Moreover, the total cost becomes

$$p_K K + \theta_0 L^{1+\theta} = [\frac{\beta_K}{\beta_L}(1+\theta) + 1]\theta_0 L^{1+\theta} = \widetilde{\theta}_0 L^{1+\theta},$$

where $\widetilde{\theta}_0 = [\frac{\beta_K}{\beta_L}(1+\theta) + 1]\theta_0$.

E.3 Details: LMS

Consider perfect competition $\epsilon = 0$. LMS solves the profit maximizatoin problem in the direction of K first

$$\max_K \quad p_0\Omega K^{\beta_K} L^{\beta_K} - p_K K - \theta_0 L^{1+\theta}$$

which gives $p_0\Omega\beta_K K^{\beta_K - 1} L^{\beta_L} = p_K$ i.e. the optimal capital choice as a function of L

$$K = [\frac{p_0\Omega\beta_K L^{\beta_L}}{p_K}]^{\frac{1}{1-\beta_K}}$$

Their problem becomes

$$\max_L \quad p_0\Omega[\frac{p_0\Omega\beta_K L^{\beta_L}}{p_K}]^{\frac{\beta_K}{1-\beta_K}} L^{\beta_L} - p_K[\frac{p_0\Omega\beta_K L^{\beta_L}}{p_K}]^{\frac{1}{1-\beta_K}} - \theta_0 L^{1+\theta}$$

where L is the only choice variable.

Collecting terms gives our new representation

$$\max_{L} \quad \frac{p_K}{\beta_K}(1-\beta_K)[\frac{p_0\Omega\beta_K}{p_K}]^{\frac{1}{1-\beta_K}} \cdot L^{\beta_L/(1-\beta_K)} - \theta_0 L^{1+\theta}$$

which defines the tilde terms. If it is monopolistic competition parameterized by $\epsilon > 0$ in KLMS, we should replace Ω, β_K, β_L with $\Omega^{1-\epsilon}, \beta_K(1-\epsilon), \beta_L(1-\epsilon)$, respectively.

E.4 Optimal labor choice

This subsection will show that further optimize in labor lead to the same results for LMS and KLMS. The optimal labor choice in KLMS is

$$p_0\widetilde{\Omega}\rho L^{-(1-\rho)} = \widetilde{\theta}_0(1+\theta)L^{\theta}$$

leading to

$$L^{\theta+(1-\rho)} = L^{[(1+\theta)(1-\beta_K)-\beta_L]} = \frac{p_0\widetilde{\Omega}\rho}{\widetilde{\theta}_0(1+\theta)} = p_0\widetilde{\Omega}\frac{(1+\theta)\beta_K+\beta_L}{[\frac{\beta_K}{\beta_L}(1+\theta)+1]\theta_0(1+\theta)} = \frac{\beta_K}{\theta_0(1+\theta)}\widetilde{\Omega}p_0$$

The optimal labor choice in LMS is determined by the FOC

$$(\frac{p_K}{\beta_K}-p_K)[\frac{p_0\Omega\beta_K}{p_K}]^{\frac{1}{1-\beta_K}}\frac{\beta_L}{1-\beta_K}L^{\frac{\beta_L}{1-\beta_K}-1} = \theta_0(1+\theta)L^{\theta},$$

which simplifies to

$$L^{(1+\theta)(1-\beta_K)-\beta_L} = [\frac{p_K\beta_L}{\beta_K\theta_0(1+\theta)}]^{1-\beta_K}\frac{p_0\Omega\beta_K}{p_K} = \frac{\beta_L}{\theta_0(1+\theta)}\widetilde{\Omega}p_0$$

For reference, recall that $\widetilde{\Omega} = \Omega[\frac{\beta_K}{\beta_L}\frac{(1+\theta)\theta_0}{p_K}]^{\beta_K}$ and $\rho = (1+\theta)\beta_K+\beta_L$. $\widetilde{\theta}_0 = [\frac{\beta_K}{\beta_L}(1+\theta)+1]\theta_0$. In summary, the optimal labor choices are identifical in LMS and KLMS.

F Online Appendix: Product Market with Perfect Competition

This section solves the second stage assuming a perfectly competitive product market ($\epsilon = 0$). Denote the competitive price as p_0. We need to distinguish two types of solutions.

In the case of an interior solution, the winner with TFP Ω maximizes the following problem

$$\max_{L} \quad p_0 \cdot [\Omega L^{\rho} - Q_1] - \theta_0 L^{1+\theta}$$

whose FOC gives the optimal labor choice $L^*_{interior}(\Omega) = [\frac{p_0\Omega\rho}{\theta_0(1+\theta)}]^{\frac{1}{\theta+1-\rho}}$ and the optimal output level $Q^*_{interior}(\Omega) = \Omega[\frac{p_0\Omega\rho}{\theta_0(1+\theta)}]^{\frac{\rho}{\theta+1-\rho}}$. Note that both optimal choices are independent of Q_1. Moreover, the optimal profit is

$$\Pi^*_{interior}(\Omega|Q_1) = (p_0\Omega)^{\frac{1+\theta}{\theta+1-\rho}}\theta_0^{-\frac{\rho}{\theta+1-\rho}}\left[(\frac{\rho}{1+\theta})^{\frac{\rho}{\theta+1-\rho}} - (\frac{\rho}{1+\theta})^{\frac{1+\theta}{\theta+1-\rho}}\right] - p_0Q_1$$

In the case of a corner solution, i.e. $Q_1 > Q^*_{interior}(\Omega)$, the firm produce only Q_1 so that $\Omega L^{\rho} = Q_1$, which gives $L^*_{corner}(\Omega) = (Q_1/\Omega)^{\frac{1}{\rho}}$. The profit is

$$\Pi^*_{corner}(\Omega|Q_1) = -\theta_0(\frac{Q_1}{\Omega})^{\frac{1+\theta}{\rho}}$$

The optimal profit achievable is defined by $\Pi^*(\Omega|Q_1) = \max\{\Pi^*_{corner}(\Omega|Q_1), \Pi^*_{interior}(\Omega|Q_1)\}$. Note that for the loser, it is always optimal to adopt the interior solution. On the other hand, the winner adopt the interior solution when Q_1 is small and the corner solution otherwise. In summary,

1. When Q_1 is small, i.e., $Q_1 < Q^*_{interior}(\Omega)$, regardless of its winning status, this firm hires the same amount of labor to equalize the marginal output and marginal labor cost. It produces the same amount of output in physical units in both cases. It makes a profit $\Pi^*_{interior}(\Omega|Q_1)$ upon winning and $\Pi^*_{interior}(\Omega|Q_1)+p_0Q_1$ upon losing. Therefore, in the bidding stage, the firm knows the opportunity cost of winning the auction is $\Delta(\Omega|Q_1 < Q^*_{interior}(\Omega)) = p_0Q_1$, where p_0Q_1 is essentially the opportunity cost of selling the same amount Q_1 to the private market.

2. When Q_1 is large, i.e., $Q_1 > Q^*_{interior}(\Omega)$, the DOT project pushes the firm to a high level of labor and a high marginal cost upon winning. In this case, the firm focuses on the DOT project and produce Q_1 upon winning the auction, which requires a labor amount s.t. $\Omega L^\rho = Q_1$. That is, $L^*_1 = (\frac{Q_1}{\Omega})^{\frac{1}{\rho}}$. On the other hand, the firm still produces at the interior solution $Q^*_{interior} = \Omega[\frac{p_0\Omega\rho}{\theta_0(1+\theta)}]^{\frac{\rho}{\theta+1-\rho}}$ using $L^*_{interior} = [\frac{p_0\Omega\rho}{\theta_0(1+\theta)}]^{\frac{1}{\theta+1-\rho}}$ upon losing the auction. Note that winning leads to more production than losing and the difference is $Q_1 - Q^*_{interior}$. Moreover, the firm hires more labor upon winning than losing. Overall, the firm knows the difference between winning and losing (opportunity cost) is $\Delta(\Omega|Q_1 > Q^*_{interior}(\Omega)) = \theta_0(\frac{Q_1}{\Omega})^{\frac{1+\theta}{\rho}} + (p_0\Omega)^{\frac{1+\theta}{\theta+1-\rho}}\theta_0^{-\frac{\rho}{\theta+1-\rho}}\left[(\frac{\rho}{1+\theta})^{\frac{\rho}{\theta+1-\rho}} - (\frac{\rho}{1+\theta})^{\frac{1+\theta}{\theta+1-\rho}}\right]$.

   The above two representations of the opportunity cost in different range of TFP gives an overall distribution of opportunity cost, i.e., $\Delta(\Omega) \sim F(\cdot)$. In the first stage, the firm chooses the optimal bid that solves the same problem as in the main text.

G Online Appendix: Solving the Model Using the Quantile Representation

Overview

For any given TFP and the outcomes from the first-stage (winner status and size of procurement projects won), the second-stage of the model can be solved to give numerical values of the firm outcomes such as profits, wages and employment, which is the information needed to perform counterfactual predictions. To solve the first-stage, we must account for equilibrium bidding behavior, which depends on the size of the procurement project, the number of bidders, and the TFP distribution. The symmetric equilibrium described in the main text involves numerical integration, which is costly since we need to solve this model many times in our counterfactual experiments. To speed up this calculation, we implement the quantile representation method of Luo (2019). Here, we provide an overview of the steps taken to solve the first-stage and second-stage problems.

Second stage: Denote the TFP quantile function as $\Omega(\alpha)$ where, for example, $\alpha = 0.10$ indicates the 10th quantile of the TFP distribution. We use a log Normal distribution to approximate the distribution of TFP, which allows for a simple mapping between Ω and α, choosing the standard deviation that matches the interquartile range of TFP (reported in Table 2). For each combination of winner status, TFP quantile, and auction size $\left(d, \alpha, \bar{Q}^G\right)$, we solve the second-stage problem for firm and worker outcomes. This is done by numerical optimization of the profit function (equation 8) subject to the labor supply curve (equation 4), the production function (equation 6), and the optimal materials condition (equation 7).

First stage: The challenge is to compute expectations of the second-stage across the distribution of outcomes from the first-stage. To solve the first-stage, note that the opportunity cost of winning an auction of size $\bar{Q}^G$ is $\Delta\left(\alpha|\bar{Q}^G\right) = \pi_0^H(\alpha) - \pi_1^H\left(\alpha|\bar{Q}^G\right)$. Since π_1^H is the winning firm's revenue in the private market net of the total cost, it follows that $\pi_0^H > \pi_1^H$ and thus $\Delta > 0$. π_1^H is decreasing in $\bar{Q}^G$, and π_0^H does not depend on $\bar{Q}^G$. Moreover, Δ is decreasing in α. In other words, higher TFP firms have lower opportunity cost of producing in the government procurement market. Since α represents quantiles of TFP, it has the standard uniform distribution. The probability that the winning quantile is less than α is the probability that it is the lowest among all I bidders' draws from the standard uniform distribution, yielding the probability α^I and associated density function $f_1(\alpha, I) = I\alpha^{I-1}$. By similar reasoning, the density function of a losing firm's TFP quantile is $f_0(\alpha, I) = \frac{I}{I-1}(1-\alpha^{I-1})$.

What do we mean by "solution"? Let $Y_d\left(\alpha|\bar{Q}^G\right)$ denote a second-stage outcome for a firm characterized by TFP quantile α bidding in an auction of size $\bar{Q}^G$. Using the distribution functions from the first stage, we compute the expected outcome as $\mathbb{E}\left[Y_d|\bar{Q}^G, I\right] = \int_0^1 Y_d\left(\alpha|\bar{Q}^G\right) f_d(\alpha, I)\, d\alpha$. For example, the probability that a bidder with TFP Ω_j wins the project is the probability that its TFP is the highest among all participating bidders, i.e, $H(\Omega_j)^I$,where H denotes the distribution of TFP. This implies that the density function of the winner's TFP is $IH(\Omega_j)^{I-1}h(\Omega_j)$. The profit function depends on who wins the auction, in particular, the TFP of the winner. The expected profit of the winner is then,

$$\bar{\pi}_{1,j} = \int \pi_{1,j}(\Omega_j|\bar{Q}^G) \times [IH(\Omega_j)^{I-1}h(\Omega_j)]d\Omega_j = \int \pi_{1,j}(\Omega_j(\alpha)|\bar{Q}^G) \times I\alpha^{I-1}d\alpha.$$

Note that this expectation depends on the combinations $\left(\bar{Q}^G, I\right)$. One possibility is to solve the model for each possible combination of $\left(\bar{Q}^G, I\right)$, and then average across them. In our setting, this is computationally infeasible. An alternative is to evaluate $\left(\bar{Q}^G, I\right)$ at representative values. In practice, we choose the values of $\left(\bar{Q}^G, I\right)$ that provide the best fit to $\left(V^W, V^F\right)$. The best fit yields a model-simulated incidence on workers of \$981, which is very close to the "sufficient statistics" estimate of \$959 in Table 2, and incidence on firms of \$3,098, which is very close to the "sufficient statistics" estimate of \$2,873 in Table 2. The implied incidence share on workers is about 24%, which is close to our main estimate of 25%. It is reassuring to find that the best fit is achieved at $I = 10$ bidders per auction, which is close to the mean observed value in the data of 8.7 bidders per auction.

Additional details

We now provide the derivation of the quantile representation of the optimal bidding strategy. Consider a procurement auction model. Following (Guerre et al., 2000), we can rewrite the first-order condition and obtain a representation of the cost as a function of observables:

$$c = b - \frac{1}{I-1}\frac{1-G(b)}{g(b)},$$

where $G(\cdot)$ and $g(\cdot)$ are the bid distribution and density, respectively. Since the bidding strategy is strictly increasing, we can further rewrite it in terms of quantiles:

$$c(\alpha) = b(\alpha) - \frac{1}{I-1}[1-\alpha]b'(\alpha),$$

where $c(\cdot)$ and $b(\cdot)$ are the cost quantile function and the bid quantile function, respectively. The boundary condition is that the least efficient firm bids the highest, i.e., $c(1) = b(1)$.

Following (Luo, 2019), we can solve this ODE and obtain the mapping from cost quantile function to bid quantile function

$$b(\alpha) = (I-1)(1-\alpha)^{1-I} \int_{\alpha}^{1} c(x)(1-x)^{I-2} dx.$$

This representation is convenient for numerically solving the first-price procurement auction model. In fact, we have

$$\begin{aligned} b'(\alpha) =& -(I-1)(1-I)(1-\alpha)^{-I} \int_{\alpha}^{1} c(x)(1-x)^{I-2} dx \\ & -(I-1)(1-\alpha)^{1-I} c(\alpha)(1-\alpha)^{I-2} \end{aligned}$$

So that

$$\begin{aligned} -\frac{1}{I-1}[1-\alpha]b'(\alpha) =& (1-I)(1-\alpha)^{1-I} \int_{\alpha}^{1} c(x)(1-x)^{I-2} dx + c(\alpha) \\ =& c(\alpha) - b(\alpha). \end{aligned}$$

Moreover, when $\alpha \to 1$, using L'Hospital's rule gives

$$b(1) = \lim_{\alpha \uparrow 1} (I-1) \frac{\int_{\alpha}^{1} c(x)(1-x)^{I-2} dx}{(1-\alpha)^{I-1}} = c(1).$$