

NBER WORKING PAPER SERIES

IMPERFECT COMPETITION AND RENTS IN LABOR AND PRODUCT MARKETS:
THE CASE OF THE CONSTRUCTION INDUSTRY

Kory Kroft
Yao Luo
Magne Mogstad
Bradley Setzler

Working Paper 27325
<http://www.nber.org/papers/w27325>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2020, Revised July 2023

The opinions expressed in this paper are those of the authors alone and do not reflect the views of the Internal Revenue Service, the US Treasury Department, or the National Bureau of Economic Research. This work is a component of a larger project on income risk in the United States, conducted through the SOI Joint Statistical Research Program. We thank Victor Aguirregabiria, David Autor, El Hadi Caoui, Lancelot Henry de Frahan, Sabrina Howell, Larry Katz, Peter Kuhn, Tom Meling, Aviv Nevo, Matt Notowidigdo, Elena Prager, Juan Carlos Suarez Serrato, John van Reenen, Owen Zidar, and participants at ASSA, Booth, CESifo, Cornell, Cowles Foundation, Hebrew University, Houston, IZA, LSE, Maryland, Microsoft Research, MIT, NBER IO, NBER Summer Institute, Northwestern, Penn State, Rice, Simon Frasier, SITE Labor, SITE IO, SUFE IO, Toronto, UBC, UCL, and Zurich for valuable comments. We thank Abdelrahman Amer, Tommaso Alba, Huishi Chen, JoonHwan Cho, Stephen Claassen, Cheok In Fok, Yingcheng Luo, Eric Mackay, Chester Madrazo, Marc-Antoine Schmidt, and Stephen Tino for valuable research assistance. Kroft and Luo acknowledge funding from the SSHRC Insight Grant. Mogstad and Setzler acknowledge funding from NSF Grant SES-1851808 and the Norwegian Research Council.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Kory Kroft, Yao Luo, Magne Mogstad, and Bradley Setzler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Kory Kroft, Yao Luo, Magne Mogstad, and Bradley Setzler

NBER Working Paper No. 27325

June 2020, Revised July 2023

JEL No. D44,J31,J42,L11

ABSTRACT

Existing work on imperfect competition typically focuses on either the labor market or the product market in isolation. In contrast, we analyze imperfect competition in both markets jointly, showing theoretically and empirically that focusing on one market in isolation may result in a limited or misleading picture of the degree and impacts of market power. Our empirical setting is the US construction industry. We develop, identify and estimate a model where construction firms imperfectly compete with one another for workers in the labor market and for projects in both the private market and the government market, where government projects are procured through auctions. Our analyses combine the universe of business and worker tax records with newly collected records from government procurement auctions. We use the estimated model to quantify the markdown of wages and the markup of prices, to show that the impacts of an increase in market power in one market are attenuated by the existence of market power in the other market, and to quantify the rents, rent-sharing, and incidence of procurements in the US construction industry.

Kory Kroft
Department of Economics
University of Toronto
150 St. George Street
Toronto, ON M5S 3G7
and NBER
kory.kroft@utoronto.ca

Magne Mogstad
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
magne.mogstad@gmail.com

Yao Luo
University of Toronto
Max Gluskin House 322
150 St. George St.
Toronto
Canada
yao.luo@utoronto.ca

Bradley Setzler
Department of Economics
Pennsylvania State University
University Park, PA 16802
and NBER
bradley.setzler@gmail.com

1 Introduction

Researchers and policymakers are keenly interested in measuring the degree of imperfect competition in the US economy and in understanding how it affects the outcomes of workers and firms. Existing work on imperfect competition typically focuses on either the labor market or the product market in isolation. In contrast, we analyze imperfect competition in both markets jointly, showing theoretically and empirically that focusing on one market in isolation may result in a limited or misleading picture of the degree and impacts of market power.

Our empirical setting is the US construction industry, where we seek to quantify the markdown of wages and the markup of prices, show that the impacts of an increase in market power in one market are attenuated by the existence of market power in the other market, and quantify the rents, rent-sharing, and incidence of government procurements for construction projects. To do so, we develop, identify and estimate a model where construction firms imperfectly compete with one another for workers in the labor market and for projects in both the private market and the government market, where government projects are procured through auctions. Our analyses combine the universe of business and worker tax records with newly collected records from government procurement auctions.

In Section 2, we present the model. 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. \(2022\)](#). Competitive labor market theory typically assumes that firms are wage-takers so that the labor supply facing a given firm is perfectly elastic. To allow the firm-specific labor supply curve to be imperfectly elastic so that the firm may have wage-setting power, we let workers have heterogeneous preferences over the non-wage job characteristics or amenities that firms offer. 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 rents to employers and workers, where rents refer to the excess return over that required to change a decision, as in [Robinson \(1933\)](#) and [Rosen \(1986\)](#).

The firm side of the model consists of two types of product markets in which the construction firms may participate: the private market and the government market, the latter of which procures projects through auctions. Incorporating both types of product markets not only gives a more accurate representation of firms’ production choices in the construction industry, but also facilitates identification of key model parameters through the use of data on the bids and outcomes of procurement auctions. The firm’s behavior is specified as a two-stage problem. In the first stage, firms bid for a government project procured through a first-price sealed-bid auction. The project specifies the amount of output that must be produced. At the end of the first stage, firms learn the outcome of the auction. If a firm wins the auction, it receives the winning bid amount as revenue and commences production. In the second stage, the firm chooses inputs to maximize profit from total production, including projects in both the private market and the government market. We allow the firm-specific demand curve in the private product market to be imperfectly elastic. As a result, the firm may earn rents due to price-setting power in the private product market and because there may be a limited number of bidders in the procurement auction.

In our model, a firm is defined to have “double market power” if it can profitably set both the price above the marginal cost and the wage below the marginal revenue product of labor.¹ In Section 3, we theoretically examine the implications of this double market power for the outcomes and behavior of workers and firms. We first characterize the determinants of the markdown of wages and markup of prices in the presence of double market power. We show that not only upward-sloping labor supply but also downward-sloping product demand are relevant and distinct sources of labor market power, as measured by the markdown of wages relative to the value of the marginal product of labor. Similarly, product market power, as measured by the markup of prices relative to the cost of production, is determined not only by the elasticity of product demand but also the elasticity of labor supply.

Next, we show that the impacts of an increase in market power in one market will be *attenuated* by the existence of market power in the other market. The intuition is

¹Our “double market power” result resembles classical double marginalization, in which consumer prices reflect the interaction of the markups that arise at each stage of production along a vertical supply chain. Our result, however, applies to a single firm that has market power both in the product market and the labor market.

straightforward: In response to an increase in labor market power, the firm wants to lower the wage it pays by reducing employment and, thereby, output. However, the firm will choose to reduce employment and output less if it is facing downward-sloping demand, as lower output increases the price it can charge and, thus, the marginal revenue product of labor. Similarly, if the firm experiences an increase in product market power, it wants to increase the price it charges by reducing output and, thereby, employment. However, the firm will choose to reduce output and employment less if it is facing upward-sloping labor supply, as lower employment decreases the wage it has to pay and, thus, the marginal cost of labor.

Motivated and guided by the theoretical results in Section 3, the remainder of the paper seeks to identify and estimate the model in order to *quantify* the double market power and its implications for the outcomes and behavior of workers and firms in the US construction industry.

As explained in Section 4, our analyses are based on a matched employer-employee panel data set, which is formed by combining the universe of US business and worker tax records for the period 2001-2015. Firm data contain information on sales, profits, intermediate inputs, and industry. Worker data contain information on the number of workers and their earnings. We merge the employer-employee panel data set with a new data set that we assembled with information on US procurement auctions. The resulting data set covers billions of dollars in procurement contracts awarded to thousands of firms. Importantly for our identification strategy, we observe the identity and bid of each firm in an auction, not only that of the winner. Comparing the firms that win a procurement auction to those that lose (but bid similarly), we assess how employment and wages change if a firm wins a procurement contract. We find that, on average, winning a procurement auction leads to a 2% increase in earnings per worker in the post-auction time period. At the same time, the number of employees in the firm increases, on average, by about 8%. The evidence that winning 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 is consistent with the notion that firms face upward-sloping labor supply curves and, therefore, have wage-setting power in the labor market, as we allow in our model.

In Section 5, we demonstrate how the model parameters are identified from the data. The primary challenge to identify the firm-specific labor supply elasticity is unobserved shifts to labor supply that may affect both employment and wages. We employ and compare several alternative approaches to overcome this identification challenge, drawing especially on institutional features of procurement auctions. One approach compares winners and losers with close bids in the (price-only) auctions in order to isolate exogenous demand shocks that shift the winning firms along their labor supply curves.² In another approach, we leverage that there is a time delay between a firm placing a bid in the procurement auction (based on estimated costs) and commencing production on the procurement project, so that the auction bid may not depend on the firm-specific labor supply shock that is realized at the time of production.³

The primary challenge to identify the relevant technology parameters is firm-specific productivity shocks, which are unobserved correlates of both inputs and output. We overcome this identification challenge by taking advantage of the monotonicity of auction bids with respect to productivity. This monotonicity allows us to use data on bids to control for unobserved productivity differences across firms in the production function estimation. To recover the product demand elasticity, we extend the price markup estimator under Leontief production, recently used by [de Loecker et al. \(2020\)](#), to account for the wage markdown. We also show that the technology and product demand parameters are over-identified and use the additional moment condition to assess the model.

In Section 6, we present the estimates of the model parameters. We find that firms have significant wage-setting power with an estimated firm-specific labor supply elasticity of about 3.5 to 4.1. This estimate indicates that, if an US construction firm aims to increase the number of employees by 10%, it needs to increase wages by

²A small number of papers have used research designs that compare winners and losers of procurement auctions. They estimate how government purchases affect employment during an economic crisis ([Gugler et al., 2020](#)) and firm dynamics and growth ([Ferraz et al., 2015](#), [Hvide and Meling, 2020](#)). None of these studies use these comparisons to draw inference about imperfect competition or rents, nor do they use these comparisons to identify and estimate an economic model of firm and worker behavior.

³Similar timing assumptions are used in the large literature on production function estimation (see the discussions by [Akerberg et al. 2015](#) and [Gandhi et al. 2020](#)).

around 2.4-2.9%. Another finding is that firms have significant price-setting power in the private product market with an estimated product demand elasticity of 7.3. The estimate suggests that, in order for a firm to increase output by 10% in the private market, it must reduce the price of its product by about 1.4%. We also find that the production function exhibits (approximately) constant returns to scale over labor and capital.

In Section 7, we use the estimated model of the construction industry to quantify key implications of double market power. In Section 7.1, we show that double market power generates a total markdown of wages of more than 30% relative to the value of the marginal product of labor, and a total markup of prices of more than 40% relative to the cost of production. To gauge the importance of analyzing imperfect competition in both markets jointly, we compare these results to the markdown and markup estimates one would obtain if one only focused on labor market power or product market power in isolation. We find that, if one assumed a perfectly competitive product market, then our estimates of the labor supply elasticity would imply that the wage is only 20% below the value of the marginal product of labor. Conversely, if one assumed a perfectly competitive labor market, then our estimates of the demand elasticity would imply that the price is only 16% above the cost of production.

Next, in Section 7.2, we use our model to perform counterfactuals that quantify the extent to which impacts of an increase in market power in one market are attenuated by the existence of market power in the other market. If the labor supply elasticity of a given firm is reduced by half, our estimates suggest the firm employs 12% fewer workers and decreases wages by 6%. By comparison, if the firm did not have price-setting power in the product market, we find that it would employ 22% fewer workers and decrease wages by 11%.

In Section 8, we conclude our empirical analyses with a quantification of the rents, rent-sharing, and incidence of procurements in the US construction industry. We find that imperfect competition leads to worker rents per year of around \$11,600 per worker, while firm rents per year (as measured by profits) amount to about \$43,100 per worker. Comparing worker rents to firm rents, we see that more than three-fourths of total rents are captured by firms. Although winning a procurement contract crowds out some private market production, it increases total output, employment, and rents.

We find that 40% of these additional rents are captured by workers.

Our paper is primarily related to a large literature on imperfect competition, rents, and inequality in the labor market, reviewed by [Manning \(2011\)](#), [Card et al. \(2018\)](#), [Card \(2022\)](#), and [Lamadon et al. \(2022\)](#), but differs in several important ways. First, our paper differs from much of the existing literature in that we fully specify an equilibrium model and identify and estimate the model parameters. This allows us to not only measure the current size and share of the rents earned by firms and workers, but also to understand the underlying mechanisms and to quantify how the outcomes and behavior of firms and workers would change if market power increased or decreased. A second salient difference from the literature is that we analyze imperfect competition in both the labor and product market jointly.⁴ Our theoretical and empirical findings highlight how studies that focus on either the labor market or the product market in isolation may result in a limited or misleading picture of the degree and impacts of imperfect competition. Third, much of this existing work is trying to measure imperfect competition in the entire labor market, paying little attention to the large heterogeneity in technology and market structure across industries.⁵ In contrast, we focus on the construction industry, paying closer attention to the structure and the functioning of the relevant markets. We leverage institutional features of the construction industry for our identification arguments, for understanding and modeling the behavior of firms, and for estimating the incidence of government procurements.

Our paper also relates and contributes in several ways to the empirical literature on auctions, reviewed by [Athey and Haile \(2007\)](#). First, we present and solve an auction model with incomplete information about unobserved productivity (rather than costs). This allows for a flexible relationship between the probability of winning the auction and other firm outcomes that depend on productivity, such as employment and output. Second, our paper also contributes by quantifying how winning a procurement auction affects the firm’s total production and whether it crowds-in

⁴An exception is [MacKenzie \(2021\)](#), who considers the gains from international trade if there is labor and product market power.

⁵Notable exceptions include [Azar et al. \(2021\)](#) and [Lamadon et al. \(2022\)](#). They study imperfect competition in the entire US labor market, but account for imperfect substitutability across markets using a nested-logit structure on preferences. See the discussion by [Card \(2022\)](#).

or crowds-out activity in the private market. Third, our work complements existing papers on auctions by taking into account how bidding behavior depends on market power, both in the labor market and in the private product market. Fourth, our paper shows how bids in procurement auctions can be used to construct a control variable that we use to address the problem of unobserved productivity in the estimation of production functions.⁶

Lastly, our paper relates to a growing body of empirical work that estimates the pass-through and incidence of firm-specific shocks, reviewed by [Card et al. \(2018\)](#). An early example is [van Reenen \(1996\)](#), who studies how innovation affects firms' profits and workers' wages. He also investigates patents as a source of variation, but finds them to be weakly correlated with profits. Building on this insight, [Kline et al. \(2019\)](#) study the incidence of patents that are predicted to be valuable and [Howell and Brown \(2020\)](#) study the incidence of R&D grants. A related literature on skill-biased technical change has examined the wage and productivity effects of the adoption of new technology in firms (see [Akerman et al., 2015](#), and the references therein).

2 A Tractable Model of Imperfect Competition in Both Labor and Product Markets

In this section, we develop a model in which construction firms compete with one another for projects in the product market and for workers in the labor market. We allow for imperfect competition, in the form of monopolistic competition in the product market and monopsonistic competition in the labor market. The production side of the model incorporates two types of product markets in which construction firms may participate: the private market and the government market, where government projects are procured through auctions.

2.1 Worker Preferences and Labor Supply

Worker i in year t has the following preferences over being employed at a firm j :

$$\mathcal{U}_{it}(j, W_{jt}) = \log W_{jt} + \log G_{jt} + \eta_{ijt}, \quad (1)$$

⁶For reviews of the literature on production function identification, see [Akerberg et al. \(2015\)](#) and [Gandhi et al. \(2020\)](#).

where W_{jt} represents earnings, G_{jt} represents the common value of firm-specific amenities, and η_{ijt} captures worker i 's idiosyncratic tastes for the amenities of firm j .⁷ Since we allow amenities to be unobserved to the analyst, they can include a wide range of characteristics, such as distance to the firm from the worker's home, flexibility in the work schedules, effort required, workplace safety, and so on.

Our specification of preferences allows for the possibility that workers view firms as imperfect substitutes. The term G_{jt} gives rise to vertical employer differentiation: some employers offer good amenities while other employers offer bad amenities. The term η_{ijt} 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 their idiosyncratic taste for a given firm. We parameterize the distribution of η_{ijt} as i.i.d. Type-1 Extreme Value (T1EV) with dispersion $\theta \geq 0$.⁸ When θ is larger, horizontal employer differentiation becomes relatively more important.

We consider an environment with monopsonistic competition in a spot market for labor, making three key assumptions. First, firms do not observe the idiosyncratic taste for amenities of any given worker η_{ijt} . 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, we assume all workers are homogenous in skill. This assumption is motivated by the fact that we find no evidence of changes in worker quality in response to winning a procurement auction.⁹ Third, we assume firms are "strategically small" in the sense that each firm views itself as infinitesimal within the market.¹⁰

⁷We specify workers' preferences as log-additive in wages and amenities. Recent work by [Dube et al. \(2022\)](#) considers a model in which log wages and amenities are non-separable. While economically interesting, non-separability is empirically challenging as the labor supply curve will no longer be iso-elastic.

⁸We only require that η_{ijt} is independently distributed across firms and workers within each cross-section t ; η_{ijt} may be arbitrarily persistent within a worker-firm pair over time.

⁹As long as worker quality does not change in response to winning a procurement auction, it is straightforward to extend the model and the empirical analysis to allow for differences in worker quality (see [Lamadon et al., 2022](#)).

¹⁰The strategically-small firm assumption has been relaxed in the models considered by [Berger et al. \(2022\)](#), [Chan et al. \(2023\)](#), and [Jarosch et al. \(2023\)](#). However, identification is difficult in models with strategic interactions in the wage-setting. Recent papers by [Roussille and Scuderi](#)

Given these assumptions, the number of workers who accept a job at firm j at time t for a posted wage offer W_{jt} is $L_{jt} = (W_{jt}G_{jt}/\Xi_t)^{1/\theta}$, where G_{jt} captures the vertical differentiation due to firm-specific amenities and Ξ_t captures aggregate labor supply factors in the relevant market.¹¹ In the baseline analysis, we consider the entire US construction industry to be the relevant labor market. We also perform several specification checks to show that our findings are robust to alternative definitions of the labor market.

Our empirical analysis focuses on the inverse labor supply curve facing firm j at time t , which is given by

$$W_{jt} = L_{jt}^\theta U_{jt}, \quad (2)$$

where $U_{jt} \equiv \Xi_t/G_{jt}$. The labor supply elasticity facing the firm is $1/\theta$, so labor supply becomes more inelastic when idiosyncratic tastes are more dispersed. Given the assumption of strategically-small firms, the marginal wage changes at one firm do not impact aggregate labor supply factors (that is, $\frac{\partial \Xi_t}{\partial W_{jt}} \approx 0$).

For the empirical analysis, it is useful to decompose $\log U_{jt}$ into an aggregate component, a firm-specific fixed component, and a firm-specific time-varying component. Denoting $u_{jt} \equiv \log U_{jt}$, $g_{jt} \equiv \log G_{jt}$, and $\xi_t \equiv \log \Xi_t$, it follows that $u_{jt} = -g_{jt} + \xi_t$. Furthermore, we can write $-g_{jt} \equiv \psi_j + \nu_{jt}$, which is without loss of generality since we can simply define $\psi_j \equiv \mathbb{E}[-g_{jt}|j]$ and $\nu_{jt} \equiv -g_{jt} - \psi_j$. Then, denoting $w_{jt} \equiv \log W_{jt}$ and $\ell_{jt} \equiv \log L_{jt}$, log wages are given by

$$w_{jt} = \theta \ell_{jt} + u_{jt} = \theta \ell_{jt} + \psi_j + \nu_{jt} + \xi_t. \quad (3)$$

Letting Δ indicate differences over time, changes in log wages are thus

$$\Delta w_{jt} = \theta \Delta \ell_{jt} + \Delta \nu_{jt} + \Delta \xi_t, \quad (4)$$

where the time-invariant component ψ_j does not appear in differences over time.

Equation (4) highlights that wages change over time for three reasons. All else equal, the firm needs to pay higher wages to hire more workers, as captured by $\theta \Delta \ell_{jt}$; the firm must pay higher wages to keep the same number of workers if it experiences

(2022) and Sharma (2022) test for strategic interactions in wage-setting, but do not find evidence of such interactions.

¹¹Formally, $\Xi_t \equiv (\bar{W}_t/\bar{L}_t)^\theta$, where \bar{L}_t is the total number of workers in the market and $\bar{W}_t \equiv \sum_{j'} (W_{j't}G_{j't})^{1/\theta}$ is the price index of labor.

a negative shock to labor supply, as captured by $\Delta\nu_{jt}$; and the firm needs to pay more to keep the same number of workers if aggregate labor supply declines or the price index of labor rises, as captured by $\Delta\xi_t$.

2.2 Firm Technology and Product Demand

The production side of the model incorporates two types of product markets in which construction firms may participate: the private market and the government market, where government projects are procured through auctions. It is important to account for the government market for three reasons. First, it provides a more accurate representation of firms' production choices. Second, it lets us draw policy implications regarding the incidence of government expenditure on construction projects. Third, variation in bidding for procurement projects will be important to identify key model parameters.

We begin by specifying the production technology, before describing the private product market, taking the outcomes of the government market as given, and then discuss optimal bidding in the government market. Following [Ackerberg et al. \(2015\)](#), the production function (in physical units) is

$$Q_{jt} = \min\{\Omega_{jt}L_{jt}^{\beta_L}K_{jt}^{\beta_K}, \beta_M M_{jt}\} \exp(e_{jt}), \quad (5)$$

where Ω_{jt} denotes total factor productivity (TFP), K_{jt} denotes capital, M_{jt} denotes intermediate inputs, and e_{jt} represents measurement error. We assume that firms can rent capital at constant price p_K . While the assumption of a rental market for capital is standard in the literature, it may be a fairly good description of the construction industry, which heavily utilizes rental equipment and machinery. We also assume the market for intermediate inputs is competitive with constant price p_M .

Our Leontief functional form in equation (5) imposes strong complementarity between labor and intermediate inputs, while allowing for substitutability between labor and capital. This assumption may be relatively reasonable for the construction industry, where greater capital expenditure (i.e., renting more efficient equipment and machinery) may substitute for labor, but labor cannot take the place of concrete, asphalt, wood, and other materials required to construct a bridge or road.¹² The

¹²The Leontief functional form appears broadly consistent with the standard construction cost

Leontief functional form implies a zero elasticity of substitution between labor and intermediate inputs. As an alternative, in Online Appendix C, we solve, identify, and estimate the model with a Cobb-Douglas production function, which has elasticity of substitution of unity and thus allows for substitutability between labor and intermediate inputs.¹³ As shown in Section 6.2, the key empirical results are broadly similar when using the Cobb-Douglas production function, implying that the Leontief functional form is not crucial to our findings.

Given the technology in equation (5), the construction firms may choose to produce in two product markets. The first is the market for private projects, which we denote H . Specifically, firm j at time t posts a price P_{jt}^H at which it is willing to produce in the market for private projects. Consumers have idiosyncratic preferences over producers. Consumer i 's utility from purchasing from firm j at time t is $U_{ijt}^H = -\log P_{jt}^H + \omega_{ijt}$. We parameterize the distribution of ω_{ijt} as i.i.d. T1EV with dispersion $0 \leq \epsilon \leq 1$. When ϵ is larger, horizontal producer differentiation becomes relatively more important, as ω_{ijt} has greater variability.

Given these assumptions, the quantity purchased from firm j at t for a posted price P_{jt}^H can be expressed as $Q_{jt}^H = (P_{jt}^H)^{-1/\epsilon} / \aleph$, where $\aleph \equiv \sum_{j'} (P_{j't}^H)^{-1/\epsilon}$ is the aggregate price index. Rearranging, $P_{jt}^H = p_H (Q_{jt}^H)^{-\epsilon}$, where $p_H \equiv \aleph^{-\epsilon}$ and $-1/\epsilon$ is the product demand elasticity.¹⁴ This implies private market revenues are

$$R_{jt}^H = P_{jt}^H Q_{jt}^H = p_H (Q_{jt}^H)^{1-\epsilon}. \quad (6)$$

Denoting $r_{jt}^H \equiv \log R_{jt}^H$ and $q_{jt}^H \equiv \log Q_{jt}^H$, it follows that

$$r_{jt}^H = \log p_H + (1 - \epsilon) q_{jt}^H, \quad (7)$$

estimation handbook (RSMMeans, 2008). This handbook provides task-specific construction cost estimates for the typical crew (labor and equipment needed) per unit of material. While crew choices may vary depending on the contractor, material input requirements are fixed for each task.

¹³Castro-Vincenzi and Kleinman (2022, p.27) review the literature estimating the elasticity of substitution between intermediate materials and labor. All of the papers find an elasticity of substitution between 0 and 1, which are the lowest (Leontief) and highest (Cobb-Douglas) elasticities that we consider.

¹⁴In the private product market, we assume firms are “strategically small” in the sense that they view themselves as infinitesimal within the market (i.e. $\frac{\partial \aleph}{\partial P_{jt}^H} \approx 0$). Thus, $\frac{\partial \log Q_{jt}^H}{\partial \log P_{jt}^H} = -1/\epsilon$, so $-1/\epsilon$ is the price elasticity of demand. This product demand curve and “small” firm assumption are assumed by Dixit and Stiglitz (1977) and a large subsequent literature on monopolistic competition.

so $1-\epsilon$ can be interpreted as the revenue elasticity of output in the private market.¹⁵

In addition to the private product market, firms may participate in the market for government projects, denoted by G . Output for the government market is denoted Q_{jt}^G . Firm j produces total output $Q_{jt} = Q_{jt}^H + Q_{jt}^G$ simultaneously across both markets using the production function in equation (5). We denote $D_{jt} = 1$ if firm j receives a procurement contract at t and $D_{jt} = 0$ otherwise. If firm j does not receive a procurement contract ($D_{jt} = 0$), it does not produce in the government market ($Q_{jt}^G = 0$). If firm j receives a procurement contract ($D_{jt} = 1$), it must produce exactly \bar{Q}^G in the government market ($Q_{jt}^G = \bar{Q}^G$), where \bar{Q}^G is set by the government. The quantity produced by firm j in the government market can then be expressed as $Q_{jt}^G = \bar{Q}^G D_{jt}$. The allocation of procurement contracts to firms as well as the revenues received from procurement projects are determined through first-price sealed-bid auctions, which we describe below.

2.3 Firm's Problem and Optimal Behavior

We model firm behavior as a two-stage problem which we solve backwards. In the first stage, a firm submits 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. In the second stage, the firm chooses inputs to maximize profit from total production, taking as given the outcome of the procurement auction. Production in both private and government projects occurs simultaneously at the end of the second stage.

We now solve for the optimal private market behavior of firm j if it receives a procurement contract in the government market as well as if it does not. Denote profit excluding procurement revenue by π_{1jt}^H if $D_{jt} = 1$ and π_{0jt}^H if $D_{jt} = 0$. In order to obtain a procurement contract, firms place bids in auctions. Denote firm j 's bid in year t by Z_{jt} . Total profit is then $\pi_{1jt} = Z_{jt} + \pi_{1jt}^H$ if the firm receives a procurement contract, and $\pi_{0jt} = \pi_{0jt}^H$ otherwise. Observed profit is $\pi_{jt} = \pi_{1jt} D_{jt} + \pi_{0jt} (1 - D_{jt})$.

¹⁵Our derivations in the text focus on $\epsilon > 0$. Online Appendix B provides derivations with perfect competition, $\epsilon = 0$. As discussed in Section 6.2, $\epsilon = 0$ is at odds with our empirical findings.

Given \bar{Q}^G and $D_{jt} = d$, the firm's second stage problem is to hire labor L_{djt} , purchase intermediate inputs M_{djt} , and rent capital K_{djt} to maximize profits,

$$\pi_{djt}^H = R_{djt}^H - W_{djt}L_{djt} - p_M M_{djt} - p_K K_{djt}, \quad (8)$$

for $d = 0, 1$, subject to the labor supply curve (equation 2), the production function (equation 5), the private market revenue curve (equation 6), the price of intermediate inputs (p_M), the price of capital (p_K), and that the government project is fulfilled by the procured firm ($Q_{1jt} \geq \bar{Q}^G$).

We now use the firm's first-order conditions to characterize the firm's private market behavior. The first-order condition for capital implies a composite production function,

$$Q_{jt} = \min\{\Phi_{jt}L_{jt}^\rho, \beta_M M_{jt}\} \exp(e_{jt}), \quad (9)$$

where $\Phi_{jt} \equiv \Omega_{jt} \left[\frac{\beta_K}{\beta_L} \frac{(1+\theta)U_{jt}}{p_K} \right]^{\beta_K}$ is composite TFP, $\rho \equiv (1+\theta)\beta_K + \beta_L$ is the composite returns to labor, and e_{jt} reflects independent measurement error in output.¹⁶ We refer to Online Appendix A.1 for the derivation of the composite production function in equation (9).

Given the composite production function, the first-order condition for intermediate inputs implies

$$X_{jt} = \frac{p_M}{\beta_M} Q_{jt} = \frac{p_M}{\beta_M} L_{jt}^\rho \Phi_{jt}, \quad (10)$$

where $X_{jt} \equiv p_M M_{jt}$ denotes expenditure on intermediate inputs. Letting $x_{jt} \equiv \log X_{jt}$ and $\phi_{jt} \equiv \log \Phi_{jt}$, and defining $\kappa_X \equiv \log(p_M/\beta_M)$, it follows that

$$x_{jt} = \kappa_X + \rho \ell_{jt} + \phi_{jt}. \quad (11)$$

Thus, our model implies a closed-form relationship between expenditure on intermediate inputs and labor, which will prove useful for identifying ρ .

Combining the product demand curve in equation (7) with the first-order condition for intermediate inputs in equation (11) yields

$$r_{jt} = \kappa_R + (1 - \epsilon) x_{jt} + (1 - \epsilon) e_{jt} \quad \text{if } D_{jt} = 0, \quad (12)$$

¹⁶One way to motivate that the measurement error e_{jt} is independent is to suppose that firms choose intermediate inputs in the second stage of period t before they observe the idiosyncratic shock to output at the end of period t , as assumed by [Akerberg et al. \(2015\)](#) and a large literature in industrial organization.

where $\kappa_R \equiv \log p_H + (1-\epsilon) \log (\beta_M/p_M)$, which shows that revenues are log-linear in intermediate input expenditures, with coefficient $1-\epsilon$, among firms that are only producing for the private market ($D_{jt} = 0$).

The first-order condition with respect to labor implies the following relationship between private market revenues and expenditures on labor and intermediate inputs:

$$R_{djt}^H (1 - \epsilon) \frac{Q_{djt}}{Q_{djt}^H} = \frac{1 + \theta}{\beta_L} B_{djt} + X_{djt}, \quad (13)$$

where $B_{jt} \equiv L_{jt}W_{jt}$ denotes the firm's wage bill. This equation will prove useful for identifying ϵ and β_L . The derivation of this equation and several important implications of the first-order conditions for labor are reported in Online Appendix A.2; we briefly summarize implications here.

One implication of the first-order condition for labor is that both winners and losers of procurement contracts always produce strictly positive output for the private market ($Q_{djt}^H > 0$, $d = 0, 1$). This follows from the fact that firms have market power ($\epsilon > 0$), 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 receives a procurement contract versus if it does not ($Q_{1jt} > Q_{0jt}$). Another implication is that the government project crowds-out private projects for firm j ($Q_{1jt}^H < Q_{0jt}^H$) if $1 + \theta > \rho$, and conversely, crowds-in private projects ($Q_{1jt}^H > Q_{0jt}^H$) 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 would induce greater private production if there is increasing return to scale (in labor and capital), $\rho > 1$. Thus, the magnitude of $1 + \theta$ relative to ρ determines if receiving a procurement contract leads to crowd-out or crowd-in of private market output for firm j .

2.4 Firm's Optimal Bidding for Government Procurements

To complete the model, we now specify how procurement contracts are allocated to firms and the determination of procurement revenues. Firms choose bids in consideration of their opportunity costs. The opportunity cost of receiving a procurement con-

tract from the government is the difference in private market profits between receiving no government contract and receiving a government contract. Formally, denote the opportunity cost by $\sigma_{u_{jt}}(\phi_{jt}) \equiv \pi_{0jt}^H - \pi_{1jt}^H$, where the notation emphasizes that firm productivity ϕ_{jt} is the only source of heterogeneity in the opportunity cost, conditional on amenities u_{jt} .¹⁷ The opportunity cost of winning a procurement contract is strictly positive, $\sigma_{u_{jt}}(\phi_{jt}) > 0$, as the firm would have received positive revenues if it sold the output quantity \bar{Q}^G to the private market instead of the government market.

In the procurement auction, bidders observe common information about the size of the project, \bar{Q}^G , the number of bidders, I , and the amenities of each bidding firm, u_{jt} . The distribution of TFP conditional on amenities, $(\phi_{jt}|u_{jt} = u) \sim \tilde{F}_u(\cdot)$, is assumed to be i.i.d. and known by all firms, and induces an i.i.d. distribution of opportunity costs $\sigma_u(\phi_{jt}) \sim F_u(\cdot)$.¹⁸ Revenue from winning the auction is the winning bid, Z_{jt} . The difference between the benefit and the opportunity cost of winning an auction with bid Z_{jt} is thus $Z_{jt} - \sigma_u(\phi_{jt})$. Conditional on amenities $u_{jt} = u$, a firm with TFP ϕ_{jt} chooses the optimal bid Z_{jt} to solve

$$\max_{Z_{jt}} \underbrace{(Z_{jt} - \sigma_u(\phi_{jt}))}_{\text{payoff}} \underbrace{\Pr(D_{jt} = 1|Z_{jt})}_{\text{probability of winning}}.$$

The first term is the payoff to winning an auction, which is increasing in Z_{jt} , while the second term is the probability of winning an auction, which is decreasing in Z_{jt} . Thus, the firm faces the usual trade-off in an auction between profits if one wins and the probability of winning.

The firm's optimal bidding strategy in the procurement auction is

$$s_u(\phi_{jt}) = \sigma_u(\phi_{jt}) \delta_u(\phi_{jt}), \text{ where } \delta_u(\phi_{jt}) \equiv 1 + \frac{\int_{\sigma_u(\phi_{jt})}^{\bar{\sigma}} [1 - F_u(\tilde{\sigma})]^{I-1} d\tilde{\sigma}}{\sigma_u(\phi_{jt}) [1 - F_u(\sigma_u(\phi_{jt}))]^{I-1}}. \quad (14)$$

We can interpret $\delta_u \geq 1$ as the bid markup relative to the opportunity cost. When $\delta_u = 1$, each firm's optimal bid equals its opportunity cost, so each firm makes zero economic profit from receiving a procurement contract. As the number of auction

¹⁷The profit function for auction winners depends also on the size of the government project \bar{Q}^G , so the opportunity cost also depends on \bar{Q}^G . For notational convenience, we suppress this dependence.

¹⁸We require that TFP is i.i.d. across firms within each cross-section t , though TFP may be arbitrarily persistent within a firm over time.

participants I declines, δ_u rises, so firms that receive procurement contracts extract greater profits in the government market when there is less competition. Since Z_{jt} exceeds $\sigma_u(\phi_{jt})$ due to finite I and the bidding strategy $s_u(\phi_{jt})$ is strictly increasing in the opportunity cost $\sigma_u(\phi_{jt})$, equation (14) defines the unique symmetric equilibrium (Milgrom and Weber, 1982; Maskin and Riley, 1984).¹⁹ The winner of the auction is determined as

$$D_{jt} = 1 \{s_u(\phi_{jt}) < s_u(\phi_{j't}), \forall j' \neq j \text{ such that } j, j' \in \mathcal{J}_t\},$$

where \mathcal{J}_t is the set of firms participating in auction t . This expression makes clear that the winner of a procurement contract is selected on TFP.

Before proceeding, it is useful to note that we have assumed auctions are symmetric. Though this assumption will not be crucial for our identification strategy discussed below, it is convenient for expositional and computational purposes, as symmetric auctions are easier to solve. Furthermore, symmetry is a standard assumption in the empirical auction literature (Athey and Haile, 2007). Nevertheless, in our empirical application, we provide a robustness check which relaxes this symmetry assumption.

2.5 Worker and Firm Rents

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 for the private market (in the absence of procurement projects) and the additional rents generated from receiving a procurement contract or, equivalently, the incidence of government procurement.

In our model, 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). The rents of worker i , V_{it} , derived from the current choice of firm, j , are defined implicitly by,

¹⁹One potential concern is that firms may collude to achieve bid revenues greater than those predicted by our first-price sealed-bid auction model. In Online Appendix Figure A.1, we apply the collusion test of Chassang et al. (2022) to each of the 28 states in our data separately, finding no evidence of collusion.

$$\mathcal{U}_{it}(j, W_{jt} - V_{it}) = \max_{j' \neq j} \mathcal{U}_{it}(j', W_{j't}).$$

This definition of rents captures that the average worker choosing firm j may be far from the margin of indifference and would maintain the same choice even if her current firm offered significantly lower wages.

Given our specification of preferences (equation 1), we prove in Online Appendix A.3 that it is possible to aggregate the rents across workers to get the measure of the total worker rents at firm j , V_{jt} , if it offers wage W , as follows,

$$V_{jt}(W) = \frac{W_{jt}L_{jt}(W)}{1 + 1/\theta} = \frac{B_{jt}(W)}{1 + 1/\theta}$$

where $L_{jt}(W)$ is the number of workers in firm j and period t if it offers wage W (determined by equation 2), and $B_{jt}(W) \equiv WL_{jt}(W)$ is the corresponding wage bill. Intuitively, V_{jt} can be interpreted as the workers' willingness-to-pay to stay at the current firm j , which is greater when the labor supply curve is steeper (i.e., when θ is greater).²⁰

When analyzing the incidence of procurements, recall that there are two potential wage offers: W_{0jt} if the firm loses the procurement auction, and W_{1jt} if the firm wins the procurement auction. Defining $V_{0jt} \equiv V_{jt}(W_{0jt})$ and $V_{1jt} \equiv V_{jt}(W_{1jt})$, the incidence of procurements on workers, $V_{\Delta jt}$, is given by,

$$\underbrace{V_{\Delta jt}}_{\text{Incidence}} \equiv \underbrace{V_{1jt}}_{\text{Rents for winners}} - \underbrace{V_{0jt}}_{\text{Rents for losers}} = \underbrace{\frac{B_{1jt} - B_{0jt}}{1 + 1/\theta}}_{\text{Incidence}},$$

where $B_{0jt} \equiv B_{jt}(W_{0jt})$ and $B_{1jt} \equiv B_{jt}(W_{1jt})$. Furthermore, we can measure the incidence for “incumbents” (i.e., workers who accept both offered wages W_{0jt} and W_{1jt} by firm j) versus “new hires” (i.e., workers who accept offered wage W_{1jt} but reject offered wage W_{0jt} by firm j). For an incumbent worker, it is clear from the definition of rents above that the incidence is simply the wage gain, $W_{1jt} - W_{0jt}$. Thus, the total incidence of procurements on all workers in the firm can be decomposed as,

²⁰Note that B_{jt} is a function of $1/\theta$, so these expressions do not imply that rents increase for workers as labor supply becomes more inelastic. When we quantify the impacts of a change in market power below, we will show that worker rents decrease when labor supply becomes more inelastic due to the wage bill decrease.

$$\underbrace{V_{\Delta jt}}_{\text{Incidence}} = \underbrace{L_{0jt}(W_{1jt} - W_{0jt})}_{\text{Incidence for incumbents}} + \underbrace{W_{1jt}(L_{1jt} - L_{0jt}) - \frac{B_{1jt} - B_{0jt}}{1 + \theta}}_{\text{Incidence for new hires}},$$

where $L_{0jt} \equiv L_{jt}(W_{0jt})$ and $L_{1jt} \equiv L_{jt}(W_{1jt})$; see Online Appendix A.3 for additional details. Thus, the incidence for incumbents is the wage change multiplied by the number of incumbent workers. The incidence for new hires is the wage bill of new hires minus the wage bill required to make them indifferent between the new and initial firm choices.

As our measure of firm rents, we use profits. There are three relevant measures of profits. First, π_{0jt} is the profit that the firm captures from production for the private market if it does not receive a procurement contract. Second, π_{1jt} is the profit the firm captures from joint production for the government and private markets if it receives the procurement contract. Third, $\pi_{\Delta jt} \equiv \pi_{1jt} - \pi_{0jt}$ is the additional rents earned by the firm from receiving the procurement contract. They are related by

$$\underbrace{\pi_{1jt}}_{\text{Firm rents for winners}} = \underbrace{\pi_{0jt}}_{\text{Firm rents for losers}} + \underbrace{\pi_{\Delta jt}}_{\text{Incidence on firms}}.$$

It is important to observe that profits do not necessarily represent ex-ante rents for the employer. Suppose, for example, that each employer initially invests in amenities offered to the workers by deciding on the firm’s location or working conditions. Workers’ heterogeneous preferences over those amenities give rise to wage-setting power, which employers can use to extract additional profits or rents. Thus, the existence of such ex-post rents could simply be returns to costly ex-ante choices of amenities. Additionally, profits from procurement projects may in part reflect a fixed cost of entry to the auction, e.g., the cost of obtaining a license. While the presence of a fixed entry cost will affect the interpretation of profits, it is possible to show that it will not affect identification of model parameters.

3 Double Market Power and its Implications

In our model, a firm is defined to have “double market power” if it can profitably set both the price above the marginal cost and the wage below the marginal revenue product of labor. We now theoretically examine the implications of this double market

power for the wage-setting and price-setting decisions of the firm. We first characterize the total markdown of wages and total markup of prices in the presence of double market power. Next, we show that the impacts of an increase in market power in one market will be attenuated by the existence of market power in the other market.

As these theoretical results may be of interest in markets other than the construction sector, we consider a version of the model from Section 2 that allows for a more flexible production function and omits the government market for procurements. In the empirical analyses in Sections 5-8, however, we return to the model of the construction industry in Section 2, quantifying the double market power and its implications.

3.1 Markdown of Wages and Markup of Prices

As in the previous section, we suppose that the firm maximizes profits (equation 8) subject to the labor supply curve (equation 2) and the product demand curve (equation 6). However, we remove the government market and allow for a flexible production function, $Q_{jt} = f_{jt}(L_{jt}, M_{jt}, K_{jt})$, that generalizes equation (5).

Consider, for now, the case where f_{jt} has positive elasticity of substitution between labor and other inputs. Denoting the marginal revenue product of labor by $\text{MRPL}_{jt} \equiv \frac{\partial(P_{jt}Q_{jt})}{\partial L_{jt}}$ and the marginal cost of labor by $\text{MCL}_{jt} \equiv \frac{\partial(W_{jt}L_{jt})}{\partial L_{jt}}$, the first-order condition is,

$$\underbrace{(1 - \epsilon) \times P_{jt} \text{MPL}_{jt}}_{\text{MRPL}_{jt}} = \underbrace{(1 + \theta) \times W_{jt}}_{\text{MCL}_{jt}} \quad (15)$$

where $\text{MPL}_{jt} \equiv \frac{\partial}{\partial L_{jt}} f_{jt}(L_{jt}, M_{jt}, K_{jt})$ denotes the marginal product of labor. Rearranging the first-order condition, we have the following result:

Lemma 1. *The optimal markdown and markup are given by,*

$$W_{jt} = \overbrace{(1 + \theta)^{-1}}^{\text{markdown}} \times \text{MRPL}_{jt} \quad \text{and} \quad P_{jt} = \overbrace{(1 - \epsilon)^{-1}}^{\text{markup}} \times \frac{\text{MCL}_{jt}}{\text{MPL}_{jt}} \quad (16)$$

Lemma 1 confirms the usual intuition that less elastic labor supply (greater θ) leads to a greater markdown of wages, and less elastic product demand (greater ϵ) leads to a greater markup of prices.

What is not directly evident from Lemma 1 is how the labor supply and product demand elasticities interact to determine the wage-setting and price-setting decisions of the firm. The following proposition highlights these interactions by expanding the MRPL and MCL in Lemma 1:

Proposition 1. *The optimal wage satisfies,*

$$W_{jt} = \underbrace{\overbrace{(1 + \theta)^{-1}}^{\text{markdown}} \times \overbrace{(1 - \epsilon)}^{\text{inverse markup}}}_{\text{double markdown}} \times \underbrace{P_{jt} \text{MPL}_{jt}}_{\text{value of MPL}} \quad (17)$$

and the optimal price satisfies,

$$P_{jt} = \underbrace{\overbrace{(1 - \epsilon)^{-1}}^{\text{markup}} \times \overbrace{(1 + \theta)}^{\text{inverse markdown}}}_{\text{double markup}} \times \underbrace{\frac{W_{jt}}{\text{MPL}_{jt}}}_{\text{prod.-adjusted wage}} \quad (18)$$

A key insight from Proposition 1 is that not only upward-sloping labor supply but also downward-sloping product demand are relevant and distinct sources of labor market power, as measured by the markdown of wages relative to the value of MPL. Similarly, product market power, as measured by the markup relative to the productivity-adjusted wage, is determined not only by the elasticity of product demand but also the elasticity of labor supply.

To provide intuition for Proposition 1, we illustrate the double markdown and double markup in Figure 1. Figures 1a-1b focus on the markdown of wages. In Figure 1a, there is no product market power ($\epsilon = 0$), so MRPL and the value of MPL are identical and there is no markup. In this case, the only markdown of the wage is the usual markdown, $(1 + \theta)^{-1}$, and the optimal employment and wage are given by (L_0, W_0) . In Figure 1b, there is product market power ($\epsilon > 0$), generating a wedge between MRPL and the value of MPL whose size is determined by the inverse markup. As a result of this wedge, MRPL intersects MCL at lower labor and wage levels (L_1, W_1) . Thus, we see that the wage is marked down further below the value of MPL if there is product market power.

Figures 1c-1d provide an analogous illustration of the double markup of prices. In Figure 1c, there is no labor market power ($\theta = 0$), so MCL and the wage are

identical and the only markup is $(1 - \epsilon)^{-1}$. In Figure 1d, there is labor market power ($\theta > 0$), generating a wedge between MCL and the wage whose size is determined by the inverse markdown, leading to an additional markup relative to the wage.

We now extend the double market power results from Proposition 1 to the case of a Leontief production function in which labor and materials are perfect complements, following equation (5). Under perfect complementarity, optimal intermediate materials will adjust in response to a change in labor, resulting in an extra term in the marginal cost that accounts for the price of and returns to intermediate materials:

$$\underbrace{(1 - \epsilon) \times P_{jt} \text{MPL}_{jt}}_{\text{MRPL}_{jt}} = \underbrace{(1 + \theta) \times W_{jt}}_{\text{MCL}_{jt}} + \underbrace{\frac{p_M}{\beta_M} \text{MPL}_{jt}}_{\text{Leontief adjustment}}$$

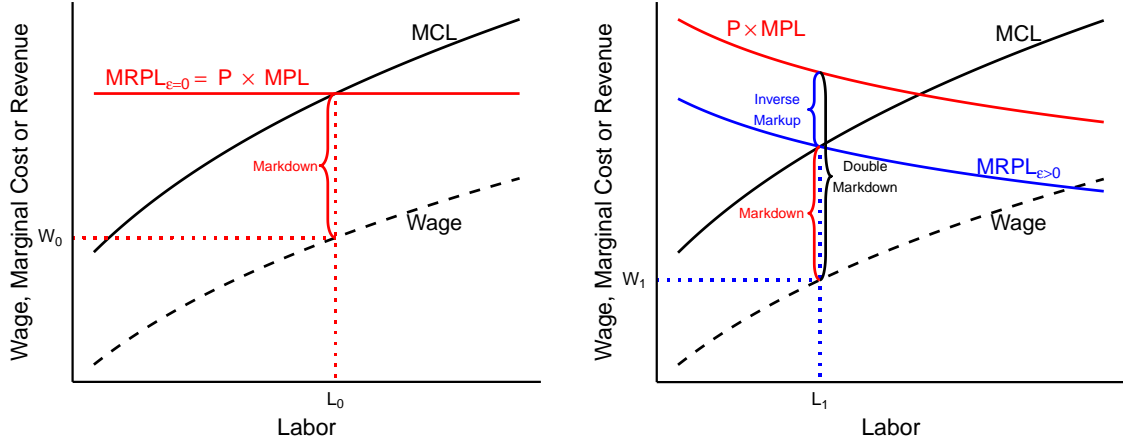
Accounting for the Leontief adjustment in the first-order condition, we have the following:

Proposition 2. *The optimal wage and price satisfy,*

$$W_{jt} = \underbrace{(1 + \theta)^{-1} \times (1 - \epsilon)^{-1}}_{\text{double markdown}} \times \underbrace{P_{jt} \text{MPL}_{jt}}_{\text{value of MPL}} - \underbrace{(1 + \theta)^{-1} \times \frac{p_M}{\beta_M} \text{MPL}_{jt}}_{\text{Leontief adjustment to wage}} \quad (19)$$

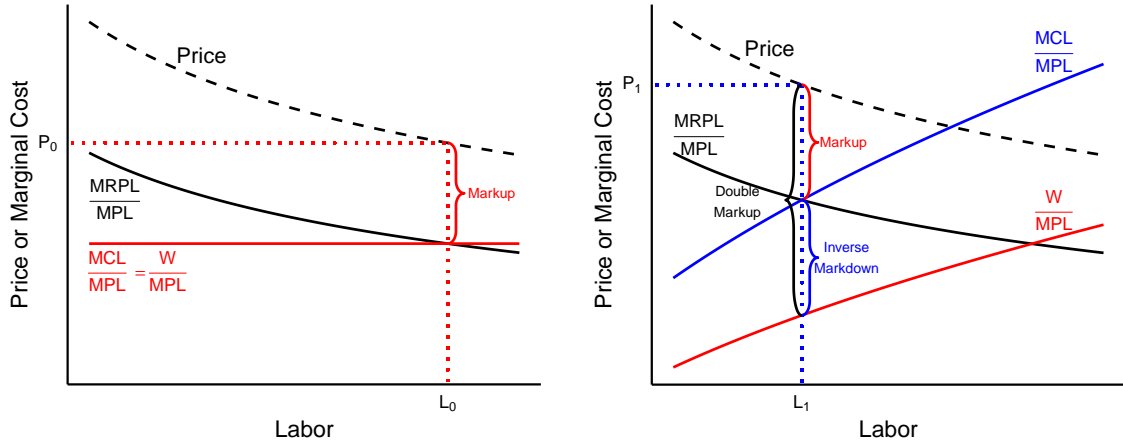
$$P_{jt} = \underbrace{(1 - \epsilon)^{-1} \times (1 + \theta)}_{\text{double markup}} \times \underbrace{\frac{W_{jt}}{\text{MPL}_{jt}}}_{\text{prod.-adjusted wage}} + \underbrace{(1 - \epsilon)^{-1} \times \frac{p_M}{\beta_M}}_{\text{Leontief adjustment to price}} \quad (20)$$

When interpreting Propositions 1 and 2, it is useful to observe that our assumptions give constant labor supply and product demand elasticities across firms within the construction sector. Allowing for heterogeneous labor supply or product demand elasticities would be economically interesting but make identification difficult. For example, such an extension may allow markups and markdowns to vary across firms and potentially be correlated. This could have important implications for analyses of misallocation: If the firms with greater markups also have greater markdowns, misallocation may be amplified.



(a) Markdown without Product Market Power

(b) Markdown with Product Market Power



(c) Markup without Labor Market Power

(d) Markup with Labor Market Power

Figure 1: Visualizing the Double Markdown and Double Markup

Notes: This figure visualizes how wages are marked down and prices are marked up due to the interaction between labor and product market power. It is constructed by simulating the solution to the firm's problem, omitting the procurement auctions. For simplicity, we parameterize the production function as $Q_{jt} = \Phi_{jt} L_{jt}$, although the results in the text hold for more general production functions. In subfigures a-b, the curve labeled Wage is the inverse labor supply curve, and the marginal cost of labor curve is related to the wage curve by $MCL = (1 + \theta) \times Wage$. The curve labeled $P \times MPL$ is the value of MPL, which is related to the MRPL by $MRPL = (1 - \epsilon) \times P \times MPL$. In subfigures c-d, the curve labeled Price is the inverse product demand curve, and the productivity-adjusted MRPL curve is related to the price curve by $\frac{MRPL}{MPL} = (1 - \epsilon) \times Price$. The curve labeled $\frac{W}{MPL}$ is the productivity-adjusted wage, which is related to the productivity-adjusted MCL by $\frac{MCL}{MPL} = (1 + \theta) \times \frac{W}{MPL}$.

3.2 Expected Impacts of Changes in Market Power

We showed, in Proposition 1, that the presence of double market power matters for both the total markdown of the wage and the total markup of the price. A natural question is whether and how the presence of double market power affects the expected impact of an increase in labor market power or an increase in product market power. To analyze this question, we need to be able to isolate the impact of a change in market power in a given market, all other things being equal.

At first glance, it may seem that the impact of product (or labor) market power on prices (or wages) can be studied by simply changing ϵ (or θ) in equation (18) (or 17), holding all other terms fixed. However, such an approach would give a misleading answer. An increase in ϵ has two distinct impacts: It both reduces the level of demand the firm is facing at any given price (a level effect, reflected by a behavioral response in the wage and MPL terms in equation 18), and it increases the profitability of pricing above marginal cost, since demand becomes less sensitive to price changes (a market power effect, changing the markup term in equation 18). Similarly, an increase in θ reduces the supply of labor the firm is facing at any given wage (a level effect, reflected by a behavioral response in the price and MPL terms in equation 17), and it raises the gains from marking down wages, since labor supply becomes less sensitive to wage changes (a market power effect, changing the markdown term in equation 17). Thus, increasing ϵ (or θ) is not the appropriate way to study a *ceteris paribus* increase in product (or labor) market power.

Instead, we will increase market power in the labor (or product) market by performing a compensated rotation of the labor supply (or product demand) curve. The compensated rotation eliminates the level effect, and thus isolates the impact of a change in market power in a given market, all other things being equal.²¹ In the labor market, this is done by increasing the inverse labor supply elasticity from an initial value θ^{Init} to a higher value θ^{New} , while making sure that labor supply at the initial wage is the same at θ^{New} as it was at θ^{Init} (by changing the location parameter in labor supply U). Similarly, the compensated rotation in the product market

²¹Our compensated rotation is closely related to Slutsky's definition of the substitution effect in demand analysis, which is defined with respect to the (hypothetical) income compensation that ensures initial consumption is still feasible after a price increase.

increases the inverse product demand elasticity from an initial value ϵ^{Init} to a higher value ϵ^{New} , while ensuring that product demand at the initial price is identical at ϵ^{New} as it was at ϵ^{Init} (by changing the location parameter in product demand p_H).

In Figures 2a-2b, we use the compensated rotation to show the expected impact of an increase in labor market power and how it depends on the presence or absence of product market power.²² The initial inverse labor supply elasticity is $\theta^{\text{Init}} > 0$, initial employment is L^{Init} , and initial wage is W^{Init} . In Figure 2a, product demand is perfectly elastic ($\epsilon = 0$). When the inverse labor supply curve rotates from the initial curve labeled $\text{Wage}^{\text{Init}}$ (corresponding to θ^{Init}) to the less elastic curve labeled Wage^{New} (corresponding to θ^{New}), the corresponding marginal cost of labor curve rises from the initial curve labeled MCL^{Init} to the new curve labeled MCL^{New} through the relationship $\text{MCL} = (1 + \theta) \times \text{Wage}$. Since the first-order condition equates MRPL and MCL, and MCL is higher for all values of labor as labor market power increases, it follows that the firm finds it profitable to reduce labor. Since the wage curve is upward sloping, a reduction in labor reduces the wage.

In Figure 2b, we perform exactly the same analysis as in Figure 2a, except that the firm now has product market power. To introduce product market power, we rotate the product demand curve by increasing ϵ from zero to a positive value, and then change the location parameter p_H such that the initial optimum is identical to Figure 2a. Next, we rotate the labor supply curve to increase labor market power. The expected impact of increased labor market power is qualitatively similar to Figure 2a: the MCL rises at all values of labor, so MRPL and MCL are equalized by a smaller choice of labor and wages. However, in contrast to Figure 2a, the output price rises as output declines in Figure 2b due to $\epsilon > 0$, so the MRPL is higher at the new optimal choice of labor in Figure 2b.

While the *qualitative* effects of increased labor market power are similar in Figures 2a and 2b, the *magnitude* of the expected decrease in labor and wages is stronger in Figure 2a due to the absence of product market power. The intuition for this result is straightforward: In response to an increase in labor market power, the firm wants to lower the wage it pays by reducing employment and, thereby, output. However,

²²The graphical results in Figures 2a-2b are accompanied by formal results in Online Appendix D.

the firm will choose to reduce employment and output less if it is facing downward-sloping demand ($\epsilon > 0$), as lower output increases the price it can charge and, thus, the marginal revenue product of labor. By contrast, in Figure 2a, the output price remains constant as output falls, giving the firm no product market power to exploit, and a greater reduction in labor is required in order for MRPL and MCL to be equalized, leading to a greater decline in the wage.

In Figures 2c-2d, we use the compensated rotation to show the expected impact of an increase in product market power and how it depends on the presence or absence of labor market power. The initial inverse product demand elasticity is $\epsilon^{\text{Init}} > 0$, initial employment is L^{Init} , and initial price is P^{Init} . In Figure 2c, labor supply is perfectly elastic ($\theta = 0$). When the inverse product demand curve rotates from the initial curve labeled $\text{Price}^{\text{Init}}$ (corresponding to ϵ^{Init}) to the less elastic curve labeled $\text{Price}^{\text{New}}$ (corresponding to ϵ^{New}), the corresponding productivity-adjusted MRPL curve falls from the initial curve labeled $\frac{\text{MRPL}^{\text{Init}}}{\text{MPL}}$ to the new curve labeled $\frac{\text{MRPL}^{\text{New}}}{\text{MPL}}$ through the relationship $\frac{\text{MRPL}}{\text{MPL}} = (1 - \epsilon) \times \text{Price}$. Since the first-order condition equates MRPL and MCL, and MRPL is lower for all values of labor as product market power increases, it follows that the firm finds it profitable to reduce labor. Since the price curve is downward sloping, a reduction in labor raises the price.

In Figure 2d, we perform exactly the same analysis as in Figure 2c, except that the firm now has labor market power. To introduce labor market power, we rotate the labor supply curve by increasing θ from zero to a positive value, and then change the location parameter U such that the initial optimum is identical to Figure 2c. Next, we rotate the product demand curve to increase product market power. The expected impact of increased product market power is qualitatively similar to Figure 2c: the MRPL falls at all values of labor, so MRPL and MCL are equalized by a smaller choice of labor and higher prices. However, in contrast to Figure 2c, the wage falls as output declines in Figure 2d due to $\theta > 0$, so the MCL is lower at the new optimal choice of labor in Figure 2d.

While the *qualitative* effects of increased product market power are similar in Figures 2c and 2d, the *magnitude* of the expected decrease in labor and increase in price relative to the initial optimum is stronger in Figure 2c due to the absence of labor market power. The intuition for this result is straightforward: if the firm experiences

an increase in product market power, it wants to increase the price it charges by reducing output and, thereby, employment. However, the firm will choose to reduce output and employment less if it is facing upward-sloping labor supply ($\theta > 0$), as lower employment decreases the wage it has to pay and, thus, the marginal cost of labor. By contrast, in Figure 2c, the wage remains constant as employment falls, giving the firm no labor market power to exploit, and a greater reduction in output is required in order for MRPL and MCL to be equalized, leading to a greater increase in prices.

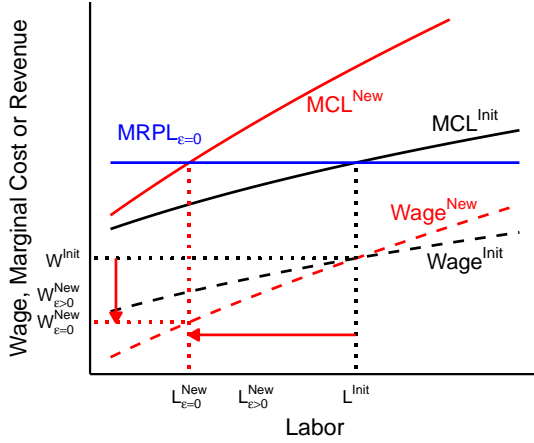
4 Data and Descriptive Evidence

4.1 Data Sources, Key Variables, and Sample Selection

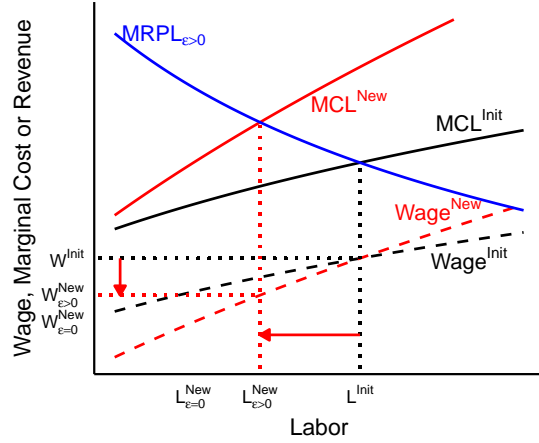
Our empirical analyses are based on a matched employer-employee panel data set for the period 2001-2015. The data set is formed by first linking business tax returns to worker-level tax returns, then merging this linked data set with procurement auction records.

Tax returns for firms and workers. Our business tax return data include balance sheet and other information from Forms 1120 (C-corporations), 1120S (S-corporations), and 1065 (partnerships). We link the business tax returns to Form W-2 (direct employee) and 1099 (independent contractor) worker-level 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 who are full-time equivalent (FTE) workers, by which we mean that their annual earnings from the primary employer are greater than the annualized full-time minimum wage in the year. Because firms sometimes use independent contractors, we also consider a broader measure of the workforce that includes any FTE independent contractors from Form 1099.²³

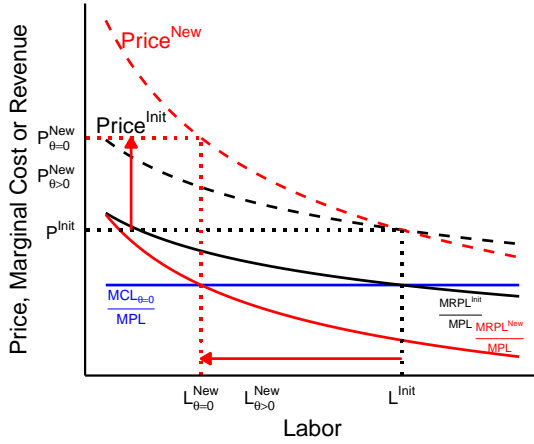
²³In Panel D of Online Appendix Table A.3, we show that the labor supply elasticity estimates do not materially change if we include or exclude independent contractors from the estimation sample.



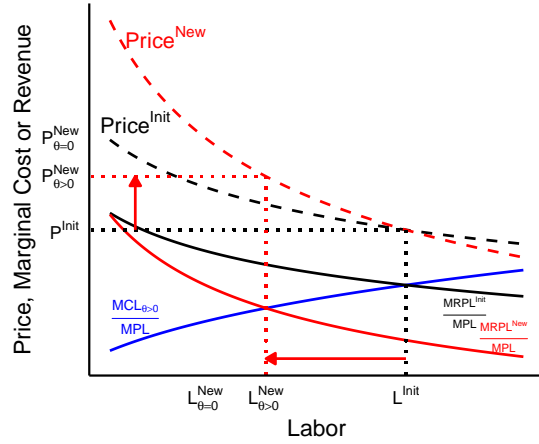
(a) Change in Labor Market Power without Product Market Power



(b) Change in Labor Market Power with Product Market Power



(c) Change in Product Market Power without Labor Market Power



(d) Change in Product Market Power with Labor Market Power

Figure 2: Expected Impacts of Changes in Market Power

Notes: Subfigures a-b present the expected impacts of increasing labor market power, while subfigures c-d present the expected impacts of increasing product market power. For simplicity, we parameterize the production function as $Q_{jt} = \Phi_{jt}L_{jt}$, although the results in the text hold for more general production functions. In both subfigures a-b, we consider the compensated rotation of the Wage curve (i.e. the inverse labor supply curve) around the initial choices (L^{Init}, W^{Init}) . In subfigures a-b, we denote the MRPL curve (in blue) by $MRPL_{\epsilon>0}$ for the case with product market power and $MRPL_{\epsilon=0}$ for the case without product market power. In both subfigures c-d, we consider the compensated rotation of the Price curve (i.e. the inverse product demand curve) around the initial choices (L^{Init}, P^{Init}) . In subfigures c-d, we denote the productivity-adjusted MCL curve (in blue) by $\frac{MCL_{\theta>0}}{MPL}$ for the case with labor market power and $\frac{MCL_{\theta=0}}{MPL}$ for the case without labor market power.

The key variables that we draw from the business tax returns are revenues, intermediate input expenditures, profits, and NAICS industry codes.²⁴ Revenues include those from business operations, excluding non-business-operation revenues such as dividends and capital gains. We follow [de Loecker et al. \(2020\)](#) in measuring intermediate input expenditures by the cost of goods sold, which includes variable costs associated with intermediate goods, transportation, and storage while excluding costs associated with overhead, durables, and labor.²⁵ Our measure of profits is earnings before interest, taxes, and depreciation (EBITD), which we construct following [Kline et al. \(2019\)](#).

The key variables we draw from worker-level tax returns are the number of employees in the firm and their earnings for the primary sample of workers. We also consider the number of employees and earnings when including independent contractors in the sample. Using the panel structure of the employer-employee data, we consider three measures of the earnings of the workers: mean earnings among all workers currently employed at the firm; mean earnings among stayers, which we define as workers employed at the same firm consistently from n years prior to the procurement auction until n years after; and mean earnings of new hires at the firm.

Information on procurement auctions. We obtain the new data set on procurement auctions from Bid Express (BidX.com), a service that facilitates online bidding for a number of states; state-specific Department of Transportation (DOT) websites; and submitting FOIA requests to state governments. The procurement projects broadly involve the construction and landscaping of local roads, bridges, and highways. Observations in this data set are at the auction-firm level, with variables on firm’s name and address as well as the firm’s bid and the auction date. In total, we recover the auction records from the DOTs of 28 states.²⁶ Construction firms bid in auctions in other states, so our auction sample includes construction firms

²⁴Detailed information about the variables in the tax data is provided in Online Data Supplement [S.3](#).

²⁵A potential concern is that firms in some industries (manufacturing and mining) include labor costs in the cost of goods sold. However, we consider firms in the construction industry, which do not (IRS Pub. 334).

²⁶Online Appendix Table [A.1](#) provides a summary of the procurement auction data sources. Online Data Supplement [S.1](#) provides step-by-step instructions on obtaining and preparing the auction records.

from nearly every state.²⁷ Our data show that these 28 DOTs allocated \$383 billion through 155,768 distinct auctions involving 16,697 bidders in 2010. There are more auctions than firms, and the same firms may compete in multiple auctions. One potential concern is that this results in collusion to achieve bid revenues greater than those predicted by our first-price sealed-bid auction model. In Online Appendix Figure A.1, we apply the collusion test of Chassang et al. (2022) to each of the 28 states in our data separately, finding no evidence of collusion in any state.

The DOTs are responsible for determining the nature of the project, including the blueprints, a 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 each DOT.

The awarding of a contract has two steps. The first step is qualification. In order to submit a bid, a firm must be pre-qualified by the DOT, which considers three aspects of the firm.²⁸ First, the DOT reviews financial statements (containing information about revenues, assets, and liabilities) to determine the firm’s ability to potentially complete the project. Second, the DOT considers the experience of the firm by reviewing information on previously completed projects and its current equipment and facilities. Third, the DOT may also deny pre-qualification if the firm has been found engaging in bid rigging or if it has systematically failed to comply with labor or safety regulations. Once approved, the firm is awarded a license to bid.

The second step is the sealed-bid auction, in which a qualified firm submits a bid without observing the bids of the other firms. Nearly all the auctions we consider are first-price auctions, where the procurement contract is awarded to the bidder with the lowest price.²⁹

²⁷In our data from BidX, we observe the firm’s business address as well as the address of the procurement project. In the BidX data, we find that 80% of bids are placed in auctions in the firm’s home state and 21% of bids are placed in auctions in the firm’s home commuting zone. In Panel E of Online Appendix Table A.3, we provide robustness checks in which we control for auction-specific or commuting zone-specific interactions in the effects, finding little sensitivity of the estimates.

²⁸Each state provides a detailed set of instructions for pre-qualification to bid, which is publicly available. For example, for Illinois, see <https://idot.illinois.gov/assets/uploads/files/doing-business/laws-&-rules/highways/construction/rules.pdf>.

²⁹A small fraction of the auctions consider dimensions in addition to price. One example is California, where 4% of the auctions consider both price and expected time to completion (Lewis and Bajari, 2011). Institutional details on procurement auctions with a secondary dimension of

We observe the bid of each firm for a given auction, not only the winner.³⁰ In the empirical applications discussed below, we will focus on recipients of procurement contracts who win a procurement auction for the first time at t . The mean procurement revenues for these first-time auction winners is \$2.7 million. We compare them to non-recipients that had never won an auction before t and placed a bid in the auction at t but lost. These sample restrictions are useful as they ensure that neither winners nor losers of auctions experience a procurement demand shock in the pre-period.

Merged data from tax returns and procurement auctions. To merge the auction data to the tax records, we use a fuzzy matching approach based on the firm’s name and location. For six states, we were able to not only obtain the name and address but also the federal Employer Identification Number (EIN) of the firm, allowing us to perform an exact match to tax records. We trained the algorithm on these six states before applying it to the other 22 states.³¹ We also provide an out-of-sample validation analysis in which we test the performance of the matching algorithm on publicly available pension data tax filings, finding that it performs well. Furthermore, we verify that our labor supply elasticity estimate does not change materially if we restrict the sample to the six states matched on EIN.

Online Appendix Table A.2 displays the sample sizes of firms and workers that participate in auctions in the year 2010. In 2010, 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. As a share of the national construction industry (as recorded in the 2010 tax records), our sample of 8,000 firms accounts for 12% of sales, 12% of employment, 10% of EBITD, 12% of intermediate input expenditures, and 13% of wage bidding are provided in Online Appendix E.4.

³⁰In the event that the firm that wins the procurement contract hires a contractor to complete the work, this will be captured in our measure of labor that includes contractors from Form 1099. In the event that the firm that wins the procurement contract passes the contract to a subcontractor, the completed work is sold to the primary contractor as intermediate inputs and thus captured in our intermediate inputs measure. Institutional details on the role of subcontracting in procurements are provided in Online Appendix E.5.

³¹Online Data Supplement S.2 explains how we trained and validated the linking algorithm used to merge the auction records to the tax returns.

payments. In a robustness check discussed in Section 6.1, we provide estimates for the entire construction industry, not only those firms that participate in procurement auctions.

4.2 Impacts of Winning a Procurement Auction

Before taking the model to the data, it is useful to describe how employment and wages change if a firm wins a procurement contract. To estimate these impacts, we run regression models which compare firms that are first-time winners of a procurement auction in year t ($D_{jt} = 1$) to firms that bid in the auction in year t but lose ($D_{jt} = 0$), before and after the auction. Let e denote an event time relative to t and y_{jt+e} denote an outcome for firm j . For each event time $e = -4, \dots, 4$, our baseline regression model can be expressed as,

$$y_{jt+e} = \underbrace{\sum_{e' \neq \bar{e}} 1_{e'=e} \mu_{te'}}_{\text{event time fixed effect}} + \underbrace{\sum_{j'} 1_{j'=j} \psi_{j't}}_{\text{firm fixed effect}} + \underbrace{\sum_{e' \neq \bar{e}} 1_{e'=e} D_{jt} \lambda_{te'}}_{\text{treatment status by event time}} + \underbrace{\epsilon_{jte}}_{\text{residual}}, \quad (21)$$

where λ_{te} is the impact of winning a procurement contract on the outcome variable for a particular pair (e, t) and $\bar{e} = -2$ is the omitted event time.³² The inclusion of firm fixed effects controls for any differences between winners and losers of auctions in their time-invariant characteristics (for example, fixed differences in composition of the workforce, amenities, or productivity), while the inclusion of event time fixed effects controls for any aggregate shocks experienced during this time interval.

Although the inclusion of fixed effects adjusts for differences across firms in levels, one could still be concerned that the auction winners in the treated group would have experienced different changes in outcomes over time than losers in the control group even in the absence of winning the auction. In particular, a firm is likely to bid less (and, thus, more likely to win the auction) if it experiences an increase in its productivity. To address this issue, it may be useful to restrict the sample used in estimating equation (21) to the control group of firms that make similar bids to the treated firms.

³²We will be estimating λ_{te} for all t and e and then average across t , using the delta method to compute standard errors (which are clustered at the firm level j to account for serial correlation). By doing so, we avoid the problem pointed out by Callaway and Sant’Anna (2020) that cohorts can be negatively weighted in pooled cohort two-way fixed-effect estimators of treatment effects.

To be concrete, define the loss margin as $\tau_{jt} \equiv (Z_{jt} - Z_{\iota}^*)/Z_{\iota}^*$ for a firm j that bids in auction ι at time t , where Z_{ι}^* is the winning bid in auction ι . When estimating equation (21), we restrict the sample to exclude the firms with large values of τ_{jt} . On the one hand, if we restrict the estimation sample so that the control group only includes firms with sufficiently small values of τ_{jt} , equation (21) can be interpreted as a local regression discontinuity design (RDD) that compares the firms that win the auction to the firms that almost win the auction.³³ On the other hand, if we include all auction losers in the control group (no restriction on τ_{jt}), then equation (21) becomes a difference-in-differences (DiD) design, comparing growth over time for the winners and losers of auctions. Empirically, restricting the control group to almost-winners (small τ_{jt} only) tends to have little effect on the point estimates, but sometimes leads to noticeably less precise estimates because the sample size becomes smaller.

Winning a procurement auction. We begin by describing the treatment associated with winning the procurement auction. Online Appendix Figure A.2 presents DiD and RDD estimates from equation (21) for two such 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 $e = 0$, so the effect is zero for $e < 0$ for both subfigures. When $e = 0$, the treated group wins a contract for the first time 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 when $e = 0$ are \$2.7 million. For $e > 0$ in subfigure a, we see that some control units win auctions, with around 15% of control units winning their first auction when $e = 1$ and around 5% when $e = 4$. This means the losers continue to bid and partially catch up to the winners in terms of the probability of having won an auction. However, as shown in subfigure b, treated units are somewhat more likely than control units to win any procurement auction on $e > 0$. Treated firms are

³³Formally, the choice of τ_{jt} plays the same role as the choice of bandwidth when using the uniform kernel implementation of RDD. In the empirical analysis, we consider a range of different choices of τ_{jt} .

around 13-21% more likely to win at least one auction at $e = 1$ and around 8-14% more likely at $e = 4$.

Taken together, the evidence in Online Appendix Figure A.2 suggests that winning the auction leads to a sharp and instantaneous increase in revenue, and that auction winning is only weakly positively correlated over time. This implies that we can reliably infer short- and long-run responses to winning a procurement contract by comparing the effect estimates across different event times.

Impacts on employment and wages across time. In Table 1, we summarize the overall impacts of winning a procurement auction on employment and earnings. To do so, we aggregate the event-specific effects estimates into a post-treatment period. This is done by averaging the event-specific estimates across event times $e \in \{0, 1, 2\}$. We compare the DiD design (no restriction on τ_{jt}) and three specifications of the RDD design (excluding firms with τ_{jt} above 0.3, 0.2, and 0.1, respectively).

In the first panel of Table 1, we find that the number of employees in the firm increases, on average, by 7-8%. The DiD and RDD estimates for earnings and employment are economically similar and never statistically distinguishable. In the second panel of Table 1, we consider the earnings for all workers in the sample. We find that, on average, winning a procurement auction leads to about a 2% increase in earnings per worker in the post-treatment period when using DiD or RDD approaches. Furthermore, both the DiD and RDD estimates show no evidence of differential trends in the wages and employment of the treated and control groups prior to the procurement auction.

The evidence suggesting that winning 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 is consistent with the notion that firms face upward-sloping labor supply curves and, therefore, have wage-setting power in the labor market. Indeed, as shown formally in the next section, we can 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 winning a procurement auction. The estimated 2% increase in earnings per worker relative to a 7-8% increase in employment is consistent with a firm-specific labor supply elasticity of 3.5-4.1.

In the above estimation, we consider all of the workers in our sample. In the third panel of Table 1, we instead present the average impacts for all stayers in the firm, defined as the sample of employees who stay in the same firm before and after the auction announcement. We find approximately the same 2% average gain in earnings when restricting the sample to stayers only. It is reassuring to find that the estimated effect on earnings barely changes when restricting the sample to stayers.

Heterogeneity across time. A natural question is how the effect estimates on earnings and employment vary over time. For example, adjustment costs may lead to smaller responses in the shorter-run than the longer-run. In Online Appendix Figure A.3, we examine this by comparing the effect estimates across event times.

In the year that the winner of the procurement auction is announced ($e = 0$), employment increases by around 6% in the winning firm (Appendix Figure A.3a). At the same time, earnings per worker increase by about 1.5% (Appendix Figure A.3b). Over time, the estimates tend to grow modestly after the year of winning, and remain relatively stable in subsequent years. By event year 2, the gain in employment is around 8% and the gain in earnings per worker is more than 2%.

Appendix Figure A.3 also reports effect estimates for each event time in the pre-period ($e < 0$), offering placebo tests for the null hypothesis that winning had no effect prior to the announcement of the winner. For none of these placebo tests are we able to reject the null hypothesis of no effect in the pre-period at the 5% level of significance.

Heterogeneity across workers. We now examine in greater detail the heterogeneity in earnings and employment impacts across workers, allowing us to assess the sensitivity of the key findings in Table 1.

In the baseline stayers estimates in Table 1, we defined stayers as workers employed by the same firm during event times $\{-2, \dots, 2\}$. In Online Appendix Figure A.4a, we vary the definition of a stayer by expanding and contracting the stayer window, always finding an increase in earnings of around 2% across definitions.

In Online Appendix Figure A.4b, we instead restrict the sample to workers who have been employed at the firm for a certain number of years prior to the auction (“tenure”). 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.

Our baseline sample of stayers consists of prime-aged W-2 employees who are FTE workers, by which we mean that their annual earnings from the primary employer are greater than the annualized full-time minimum wage in the year. In Online Appendix Figure A.4c, we investigate the sensitivity of the estimated earnings effects to changing the FTE definition. When we strengthen the FTE restriction up to 150% of the baseline definition, we still find about a 2% increase in earnings per worker.

In addition to keeping worker composition fixed by focusing on stayer or tenured samples, it is also possible to directly investigate the effect on the mean earnings of new hires after they are hired in the new firm. We do this by applying the DiD and RDD estimators of equation (21) to a panel of the mean earnings paid to new hires in a given year. As shown in Online Appendix Table A.4, we find about a 1-2% increase in the earnings of new hires in the post-treatment period for the winning firms. These estimates are comparable to the earnings effects for incumbents, and the differences across these groups are never statistically different. However, the statistical precision is limited due to the small number of new hires, so we are reluctant to make strong claims about the exact impact on new hires.

In principle, it is also possible to directly investigate composition changes for new hires by using the prior earnings of new hires at their previous firms as a proxy for worker quality.³⁴ As shown in Online Appendix Table A.5, both our DiD and RDD estimates find that worker quality does not significantly change in the post-treatment period for the winning firms. However, the standard errors are too large to rule out meaningful changes in worker composition.

5 Identification of Model Parameters

We now show how the model laid out in Section 2 can be taken to the data described in Section 4.

³⁴In theory, the wages in the previous firm reflect both worker quality and the amenities the firm offers. Thus, the earnings in the previous firm are an imperfect proxy for worker quality.

Design:	DiD	RDD		
Proximity:	Any	0.3	0.2	0.1
Log Employment				
Impact: Before Treatment	-0.015 (0.016)	-0.018 (0.017)	-0.024 (0.017)	-0.021 (0.020)
Impact: After Treatment	0.083 (0.019)	0.079 (0.020)	0.079 (0.021)	0.065 (0.025)
Unique Firms (at $e = 0$)	6,033	5,246	4,865	4,274
Log Earnings per Worker: All Workers				
Impact: Before Treatment	0.002 (0.007)	0.002 (0.007)	0.001 (0.008)	-0.002 (0.009)
Impact: After Treatment	0.020 (0.008)	0.022 (0.008)	0.020 (0.009)	0.019 (0.010)
Unique Firms (at $e = 0$)	6,033	5,246	4,865	4,274
Log Earnings per Worker: Stayers				
Impact: Before Treatment	0.003 (0.006)	0.002 (0.006)	0.001 (0.006)	0.002 (0.007)
Impact: After Treatment	0.023 (0.006)	0.023 (0.006)	0.021 (0.007)	0.019 (0.007)
Unique Firms (at $e = 0$)	5,492	4,784	4,453	3,929

Table 1: Impacts of Winning a Procurement Auction

Notes: This table presents estimates of the impacts of winning a procurement auction on employment and earnings per worker in the post-treatment and pre-treatment time periods. Event time $\bar{e} = -2$ is the omitted event time. The pre-treatment and post-treatment periods are formed by averaging the event-specific estimates across event times $e \in \{-4, -3\}$ and $e \in \{0, 1, 2\}$, respectively. Employment and earnings per worker are measured in log units. Earnings per worker are either measured for the sample of all workers in the firm or the sample of stayers, where stayers are defined as workers employed by the same firm during $e \in \{-2, \dots, 2\}$. Proximity refers to the largest value of τ_{jt} permitted in the sample. Additional specifications are provided in Online Appendix Table A.3.

5.1 Labor Supply Elasticity

Recall from equation (3) that the inverse labor supply curve is given by $w_{jt} = \theta \ell_{jt} + u_{jt}$, where we decompose the unobservable determinants of labor supply as $u_{jt} = \psi_j +$

$\xi_t + \nu_{jt}$. Our goal is to identify the labor supply elasticity, $1/\theta$.

To see why this is challenging, consider a cross-sectional regression of w_{jt} on ℓ_{jt} . This regression may result in a biased estimate of θ because both w_{jt} and ℓ_{jt} depend on time-invariant firm-specific determinants of labor supply (ψ_j). To address this issue, one might consider eliminating ψ_j by taking differences over time, denoted by Δ . Furthermore, the aggregate labor supply shocks ($\Delta\xi_t$) can be eliminated by including time fixed effects in this regression.³⁵ However, the resulting estimates of θ may still be biased due to firm-specific unobserved labor supply shocks ($\Delta\nu_{jt}$).

To address the possibility that $\Delta\nu_{jt}$ co-varies with Δw_{jt} and $\Delta\ell_{jt}$, we take advantage of the data on procurement auctions. Given this data, one possible identifying assumption is that $\Delta\nu_{jt}$ does not co-vary with the probability of winning an auction, motivating the following DiD estimand:

Proposition 3. *Under the exogeneity assumption that the auction winner D_{jt} is independent of the labor supply shock $\Delta\nu_{jt}$, and the rank condition $\text{Cov}[\Delta\ell_{jt}, D_{jt}] \neq 0$,*

$$\theta_{\text{DiD}} \equiv \frac{\text{Cov}[\Delta w_{jt}, D_{jt}]}{\text{Cov}[\Delta\ell_{jt}, D_{jt}]} = \theta.$$

Proof. By equation (4),

$$\theta_{\text{DiD}} = \underbrace{\frac{\text{Cov}[\theta\Delta\ell_{jt}, D_{jt}]}{\text{Cov}[\Delta\ell_{jt}, D_{jt}]}}_{=\theta} + \underbrace{\frac{\text{Cov}[\Delta\nu_{jt}, D_{jt}]}{\text{Cov}[\Delta\ell_{jt}, D_{jt}]}}_{=0} + \underbrace{\frac{\text{Cov}[\Delta\xi_t, D_{jt}]}{\text{Cov}[\Delta\ell_{jt}, D_{jt}]}}_{=0} = \theta,$$

where the second term on the right-hand side is zero because D_{jt} is independent of $\Delta\nu_{jt}$, and the third term is zero because $\Delta\xi_t$ is the same for all firms in each period t . \square

The identifying assumption of Proposition 3 may hold if firm-specific labor supply shocks are unpredictable at the time of bidding, and therefore do not influence the probability of winning. In our context, there is typically a time delay between the procurement solicitation for bids and the commencement of production on the government project. If the firm cannot predict its firm-specific labor supply shock from

³⁵In a specification check presented in Section 6.1, we fully interact the event time with a local market identifier, which allows the aggregate labor supply shocks to vary across firms in different local markets.

the time of the procurement solicitation to the time that it actually needs to hire workers, $\Delta\nu_{jt}$ should be independent of Z_{jt} and therefore also independent of D_{jt} .

It is useful to observe what is and is not restricted under Proposition 3. First of all, the proposition does not impose any additional restrictions on the relationships among the variables $(Z_{jt}, D_{jt}, \phi_{jt}, \psi_j, \xi_t)$. In other words, it does not restrict how bids, TFP, time-invariant firm-specific labor supply determinants, and market-wide labor supply shocks co-vary. Second, the proposition permits $\text{Var}[\Delta\nu_{jt}] > 0$, so it allows for firm-specific labor supply shocks. This is less restrictive than much of the existing literature on identifying the labor supply curve, which requires that firm-specific labor supply determinants are constant within the estimation window (see the discussion by Lamadon et al. 2022). Third, since both the control and treated firms participate in the auction, the proposition still holds if bidders choose to participate due to (expectations over) firm-specific labor supply shocks. For example, both the treated and control firms had to qualify to bid in the auctions, which may, in part, depend on the amenities they offer, as discussed in Section 4.1. However, any qualification criteria apply to the treated and control firms, so they should not generate differences in amenities between winning and losing auction bidders.

A key restriction in the assumptions underlying Proposition 3 is that, even though both treated and control firms meet the qualifications for bidding, the treated firms may bid a lower amount because they (expect to) have an increase in labor supply relative to control firms. To address this concern, we will use the RDD estimator that compares firms that win the price-only auction to firms that make similar bids and almost win the auction. If differences in bids reflect differences in $\Delta\nu_{jt}$, then we would expect firms that make similar bids to also have similar $\Delta\nu_{jt}$, which motivates the following RDD estimand:

Proposition 4. *Consider the ratio of RDD estimators defined by,*

$$\theta_{RDD}(\bar{\tau}) \equiv \frac{\mathbb{E}[\Delta w_{jt} | \tau_{jt} = 0] - \mathbb{E}[\Delta w_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]}{\mathbb{E}[\Delta \ell_{jt} | \tau_{jt} = 0] - \mathbb{E}[\Delta \ell_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]},$$

where $\bar{\tau}$ specifies the maximum proximity to the discontinuity and the conditioning on ι is implicit. Under the rank condition $\mathbb{E}[\Delta \ell_{jt} | \tau_{jt} = 0] \neq \lim_{\bar{\tau} \rightarrow 0^+} \mathbb{E}[\Delta \ell_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]$, $\lim_{\bar{\tau} \rightarrow 0^+} \theta_{RDD}(\bar{\tau})$ recovers θ .

Proof. Since bids of winners and losers converge as $\bar{\tau} \rightarrow 0^+$,

$$\lim_{\bar{\tau} \rightarrow 0^+} \theta_{RDD}(\bar{\tau}) = \theta + \lim_{\bar{\tau} \rightarrow 0^+} \frac{\mathbb{E}[\Delta\nu_{jt}|\tau_{jt} = 0] - \mathbb{E}[\Delta\nu_{jt}|0 < \tau_{jt} \leq \bar{\tau}]}{\mathbb{E}[\Delta\ell_{jt}|\tau_{jt} = 0] - \mathbb{E}[\Delta\ell_{jt}|0 < \tau_{jt} \leq \bar{\tau}]} = \theta,$$

where the second term is zero because, by independence of D_{jt} and $\Delta\nu_{jt}$ conditional on Z_{jt} , $\mathbb{E}[\Delta\nu_{jt}|D_{jt} = 1, Z_{jt}] = \mathbb{E}[\Delta\nu_{jt}|D_{jt} = 0, Z_{jt}]$, which implies $\mathbb{E}[\Delta\nu_{jt}|\tau_{jt} = 0] = \lim_{\bar{\tau} \rightarrow 0^+} \mathbb{E}[\Delta\nu_{jt}|0 < \tau_{jt} \leq \bar{\tau}]$. \square

While the RDD estimand relies on weaker assumptions than the DiD estimand, one still may be concerned about biases. One possible source of bias is simultaneity: winning the auction may have a causal impact on the labor supply the firm is facing. For example, as a result of winning the auction, firms may not only bid up wages to hire more workers but also purchase or produce better amenities, which may in and of itself attract more workers. To address this concern, Online Appendix H provides a sensitivity check where we examine how the key conclusions regarding the labor supply curve would change if amenities became a greater or smaller share of total compensation in response to winning a procurement contract.

5.2 Firm Technology

Our goal in this subsection is to identify the composite returns to labor ρ in the production function (equation 9). Given ρ , it will be straightforward to recover the other technology parameters, as discussed below. We do not observe output quantity Q_{jt} , making it difficult to estimate the production function in equation (9) directly. Our approach instead makes use of the first-order condition $x_{jt} = \kappa_X + \rho\ell_{jt} + \phi_{jt}$ from equation (11), which shows that ρ relates log intermediate expenditure x_{jt} and log labor ℓ_{jt} . The identification challenge is that log TFP, ϕ_{jt} , is an unobserved determinant of intermediate input expenditures in equation (11), and ℓ_{jt} depends directly on TFP, as shown in Online Appendix A.2. Thus, a regression of x_{jt} on ℓ_{jt} will fail to recover ρ due to the correlated unobservable ϕ_{jt} .

To address this identification challenge, we will invert the bidding strategy to control for TFP. Given the equilibrium bidding strategy $Z_{jt} = s_{u_{jt}}(\phi_{jt})$ from equation (14), where s is monotonic in ϕ_{jt} given u_{jt} , we can write the inverse equilibrium bidding strategy as $\phi_{jt} = s_{u_{jt}}^{-1}(Z_{jt})$. Monotonicity ensures that $s_{u_{jt}}^{-1}$ is unique and ϕ_{jt}

is pinned down by the bids Z_{jt} , conditional on the amenities u_{jt} . Given any consistent estimator $\widehat{\theta}$ of θ , an estimator for u_{jt} is

$$\widehat{u}_{jt} = w_{jt} - \widehat{\theta}l_{jt}.$$

We then have the following result:

Proposition 5. *Consider a regression of x_{jt} on l_{jt} controlling for $(\widehat{u}_{jt}, Z_{jt})$, i.e.,*

$$\widehat{\rho} \equiv \frac{\text{Cov}[x_{jt}, l_{jt} | \widehat{u}_{jt}, Z_{jt}]}{\text{Var}[l_{jt} | \widehat{u}_{jt}, Z_{jt}]}.$$
 (22)

Given that $\widehat{\theta}$ recovers θ and the rank condition $\text{Var}[l_{jt} | \widehat{u}_{jt}, Z_{jt}] > 0$, $\widehat{\rho}$ recovers ρ .

Proof. By equation (11),

$$\widehat{\rho} = \underbrace{\frac{\text{Cov}[\rho l_{jt}, l_{jt} | \widehat{u}_{jt}, Z_{jt}]}{\text{Var}[l_{jt} | \widehat{u}_{jt}, Z_{jt}]}}_{=\rho} + \underbrace{\frac{\text{Cov}[\phi_{jt}, l_{jt} | \widehat{u}_{jt}, Z_{jt}]}{\text{Var}[l_{jt} | \widehat{u}_{jt}, Z_{jt}]}}_{=0} = \rho,$$

where the second term is zero because $\text{Cov}[\phi_{jt}, l_{jt} | \widehat{u}_{jt}, Z_{jt}] = \text{Cov}[\phi_{jt}, l_{jt} | u_{jt}, \phi_{jt}]$ by the uniqueness of $\phi_{jt} = s_{u_{jt}}^{-1}(Z_{jt})$ and that u_{jt} recovers \widehat{u}_{jt} since $\widehat{\theta}$ recovers θ . To see why the rank condition is expected to hold in our context (and, indeed, we find that it holds in the data), note that if $\epsilon > 0$, then l_{jt} depends directly on D_{jt} in the firm's problem, all else equal. From equation (14), D_{jt} depends not only on u_{jt}, ϕ_{jt} , but also on the realized competitors' bids in the auction. Competitors' bids are unknown to the firm at the time it bids and thus not captured by the optimal bidding function. \square

In practice, we implement this control function approach by controlling for auction fixed effects as well as third-order polynomials in $\log Z_{jt}$ and \widehat{u}_{jt} . The results do not change materially if we increase the polynomial order.

5.3 Product Demand Curve

Our goal in this subsection is to identify the elasticity of the product demand curve, $-1/\epsilon$, introduced in equation (6). To do so, we rely on the first-order conditions for labor (equation 13) and for intermediate inputs (equation 12) among the firms in the

construction industry that do not currently receive procurement contracts ($D_{jt} = 0$ firms).

Rearranging equation (13) and conditioning on $D_{jt} = 0$ firms, the first-order condition for labor implies,

$$\overbrace{(1 - \epsilon)}^{\text{inverse markup}} = \overbrace{(1 + \theta)}^{\text{inverse markdown}} \frac{s_L}{\beta_L} + s_M, \quad (23)$$

which allows us to express the inverse price markup in terms of the inverse wage markdown, the labor share of revenue ($s_L \equiv B_{jt}/R_{jt}$), and the intermediates share of revenue ($s_M \equiv X_{jt}/R_{jt}$), and the output elasticity of labor (β_L).

Rearranging equation (12) and conditioning on $D_{jt} = 0$ firms, the first-order condition for intermediate inputs implies,

$$\overbrace{(1 - \epsilon)}^{\text{inverse markup}} = \frac{\text{Cov}[r_{jt}, x_{jt} | D_{jt} = 0]}{\text{Var}[x_{jt} | D_{jt} = 0]}, \quad (24)$$

Equations (23) and (24) provide two equations in two unknowns, ϵ and β_L , allowing us to simultaneously recover ϵ and β_L .

The above identification argument for ϵ follows the logic of the markup estimator of [de Loecker et al. \(2020\)](#) for the case where the production function is Leontief, except our expression accounts for imperfect competition in the labor market ($\theta > 0$). This shows that the markup estimator of [de Loecker et al. \(2020\)](#) can be easily corrected for labor market power in the Leontief case using a plug-in estimate of θ .

5.4 Over-identifying Restriction

Above, we demonstrated how $(\rho, \beta_L, \epsilon)$ can be exactly identified. We now show how an additional moment can be used to over-identify $(\rho, \beta_L, \epsilon)$ and, thus, directly examine the validity of the model. Among firms that receive procurement contracts, the firm's first-order condition in equation (13) implies

$$\Lambda_{jt} = \kappa_\Lambda + \rho \ell_{jt} + \phi_{jt} + e_{jt} \quad \text{if } D_{jt} = 1. \quad (25)$$

where $\Lambda_{jt} \equiv \frac{\epsilon}{1-\epsilon} r_{jt}^H + \log\left(\frac{1+\theta}{\beta_L} B_{jt} + X_{jt}\right)$ and $\kappa_\Lambda \equiv \log(1-\epsilon) + \frac{\log \rho H}{1-\epsilon}$. The derivation is provided in Online Appendix A.4. Given θ , for any candidate values of (β_L, ϵ) , we can construct the left-hand side variable. Furthermore, note that for any candidate

value of ρ , we can rearrange equation (11) to recover log TFP as

$$\widehat{\phi}_{jt}(\rho) \equiv x_{jt} - \kappa_X - \rho \ell_{jt}.$$

We can then construct the covariance between ℓ_{jt} and the left-hand side of equation (25), which is a moment equation that depends only on the unknown parameters $(\rho, \beta_L, \epsilon)$, the TFP estimates $\widehat{\phi}_{jt}(\rho)$, and the data. Thus, in addition to equations (22), (23), and (24), equation (25) gives us a fourth equation that must be satisfied by the true values of $(\rho, \beta_L, \epsilon)$. In practice, we estimate $(\rho, \beta_L, \epsilon)$ simultaneously in equations (22), (23), (24), and (25) using GMM, and then check how the estimates change if we do not use (25).

5.5 Identification of Remaining Parameters

We now show how to recover β_K , β_M/p_M , and p_H given that we already know θ , ρ , β_L , and ϵ . From the definition of ρ , we recover β_K as $\beta_K = (\rho - \beta_L) / (1 + \theta)$. From equation (11), we recover β_M/p_M as

$$\log(\beta_M/p_M) = \mathbb{E}[\rho \ell_{jt} - x_{jt}], \quad (26)$$

where we normalize $\mathbb{E}[\phi_{jt}] = 0$ without loss of generality. Rearranging equation (12),

$$\log p_H = \mathbb{E}[r_{jt} - (1-\epsilon)x_{jt} | D_{jt} = 0] - (1-\epsilon) \log(\beta_M/p_M), \quad (27)$$

where we normalize $\mathbb{E}[e_{jt} | D_{jt} = 0] = 0$ without loss of generality.

5.6 Rents and Incidence

We now show how to recover the total rents, baseline rents, and the incidence of procurement on firms and workers, focusing on the sample of firms with procurement contracts ($D_{jt} = 1$). From the expressions in Section 2.5 and given that we identified θ above, we can characterize rents and incidence for workers if we recover B_{0jt} and B_{1jt} , as well as rents and incidence for firms if we recover π_{0jt} and π_{1jt} , for each (j, t) . For firms with $D_{jt} = 1$, we observe B_{1jt} and π_{1jt} , so the only remaining challenge is to recover B_{0jt} and π_{0jt} .

To recover B_{0jt} and π_{0jt} , we use the profit-maximizing first-order condition with

respect to labor for the case in which the firm loses the procurement contract ($d = 0$), which is

$$\rho p_H (1-\epsilon) \Phi_{jt}^{1-\epsilon} L_{0jt}^{\rho(1-\epsilon)-1} = (1 + \theta) \kappa_U U_{jt} L_{0jt}^\theta + \rho \frac{p_M}{\beta_M} \Phi_{jt} L_{0jt}^{\rho-1}, \quad (28)$$

where $\kappa_U \equiv \frac{\beta_K}{\beta_L} (1 + \theta) + 1$. The derivation of equation (28) is provided in Online Appendix A.2. We identify $(\theta, \rho, \epsilon, p_H, \beta_M/p_M)$ above, and recover U_{jt} using \hat{u}_{jt} and Φ_{jt} using $\hat{\phi}_{jt}(\rho)$. Thus, we can numerically solve equation (28) to obtain L_{0jt} for each (j, t) . Given L_{0jt} , it is straightforward to recover B_{0jt} and π_{0jt} using the firm's constraints (equations 2, 8, 9, 10).

6 Estimates of Model Parameters

In this section, we combine the identification arguments in Section 5 with the data described in Section 4 to estimate the parameters that govern labor supply, firm technology, and product demand.

6.1 Labor Supply Elasticity and Wage Markdown

We now implement the estimators θ_{DiD} and θ_{RDD} described in Section 5.1. For each of these estimators, the numerator (or denominator) is given by the estimated impact of winning a procurement auction on log earnings per worker (or log number of employees). These estimated impacts are reported in the first and fourth columns of Table 1 for the DiD and RDD estimators, respectively. In Figure 3, we present the labor supply elasticity, $1/\theta$, and the wage markdown relative to the MRPL, $(1 + \theta)^{-1}$, implied by these estimated impacts.

Figure 3 shows that the estimates of the labor supply elasticity range from 3.5 to 4.1 depending on whether we use the RDD or DiD approach. The labor supply elasticity estimates indicate that, if a US construction firm aims to increase the number of its employees by 10%, it needs to increase wages by 2.4-2.9%. Our estimates are broadly comparable to existing work. Lamadon et al. (2022) estimate a labor supply elasticity of 4.6 and Suárez Serrato and Zidar (2016) estimate a labor supply elasticity of 4.2, while Card et al. (2018) pick 4.0 as the preferred value in their calibration exercise. A related literature using experimentally-manipulated wage offers for small tasks or survey experiments typically finds labor supply elasticities ranging from 3.0

to 5.0 (Caldwell and Oehlsen, 2018; Dube et al., 2020; Sokolova and Sorensen, 2021).


Specification checks. We now show that our conclusions about the labor supply curve are robust to a number of specification checks.

Our model assumes that the labor supply elasticity is the same for all construction firms. One potential concern is that firms that bid in auctions differ from the rest of the construction industry. In order to investigate this possibility, we consider the estimator of Lamadon et al. (2022, LMS), which uses the lagged value-added change as the instrument in the DiD estimator in equation (21). The LMS estimator can be justified either if firm-specific unobserved labor supply determinants are fixed over the estimation window, or if TFP shocks are more persistent than firm-specific labor supply shocks.³⁶ Importantly, value-added shocks are defined for all firms in the construction industry, not only for firms that participate in auctions, so they can be used to assess the representativeness of the subsample of firms that participate in auctions. The bars labeled “LMS Design” and “LMS Design, Control CZ” in Online Appendix Figure A.5 show that, when applying the LMS estimator to the entire construction industry (with or without controlling for local labor market shocks, as proxied by commuting zone), the labor supply elasticity and markdown estimates are nearly the same as our RDD and DiD estimates.

Another potential concern is that, 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 labor supply elasticity. Our estimates in Figure 3 report averages across three event years. Online Appendix Figure A.5 provide estimates separately across event years. We find approximately the same labor supply elasticity estimate in each event time, suggesting adjustment costs are relatively unimportant in our setting.

Another challenge when estimating the labor supply curve is skill-upgrading. Earnings per worker could increase for the winning firm in part because they hire more productive workers. The results presented in Figure 3 suggest that skill-upgrading is empirically unimportant for our estimates of the labor supply elasticity: we find little difference between estimates using stayers versus all workers in the firm. In Online

³⁶Proposition 6 of Online Appendix F.1 provides the formal conditions for identification using the LMS estimator.



(a) Labor Supply Elasticity, $1/\theta$

(b) Wage Markdown relative to MRPL, $(1 + \theta)^{-1}$

Figure 3: Main Estimates of the Labor Supply Elasticity and Wage Markdown

Notes: This figure presents the main estimates of the labor supply elasticity, $1/\theta$, and wage markdown relative to MRPL, $(1 + \theta)^{-1}$. It uses the DiD and RDD estimands defined in Propositions 3 and 4, respectively. In the estimation, earnings per worker are either measured for the sample of all workers in the firm or the sample of stayers. The largest proximity permitted in the RDD estimation is 0.1. Specification details, sample definitions, and sensitivity checks are discussed in the text.

Appendix Figure A.6a, we show that this conclusion is insensitive to the number of years that the worker is a stayer.

A final specification check is motivated by the potential concern that procurement contracts might be awarded in local labor markets with worse labor supply shocks, which could induce a correlation between local aggregate labor supply shocks and procurement contracts.³⁷ In a specification check presented in Figure A.5, we fully interact the event time with a local market identifier, which allows the aggregate labor supply shocks (ξ_t) to vary across firms in different markets. We define the set of firms in the same commuting zone or that bid against each other in the same auction as belonging to the same market.³⁸ We find similar labor supply elasticity estimates in either case, indicating local labor market shocks are not an important confounder in

³⁷Note that the estimates based on Proposition 4 are not subject to this concern, as local aggregate labor supply shocks are expected to be the same on average between winners and almost winners of auctions.

³⁸Online Appendix F.3 provides implementation details for the specifications that are interacted with the commuting zone or auction identifier.

our estimation.

Regulations on amenities and pay in the construction industry. There are a number of regulations governing the construction industry that may raise concerns for our research design. We now consider these concerns in turn.

One potential concern is that some states have prevailing wage laws that could force firms that win a state government procurement auction to increase the wages of incumbent workers, independently of whether or not they hire new workers. State prevailing wage laws require that construction workers employed by private construction firms on state-funded construction projects be paid at least the wages and benefits paid to similar workers in the same location where the project is located. We investigate if these laws affect our estimates in two ways.

First, we examine whether the estimates are substantially different among states that have prevailing wage laws for state contracts. In Online Appendix Table A.6, we find for this sample very similar estimated impacts of winning a procurement contract on employment as well as on earnings per worker, both for stayers and for all workers in the firm. As a result, the implied labor supply elasticity and markdown estimates are not materially different for states with prevailing wage laws than for the population of firms at large.

Second, we use repeals of state prevailing wage laws in a difference-in-differences analysis at the state-level to examine how wage and non-wage compensation are impacted by prevailing wage laws. For outcome measures, we use state-level survey data on the components of worker compensation from the Economic Census of the Construction Sector, which is available at five-year frequency from 1977-2017. These components include both wages and non-wage fringe benefits.³⁹ During this time interval, 15 states repealed their prevailing wage laws. In Online Appendix E.2, the difference-in-differences estimates for repeals suggests that prevailing wage laws have little effect on total compensation, wages, non-wage fringe benefits, or the share of total compensation from non-wage fringe benefits.

Other than prevailing wage laws, another potential concern is that winning a

³⁹The Economic Census of the Construction Sector data are described in Online Appendix E.1. Voluntary fringe benefits – which is the component of non-wage compensation that the firm can in principle adjust – account for only one-tenth of total compensation in the construction industry.

procurement contract improves workplace safety, which is arguably a key non-wage attribute of jobs in the construction industry. For two reasons, we do not believe this is a particularly relevant concern in practice. First of all, a review of workplace safety standards from the key governing body, the Occupational Safety and Health Administration (OSHA), indicates that construction workplaces are highly-regulated for worker safety, with numerous licenses and regulations related to the type of tasks being performed. However, as discussed in Online Appendix E.3, these regulations apply to all firms in the construction industry, not specifically to firms that win a procurement auction.

Consistent with the fact that safety regulations are not specific to firms that win procurement contracts, we find no effect of winning a procurement auction on measures of violations of safety regulations. To estimate these effects, we run the same regression specification as in equation (21), but now using safety violations and safety investigations as the outcome variable, as measured in publicly available data from OSHA. As shown in Online Appendix Table A.7, we find fairly precisely estimated zero effects of winning a procurement auction on the probability of a safety violation and on the probability of a safety investigation.

While the above specification checks find that the amenities that are measured in these datasets do not change materially in response to winning an auction, we admittedly cannot rule out changes to amenities that are not observed in our data. Therefore, Online Appendix H provides a sensitivity check where we examine how the key conclusions would change if winning a procurement auction actually did have a causal effect on amenities. We show that, if amenities increase by the same percent as earnings, then amenity changes do not introduce a bias in the estimate of θ . This is because log compensation and log earnings increase by the same amount if amenities increase by the same percent as earnings, so log earnings changes are a valid proxy for log compensation changes. However, log compensation could in reality increase more (or less) than log earnings, so the estimate of θ using earnings may be downward-biased (or upward-biased). In Online Appendix Figure A.11, we examine how differential changes in earnings and total compensation may bias our estimates of the labor supply elasticity. We find that, even in the rather extreme case in which log compensation increases by 20% more (or less) than log earnings due to amenity

responses, the labor supply curve is upward sloping with an elasticity that would be close to our baseline estimates in Figure 3.

Sensitivity to labor hours and full-time status. A potential concern in our data is that we only observe earnings changes, not changes in hourly wages. If part of the earnings change is due to an increase in hours, then we overstate the wage increase relative to the labor increase, leading to a downward-biased estimate of the labor supply elasticity, $1/\theta$. We now investigate this possibility.

One natural way to increase hours is to promote part-time workers to full-time status. However, part-time employment is rare in the construction industry: in 2015, 13.9% of all US private sector workers were part-time but only 4.6% of construction industry workers were part-time (BLS, 2021). To better understand the relationship between full-time employment status and annual earnings, we analyze CPS ASEC cross-sectional random samples of the US labor force, which include measures of hours worked and annual earnings. In the 2001-2015 ASEC samples, 6.4% of construction workers are part-time employed, compared to 11.9% across all industries. When imposing in the ASEC samples our FTE restriction that annual earnings from the primary employer are greater than the annualized full-time minimum wage, only 4.4% of the remaining construction sample is part-time employed, compared to 7.1% across all industries. Thus, the FTE restriction removes about 30% of part-time employees from the ASEC samples.

Nevertheless, one may worry that promotions from part-time to full-time employment in response to winning a procurement contract explain our estimated increase in earnings per worker. To investigate this, we consider again the stayers sample described above. Since stayers are defined as workers who were already FTE before the procurement contract was received and remained FTE after, stayers could not have been promoted from non-FTE to FTE status in response to winning the procurement contract. We find almost the same estimates for the stayers sample as we do for the full sample of workers, suggesting that promotion from part-time to full-time status does not drive our results. However, it could be that our baseline FTE sample includes some part-time workers with relatively high hourly wages. In the bars labeled “125% FTE” and “150% FTE” in Online Appendix Figure A.5, we strengthen the FTE restriction up to 125% and 150% of the baseline definition, respectively,

with additional choices provided in Online Appendix Figure A.6e. Transitions from part-time to full-time status should become less likely as the FTE restriction rises. We find that the estimate is insensitive to raising the FTE restriction, suggesting we have successfully ruled out part-time to full-time promotions with the baseline FTE restriction. It is not surprising that part-time employment does not confound our estimates, given how rare part-time employment is in the construction industry.

The other natural way to increase hours is over-time pay for incumbent full-time workers. When thinking about the plausibility that our estimates are driven by over-time pay, it is useful to observe that the effects of receiving a procurement contract persist over several years. In Figure A.3, we find that the change in earnings due to receiving a procurement contract (which is the numerator of θ_{DiD} and θ_{RDD}) is positive, statistically significant, and relatively stable over the four years after the firm wins the auction, whereas the typical procurement project lasts for less than one year. Thus, it is unlikely that over-time pay to meet a short-lived increase in product demand explains our estimated increase in earnings.

While the evidence from the US data indicate that the increase in earnings is not due to increased hours worked, the most compelling evidence would come from directly estimating annual earnings versus hourly wage responses in data with administrative measures of each worker’s labor hours. Labor hours data is not available from the IRS, nor is it available in other nationally representative employer-employee panel data from the US (e.g., LEHD).⁴⁰ To overcome this challenge, we consider data from Norway. Norway provides a rare opportunity, as it is one of the few countries where the hours worked by each employee are reported to the government. We restrict the Norwegian sample to the construction industry and workers who satisfy the same FTE restriction as we impose in the US data.⁴¹

To determine whether or not labor hours responses confound wage responses to

⁴⁰Only Minnesota, Oregon, Rhode Island, and Washington State collect labor hours data as part of their unemployment insurance records, representing a small fraction of the US workforce and a small fraction of states covered by our procurement auction records. This hourly wage data is not available through the IRS. For other papers using IRS data (without information on hours) to study wage-setting and labor markets, see e.g., [Kline et al. \(2019\)](#), [Song et al. \(2019\)](#), [Yagan \(2019\)](#), and [Güvenen et al. \(2021\)](#).

⁴¹Online Data Supplement S.4 provides details on the Norwegian data sources and sample construction.

firm shocks in Norway, we apply the LMS estimator (discussed above) to recover the pass-through from revenue shocks to annual earnings and hourly wages. We find that the elasticity of annual earnings and hourly wages to revenue shocks are 0.092 and 0.091, respectively, while the elasticity of hours to revenue shocks is 0.001. Thus, the labor supply elasticity is nearly identical when estimated using annual earnings versus hourly wages.

6.2 Firm Technology and Product Demand Parameters

We use GMM to jointly estimate $(\rho, \beta_L, \epsilon)$ based on equations (22), (23), (24), and (25). The estimates and standard errors are reported in Panel B of Table 2.

We estimate β_L to be 0.50 and ρ to be 1.09. The value of β_L implies that a 100% increase in a firm's employment results in 50% more output, all else equal. The value of ρ implies that, if a firm has 100% more labor than another firm, we expect it to produce 109% more output, not holding all else equal. The larger firm will optimally have greater utilization of capital and intermediate inputs. Since $\rho \equiv (1 + \theta)\beta_K + \beta_L$, these estimates imply that β_K is about 0.47 (see Panel C in Table 2). The returns to scale over labor and capital, $\beta_L + \beta_K$, is about 1.0, which is comparable to the range of estimates from 1.0 to 1.2 by [Levinsohn and Petrin \(2003\)](#).

As discussed in Section 2.3, the government project crowds-out private projects ($Q_{1jt}^H < Q_{0jt}^H$) since we estimate that $1 + \theta > \rho$. To see why this is the case, note that winning a government project increases the total output level. In turn, more employment is required 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 greater private production under composite economies of scale, $\rho > 1$. Since we estimate $1 + \theta > \rho$, it is optimal for a firm to cut its production for the private market if it receives a procurement contract. We quantify the crowd-out effect when discussing incidence in the next section.

In the private product market, we estimate that ϵ is 0.14, so the product demand elasticity $-1/\epsilon$ is about -7.3 . This implies that, in order for a firm to increase output by 10%, it must reduce its price by about 1.4%. The price markup $(1 - \epsilon)^{-1}$ is about 1.16, which implies that the price is 16% more than the relevant marginal cost reflecting both labor and intermediate inputs. Online Appendix Figure A.8

Panel A. Worker Preference Dispersion and Wage Markdown		
	Parameter and Identifying Moments	Data
Inverse Labor Supply Elasticity θ :		
DiD Estimand (Prop 3)	θ_{DiD}	0.245 (0.086)
RDD Estimand (Prop 4)	θ_{RDD}	0.286 (0.092)
Wage Markdown $1/(1 + \theta)$		
DiD Estimand (Prop 3)	$1/(1 + \theta_{\text{DiD}})$	0.803 (0.055)
RDD Estimand (Prop 4)	$1/(1 + \theta_{\text{RDD}})$	0.777 (0.059)
Panel B. Technology and Product Demand Parameters		
Baseline Estimates using Over-identified GMM		
	Parameters	Data
Composite returns to labor (eqs 22-25)	ρ	1.089 (0.017)
Marginal returns to labor (eqs 22-25)	β_L	0.499 (0.192)
Inverse product demand elasticity (eqs 22-25)	ϵ	0.137 (0.015)
Alternative Estimates using Exactly-identified System		
	Parameters	Data
Control function estimation using bids (eq 22)	ρ	1.057 (0.015)
Markup and labor share relationship (eq 23)	β_L	0.514 (0.209)
Intermediate inputs to revenues ratio (eq 24)	ϵ	0.137 (0.008)
Panel C. Remaining Parameters for Price, Scale, and TFP		
	Parameter and Identifying Moments	Data
Scale of log output price (eq 27)	$\log p_H = \mathbb{E}[r_{jt} - (1 - \epsilon)(\log \frac{\beta_M}{p_M} + x_{jt}) D = 0]$	12.801 (0.053)
Scale of log amenities (\hat{u}_{jt})	$\mathbb{E}[u_{jt}] = \mathbb{E}[b_{jt}] - (1 + \theta)\mathbb{E}[\ell_{jt}]$	10.075 (0.000)
Scale term for intermediates (eq 26)	$\log \frac{\beta_M}{p_M} = \rho\mathbb{E}[\ell_{jt}] - \mathbb{E}[x_{jt}]$	-11.722 (0.047)
Marginal returns to capital	$\beta_K = (\rho - \beta_L)/(1 + \theta)$	0.474 (0.161)
Interquartile range of log TFP ($\hat{\phi}_{jt}$)	$\text{IQR}(\phi_{jt}) = \text{IQR}(x_{jt} - \rho\ell_{jt})$	0.918 (0.001)

Table 2: Estimates of the Structural Parameters

Notes: This table summarizes identifying equations and provides estimates of several model parameters. DiD and RDD estimates of the inverse labor supply elasticity, θ , and wage markdown relative to MRPL, $(1 + \theta)^{-1}$, are provided in Panel A. Over-identified and exactly-identified estimates of $(\rho, \beta_L, \epsilon)$ are provided in Panel B, where we equally weight the moments in over-identified GMM. Estimates of the remaining parameters are provided in Panel C. Specification details, sample definitions, and sensitivity checks are discussed in the text.

estimates heterogeneity in $1 - \epsilon$ across Census regions, finding little variation. 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 \(1999\)](#) estimate demand elasticities for German beer to be -2.3 to -15.4 . [Goldberg and Verboven \(2001\)](#) estimate demand elasticities for foreign cars to be -4.5 to -6.5 .

Sensitivity analyses. We now apply several sensitivity checks to our GMM estimates of $(\rho, \beta_L, \epsilon)$ to verify that our results are not driven by model assumptions. In Panel B of Table 2, we provided baseline estimates from over-identified GMM. In order to directly examine the validity of the model, we can instead estimate the parameters using direct estimation of the exactly-identified system, dropping the over-identifying restriction discussed in Section 5.4.⁴² Using this approach, we estimate that ρ is 1.06, β_L is 0.51, and ϵ is 0.14, which are nearly the same as the baseline estimates.

Our baseline analyses have assumed that the price parameters (p_H, p_K, p_M) do not vary over time. While time-varying price parameters would lead to the same model equations in Section 2, the regression intercepts in equations (22), (24), and (25) would have year subscripts, which suggests controlling for year fixed effects in these regressions. When doing so, we find that the GMM estimates of $(\rho, \beta_L, \epsilon)$ remain identical, so accounting for time-variation in the price parameters does not affect our results.

The Leontief production function is motivated by institutional features of the construction industry and allows us to derive a linear relationship between x_{jt} and ℓ_{jt} (equation 11), but one may worry that it is misspecified. Although misspecification in the production function would not affect the estimated labor supply elasticity or total rents, it could affect the estimates of ρ and ϵ as well as analyses of the incidence of procurement. In Online Appendix C, we solve, identify, and estimate the model with a Cobb-Douglas production function. In this case, we estimate that ρ is 1.08 and $-1/\epsilon$ is -5.04 , which are similar to the estimates based on the Leontief production function. We also find similar estimates of the incidence of procurement on firms and workers. Thus, the functional form of the production function does not drive our results.

In Section 2.3, we discussed the assumption that auctions are symmetric, which leads to a closed-form expression for the optimal bidding strategy. While the identification strategies for θ , β_L , and ϵ do not rely on auction symmetry, we used symmetry to recover ρ in Proposition 5. If auctions were not symmetric, inverting the bid-

⁴²In particular, equation (22) provides a direct estimate of ρ , equation (24) provides a direct estimate of ϵ , and equation (23) provides a plug-in estimate of β_L given the direct estimates of ϵ and ρ .

ding function to control for TFP would require not only controlling for a firm’s own amenity term and bid but also the amenities of all other bidders in the auction. As a sensitivity check, we re-estimate ρ when controlling for a polynomial in the average amenities of all other firms in the same auction, finding a nearly identical estimate of ρ . Thus, the assumption that auctions are symmetric does not drive estimates of key model parameters.

Remaining model parameters. For the few remaining model parameters, the identifying equations and estimates are provided in Panel C of Table 2. These include the private market price index parameter p_H (using equation 27), the scale of amenities $\mathbb{E}[u_{jt}]$ (using \hat{u}_{jt}), the returns to intermediate inputs relative to marginal cost β_M/p_M (using equation 26), the returns to capital β_K (using the definition of ρ), and the interquartile range of TFP (using $\hat{\phi}_{jt}(\rho)$).⁴³ Although the magnitudes of these parameters are perhaps not of interest on their own, they are needed to perform model simulations.⁴⁴

7 Quantifying the Importance of Double Market Power

We now use the estimated model to quantify double market power and its implications for the outcomes and behavior of workers and firms in the US construction industry.

7.1 The Double Wage Markdown and Double Price Markup

We use the estimates in Table 2 to quantify the double markdown of wages and double markup of prices. The results are presented in Table 3. We consider both RDD and DiD estimates of θ . Drawing on equation (19), the double markdown of the wage is about 0.69 when using the DiD estimator for θ and about 0.67 when using the RDD estimator for θ . Thus, the wage is 31-33% below the value of MPL. Drawing on

⁴³The estimates of parameters β_K and $\mathbb{E}[u_{jt}]$ depend on the estimate of θ . In Panel C of Table 2, we present the estimates using θ_{DiD} . When instead using θ_{RDD} , the estimate of β_K is 0.46 instead of 0.47 and $\mathbb{E}[u_{jt}]$ is 9.96 instead of 10.08, which are very similar.

⁴⁴For example, the simulation exercises below use the TFP interquartile range when simulating optimal bids as a function of TFP dispersion from equation (14). One potential concern is that the distribution of TFP varies over time, so our interquartile range estimate may be overstated due to pooling across years. Online Appendix Figure A.7 estimates the TFP interquartile range separately by calendar year, finding little evidence of changes over time.

Panel A.	Components of the Double Markdown of the Wage		
	Markdown $(1 + \theta)^{-1}$	Inverse Markup $(1 - \epsilon)$	Double Markdown $(1 + \theta)^{-1}(1 - \epsilon)$
Using θ_{DiD} :	0.803		0.693
Using θ_{RDD} :	0.777	0.863	0.671
Panel B.	Components of the Double Markup of the Price		
	Markup $(1 - \epsilon)^{-1}$	Inverse Markdown $(1 + \theta)$	Double Markup $(1 - \epsilon)^{-1}(1 + \theta)$
Using θ_{DiD} :		1.245	1.443
Using θ_{RDD} :	1.159	1.286	1.491

Table 3: Estimates of the Double Markdown and Double Markup

Notes: This table summarizes our estimates of double market power. The first column of Panel A provides our estimate of the wage markdown relative to MRPL, $(1 + \theta)^{-1}$, using either the DiD or RDD estimand, while the first column of Panel B provides our estimate of the price markup relative to marginal cost, $(1 - \epsilon)^{-1}$. The second column provides the inverse of these markup and markdown terms. The third column uses Proposition 2 to define the double markdown in Panel A and the double markup in Panel B.

equation (20), the double markup of the price is 1.44 when using the DiD estimator for θ and 1.49 when using the RDD estimator for θ . Thus, the price is 44-49% above the (productivity-adjusted) wage.

To gauge the empirical relevance of the double markdown, it is useful to compare it to the markdown estimate that one would obtain if ignoring product market power. If one assumed a perfectly competitive product market, then the only markdown would be $(1 + \theta)^{-1}$ and one would conclude that the wage is only 20-22% below the value of MPL. Conversely, if one assumed a perfectly competitive labor market, then the only markup would be $(1 - \epsilon)^{-1}$ and one would conclude that the price is only 16% above the (productivity-adjusted) wage.

7.2 Estimated Impacts of Changes in Market Power

We quantify the impacts of a change in market power in a given market, all other things being equal, using the approach laid out in Section 3.2 and the model estimates from Table 2. We focus on a typical firm, by which we mean the median-TFP firm.

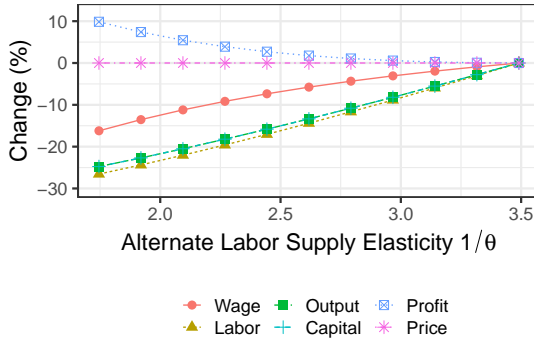
In Figure 4a, we predict the impacts of increased labor market power if the firm did not have product market power. This is done by increasing the inverse labor supply elasticity from our estimated value to a higher value, while making sure that labor supply at the initial wage is the same (by changing the location parameter in labor supply U). For each outcome, we find a monotonic relationship across values of the labor supply elasticity. When the labor supply elasticity of a given firm is reduced by half, the firm employs 27% fewer workers and decreases wages by 16%. Expenditure on inputs declines by 39% for labor, 25% for capital, and 25% for intermediate materials.⁴⁵ Output is reduced by 25% and prices are constant (since $\epsilon = 0$), so revenue declines by 25%. Despite these reductions in the firm's activities, profit rises 10% due to the increased markdown of wages.

In Figure 4b, we predict the impacts of increased labor market power given our estimate of product market power, thus providing our best prediction of the impacts that increased labor market power would have on the outcomes of the US construction industry. As before, outcomes are monotonic across values of the labor supply elasticity. However, the impacts of the increase in labor market power are substantially attenuated by taking into account that the firm has product market power and, thus, it was profitable to initially set the price above the marginal cost.


When the labor supply elasticity of a given firm is reduced by half, the firm is less willing to cut wages by lowering employment. Due to its product market power, the firm only reduces employment by 15% and only cuts wages by 9%. Expenditures decrease by 23% for labor, 6% for capital, and 11% for intermediate materials. Output decreases by 11%, prices rise by 2%, revenue decreases by 9%, and profits increase by 1%.

Figures 4c-4d perform the same exercises as in Figures 4a-4b, except we now increase product market power. This is done by increasing the inverse product demand elasticity from our estimated value to a higher value, while making sure that product demand at the initial price is the same (by changing the location parameter p_H). In Figure 4c, we assume the firm does not have labor market power. When the product demand elasticity of a given firm is reduced by half, the firm will then employ 47%


⁴⁵Intermediate material expenditure is not shown in the figure because it must change in proportion to output due to the Leontief production function.



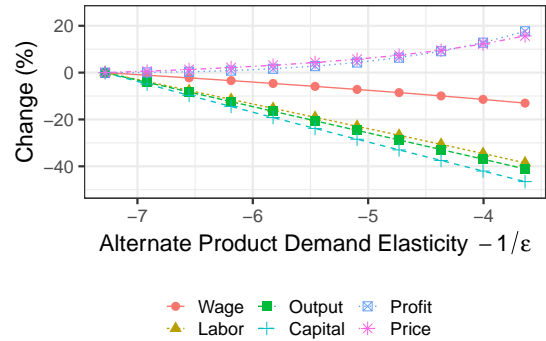
(a) Impacts of Labor Market Power without Product Market Power ($\epsilon = 0$)



(b) Impacts of Labor Market Power with Product Market Power (True ϵ)



(c) Impacts of Product Market Power without Labor Market Power ($\theta = 0$)



(d) Impacts of Product Market Power with Labor Market Power (True θ)

Figure 4: Estimated Impacts of Changes in Market Power

Notes: This figure presents expected impacts of increased labor or product market power, using the approach defined in Section 3.2. To do so, it simulates from the model defined in Section 2 for the typical firm, evaluated at the parameter estimates provided in Table 2. In subfigures (a,b), we increase labor market power through compensated changes in θ , while in subfigures (c,d), we increase product market power through compensated increases in ϵ . In subfigures (a,c), we assume there is only market power in the indicated market, while in subfigures (b,d), we assume market power in both labor and product markets is characterized by our preferred estimates. Each simulated outcome is expressed as a percent change relative to the value observed in the data.

fewer workers while wages are unchanged (since $\theta = 0$). By comparison, the use of labor only declines by 39% and wages only fall by 13% when we, in Figure 4d, take into account that the firm has labor market power and, thus, was profitably setting wages below MRPL. Incorporating labor market power also leads to a modest atten-

uation of the predicted impact of increased product market power on output, prices, intermediate materials, and profits.

8 Rents and Incidence of Government Procurements

In this section, we use the estimates of model parameters to infer the rents, rent-sharing, and incidence of procurements in the US construction industry. We here focus on firms that currently have a procurement contract so that we can characterize the incidence of procurements for the firms that are actually impacted.

Table 4 provides our main results on rents and incidence. The first column of Table 4 presents the actual firm and worker outcomes. We focus on firms that receive a procurement contract ($D_{jt} = 1$) and provide estimates for the typical firm, by which we mean the median-TFP firm. The typical firm employs about 25 workers and pays them an annual wage of \$59,100. This amounts to an annual wage bill of about \$1.5 million. Using the expression for incidence on worker rents in Section 2.5, this implies that worker rents are about \$11,600 per worker, which amounts to 20% of their average earnings. Comparing revenues to expenditures on all inputs, firm rents (i.e. profits) amount to about \$43,100 per worker. Comparing worker rents to firm rents, 79% of total rents are captured by firms.

We now investigate how these results would change if the median-TFP firm instead had above-median or below-median TFP. To do so, we assign alternative TFP quantiles to this firm without changing any other primitives of the model, then resolve the model to obtain this firm’s alternative outcomes and rents.⁴⁶ The results are provided in Figure 5. The x-axis displays the alternative draw of TFP assigned to the firm (as a percentile in the population TFP distribution). In Figure 5(a), the y-axis presents the firm’s labor, wage, wage bill, output, and profits, and in Figure 5(b), the y-axis presents the total rents of the firm, total rents of its workers, and workers’ share of total rents. Each y-axis value is expressed as a percent change relative to the actual value for the median-TFP firm (reported in the first column of Table 4).

⁴⁶To solve the model for firms that currently receive procurement contracts, we must account for the optimal markups in the bids, requiring that we integrate across the distribution of opportunity costs (equation 14). To overcome the computational challenge, we implement the quantile representation method proposed by Luo (2020); implementation details are provided in Online Appendix I.

			Actual	Counterfactual	Change due to procurements	
			outcomes	no procurements	Level	Relative
Labor market						
L_{jt}	Employment	(workers)	24.7	12.8	11.9	92.7%
W_{jt}	Wage	(\$1,000)	59.1	50.4	8.8	17.4%
B_{jt}	Wage bill	(\$1,000)	1,459.6	645.2	814.4	126.2%
Intermediate markets						
X_{jt}	Intermediate inputs	(\$1,000)	4,715.1	2,308.6	2,406.5	104.2%
$p_K K_{jt}$	Capital rentals	(\$1,000)	1,724.7	762.4	962.3	126.2%
Total production						
Q_{jt}	Output	(quantity)	38.3	18.7	19.5	104.2%
R_{jt}	Revenue	(\$1,000)	8,962.1	4,541.6	4,420.5	97.3%
Private production						
Q_{jt}^H	Output	(quantity)	13.7	18.7	-5.1	-27.0%
R_{jt}^H	Revenue	(\$1,000)	3,460.7	4,541.6	-1,080.9	-23.8%
Rents						
V_{jt}	Worker rents	(\$1,000/worker)	11.6	5.1	6.5	126.2%
π_{jt}	Firm rents or Profits	(\$1,000/worker)	43.1	33.4	9.6	28.7%

Table 4: Outcomes of Firms and Workers and the Incidence of Procurement

Notes: For the median-TFP firm in the sample of firms that received procurement contracts ($D_{jt} = 1$), this table presents the observed values of various outcomes (column 1) as well as outcomes that would have been experienced if the firm had not received a procurement contract (column 2) using the approach of Section 5.6. It presents the differences between columns 1 and 2 in levels (column 3) and in percent changes (column 4).

We find that, when the firm is more productive (above-median TFP), it chooses to produce more output, which requires hiring more workers. Since the labor supply curve is upward-sloping, it must bid up wages to increase employment, which also increases the wage bill. If TFP is set to the 75th percentile, the firm employs 12% more labor, pays 3% higher wages, and spends 15% more on labor. It produces 65% more output and earns 74% more profits. By contrast, if TFP is set to the 25th percentile, it produces 26% less output and earns 37% lower profit. If TFP is set below the 25th percentile, the firm also hires more workers and pays greater wages than with median TFP. This is because it needs to produce the minimum output specified by the government in the procurement contract, \bar{Q}^G , and must compensate for low productivity by hiring more labor than it would with median TFP.⁴⁷ Since firm

⁴⁷Online Appendix Figure A.9 provides a similar analysis but for the majority of firms that do not currently have a procurement contract and thus have no incidence of procurements. We find that wages, employment, the wage bill, and rents are monotonically increasing in TFP when shutting

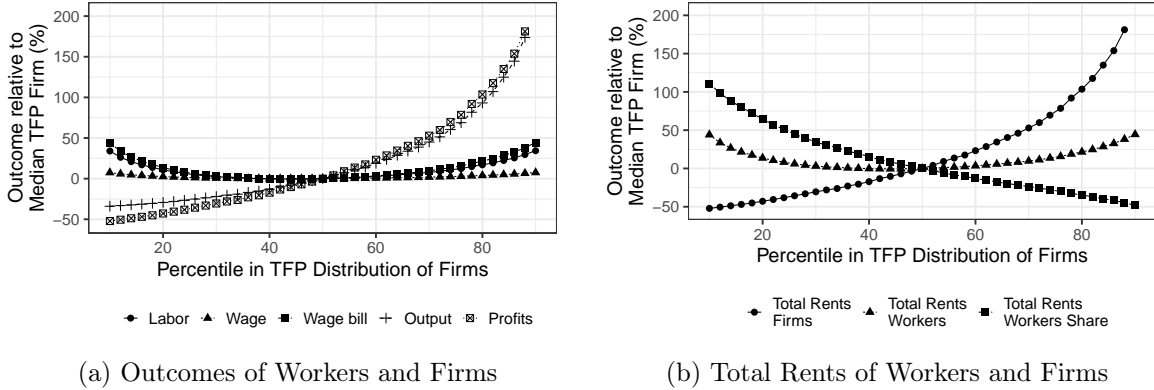


Figure 5: Outcomes and Rents for Alternative TFP Percentiles

Notes: In this figure, we assign alternative TFP quantiles to the median-TFP firm in the $D_{jt} = 1$ sample without changing any other primitives of the model, then re-solve the model to obtain this firm’s alternative outcomes and rents. The x-axis displays the alternative TFP assigned to the firm (as a percentile in the population TFP distribution). Each y-axis value is expressed as a percent change relative to the actual value for the median-TFP firm (reported in the first column of Table 4).

rents increase more than worker rents as TFP increases, the share of rents captured by workers is decreasing in TFP. This result complements the recent literature on product market competition which has found that more productive firms have higher markups and lower labor shares (Autor et al., 2020; de Loecker et al., 2020). We account for both labor and product market power in a constant-elasticity framework and find a lower rent share to workers at more productive firms.

Last, we investigate the incidence of government procurements on firms and workers. The second column of Table 4 provides our estimates of the outcomes that would have prevailed if the firm had not received a procurement contract using the approach of Section 5.6. The difference between columns 1 and 2 is the incidence of a procurement contract on the outcomes of the firm and its workers, which are both presented in absolute level changes (column 3) and changes relative to the case in which the firm does not receive a procurement contract (column 4). We find that receiving a procurement contract induces it to hire 12 more workers (nearly doubling the firm’s

down the government market, confirming that the non-monotonicity in these outcomes in Figure 5 is due to the constraint that $Q_{1jt} \geq \bar{Q}^G$.

workforce) and pay each of its workers about \$8,800 more in wages (a 17% increase), increasing its wage bill by about \$0.8 million. Worker rents increase by about \$6,500 per worker (more than double the baseline rents). Using the decomposition of incidence in Section 2.5, 70% of worker rents generated by the procurement contract accrue to incumbent workers rather than new hires.

In response to receiving the procurement contract, the median firm increases expenditure on intermediate inputs by \$2.4 million (about double the baseline) and capital rental by nearly \$1.0 million (more than double the baseline). Total output approximately doubles. Government demand crowds out private market production, with private market output decreasing by about 27%. Total revenues increase by \$4.4 million, while private market revenues decrease by about \$1.0 million. Comparing revenues to expenditures on all inputs, firm rents (i.e. profits) increase by \$9,600 per worker. Thus, 40% of the rents generated by government procurement is captured by workers, and workers receive a larger share of the rents for the marginal procurement contract than for the baseline output if only operating in the private product market.

9 Conclusion

Existing work on imperfect competition typically focuses on either the labor market or the product market in isolation. In contrast, we analyzed imperfect competition in both markets jointly, showing theoretically and empirically how the impacts of market power in one market depend on market power in the other market. Our context was the US construction industry. We developed, identified and estimated a model where construction firms imperfectly compete with one another for workers in the labor market and for projects in both the private market and the government market, where government projects are procured through auctions. Our analyses combined the universe of business and worker tax records with newly collected records from government procurement auctions. We used the estimated model to quantify the markdown of wages and the markup of prices, to show that the impacts of an increase in market power in one market are attenuated by the existence of market power in the other market, and to quantify the rents, rent-sharing, and incidence of procurements in the US construction industry.

References

- ACKERBERG, D. A., K. CAVES, AND G. FRAZER (2015): “Identification properties of recent production function estimators,” *Econometrica*, 83, 2411–2451.
- AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2015): “The skill complementarity of broadband internet,” *The Quarterly Journal of Economics*, 130, 1781–1824.
- ATHEY, S. AND P. A. HAILE (2007): “Nonparametric approaches to auctions,” *Handbook of econometrics*, 6, 3847–3965.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): “The fall of the labor share and the rise of superstar firms,” *The Quarterly Journal of Economics*, 135, 645–709.
- AZAR, J., S. BERRY, AND I. E. MARINESCU (2021): “Estimating labor market power,” .
- BALAT, J., T. KOMAROVA, AND E. KRASNOKUTSKAYA (2017): “Ex-ante and Ex-post Subcontracting in Highway Procurement Markets,” Tech. rep., Working Paper, Johns Hopkins University.
- BERGER, D., K. HERKENHOFF, AND S. MONGEY (2022): “Labor market power,” *American Economic Review*, 112, 1147–93.
- BHASKAR, V., A. MANNING, AND T. TO (2002): “Oligopsony and Monopsonistic Competition in Labor Markets,” *Journal of Economic Perspectives*, 16, 155–174.
- BLS (2021): “Labor Force Statistics from the Current Population Survey. 2002 - 2021 Annual Averages: Persons at work in nonagricultural industries by class of worker and usual full- or part-time status.” *Data Tables, U.S. Bureau of Labor Statistics*.
- BOAL, W. M. AND M. R. RANSOM (1997): “Monopsony in the labor market,” *Journal of Economic Literature*, 35, 86–112.
- BONHOMME, S., K. HOLZHEU, T. LAMADON, E. MANRESA, M. MOGSTAD, AND B. SETZLER (2023): “How Much Should we Trust Estimates of Firm Effects and Worker Sorting?” *Journal of Labor Economics*.
- CALDWELL, S. AND E. OEHLSEN (2018): “Monopsony and the Gender Wage Gap: Experimental Evidence from the Gig Economy,” *MIT Working Paper*.
- CALLAWAY, B. AND P. H. SANT’ANNA (2020): “Difference-in-differences with multiple time periods,” *Journal of Econometrics*.
- CARD, D. (2022): “Who Set Your Wage?” *American Economic Review*, 112, 1075–90.
- 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.
- CASTRO-VINCENZI, J. M. AND B. KLEINMAN (2022): “Intermediate input prices and the labor share,” *Unpublished manuscript, Princeton University*.

- CBO (2007): “The Impact of Unauthorized Immigrants on the Budgets of State and Local Governments,” Tech. rep., Congressional Budget Office.
- CHAN, M., K. KROFT, E. MATTANA, AND I. MOURIFIÉ (2023): “An Empirical Framework for Matching with Imperfect Competition,” *Unpublished manuscript*.
- CHASSANG, S., K. KAWAI, J. NAKABAYASHI, AND J. M. ORTNER (2022): “Robust Screens for Noncompetitive Bidding in Procurement Auctions,” *Econometrica*, 90, 315–346.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): “The rise of market power and the macroeconomic implications,” *The Quarterly Journal of Economics*, 135, 561–644.
- DIXIT, A. K. AND J. E. STIGLITZ (1977): “Monopolistic competition and optimum product diversity,” *The American Economic Review*, 67, 297–308.
- DUBE, A., A. MANNING, AND S. NAIDU (2020): “Monopsony and employer mis-optimization explain why wages bunch at round numbers,” *Unpublished manuscript*.
- DUBE, A., S. NAIDU, AND A. D. REICH (2022): “Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers,” *Unpublished manuscript*.
- DUNCAN, K. AND R. ORMISTON (2019): “What does the research tell us about prevailing wage laws?” *Labor Studies Journal*, 44, 139–160.
- FERRAZ, C., F. FINAN, AND D. SZERMAN (2015): “Procuring firm growth: the effects of government purchases on firm dynamics,” NBER Working Paper w21219.
- GANDHI, A., S. NAVARRO, AND D. A. RIVERS (2020): “On the identification of gross output production functions,” *Journal of Political Economy*, 128, 2973–3016.
- GOLDBERG, P. K. AND M. M. KNETTER (1999): “Measuring the intensity of competition in export markets,” *Journal of International Economics*, 47, 27–60.
- GOLDBERG, P. K. AND F. VERBOVEN (2001): “The evolution of price dispersion in the European car market,” *The Review of Economic Studies*, 68, 811–848.
- GUGLER, K., M. WEICHSELBAUMER, AND C. ZULEHNER (2020): “Employment behavior and the economic crisis: Evidence from winners and runners-up in procurement auctions,” *Journal of Public Economics*, 182, 104112.
- GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2021): “What do data on millions of US workers reveal about lifecycle earnings dynamics?” *Econometrica*, 89, 2303–2339.
- HOWELL, S. T. AND J. D. BROWN (2020): “Do Cash Windfalls Affect Wages? Evidence from R&D Grants to Small Firms,” CES Working Paper 20-06, U.S. Census Bureau.
- HVIDE, H. K. AND T. MELING (2020): “Do Temporary Demand Shocks Have Long-Term Effects for Startups?” *SSRN 3437270*.

- JAROSCH, G., J. S. NIMCZIK, AND I. SORKIN (2023): “Granular Search, Market Structure, and Wages,” NBER Working Paper w26239.
- KESSLER, D. P. AND L. F. KATZ (2001): “Prevailing Wage Laws and Construction Laborc Markets,” *ILR Review*, 54, 259–274.
- 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 (2022): “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market,” *American Economic Review*, 112, 169–212.
- LEVINSOHN, J. AND A. PETRIN (2003): “Estimating production functions using inputs to control for unobservables,” *The Review of Economic Studies*, 70, 317–341.
- LEWIS, G. AND P. BAJARI (2011): “Procurement Contracting With Time Incentives: Theory and Evidence,” *The Quarterly Journal of Economics*, 126, 1173–1211.
- LUO, Y. (2020): “Unobserved heterogeneity in auctions under restricted stochastic dominance,” *Journal of Econometrics*, 216, 354–374.
- MACKENZIE, G. (2021): “Trade and market power in product and labor markets,” .
- MANNING, A. (2003): *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press.
- (2011): “Imperfect competition in the labor market,” *Handbook of Labor Economics*, 4, 973–1041.
- MASKIN, E. AND J. RILEY (1984): “Optimal auctions with risk averse buyers,” *Econometrica*, 1473–1518.
- MILGROM, P. R. AND R. J. WEBER (1982): “A theory of auctions and competitive bidding,” *Econometrica*, 1089–1122.
- ROBINSON, J. (1933): *The Economics of Imperfect Competition*, Macmillan and Co.
- ROSEN, S. (1986): “The theory of equalizing differences,” *Handbook of Labor Economics*, 1, 641–692.
- ROUSSILLE, N. AND B. SCUDERI (2022): “Bidding for Talent,” .
- RSMEANS (2008): *Building Construction Cost Data*, 66 ed.
- SHARMA, G. (2022): “Monopsony and Gender,” .
- SOKOLOVA, A. AND T. SORENSEN (2021): “Monopsony in labor markets: A meta-analysis,” *ILR Review*, 74, 27–55.
- SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): “Firming up inequality,” *The Quarterly Journal of Economics*, 134, 1–50.

- 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.
- TAKAHASHI, H. (2018): “Strategic design under uncertain evaluations: structural analysis of design-build auctions,” *The RAND Journal of Economics*, 49, 594–618.
- VAN REENEN, J. (1996): “The creation and capture of rents: wages and innovation in a panel of UK companies,” *The Quarterly Journal of Economics*, 111, 195–226.
- YAGAN, D. (2019): “Employment hysteresis from the great recession,” *Journal of Political Economy*, 127, 2505–2558.

Online Appendix to “Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry”

A Model Derivations

A.1 The Composite Production Function

In this appendix, we derive equation (9). To do so, we express revenues and costs as functions of Q_{jt} so as to separate the joint maximization into two steps: In the first step, we find the optimal combination (K_{jt}, L_{jt}) for each Q_{jt} . In the second step, we solve for the optimal Q_{jt} .

Recall that firms can rent capital at price p_K and hire labor at price $W_{jt} = U_{jt}L_{jt}^\theta$. The production function (in physical units) satisfies

$$Q_{jt} = \Omega_{jt} K_{jt}^{\beta_K} L_{jt}^{\beta_L}. \quad (29)$$

Intermediate inputs have constant price p_M and, due to the Leontief functional form, must satisfy $M_{jt} = Q_{jt}/\beta_M$. Given any production level Q_{jt} , the firm can find the most cost efficient combination (K_{jt}, L_{jt}) by solving the Hicksian cost-minimization problem,

$$\min_{(K_{jt}, L_{jt}): Q_{jt} = \Omega_{jt} K_{jt}^{\beta_K} L_{jt}^{\beta_L}} p_K K_{jt} + U_{jt} L_{jt}^{1+\theta}, \quad (30)$$

where $p_K K_{jt} + U_{jt} L_{jt}^{1+\theta}$ is the total cost of capital and labor. We now solve for the Hicksian demand for capital and labor using the Lagrangian,

$$\mathcal{L}_{jt} \equiv p_K K_{jt} + U_{jt} L_{jt}^{1+\theta} + \lambda_{jt} (Q_{jt} - \Omega_{jt} K_{jt}^{\beta_K} L_{jt}^{\beta_L}), \quad (31)$$

where λ_{jt} is the Lagrange multiplier. The first-order conditions for capital and labor, respectively, are as follows:

$$p_K = \lambda_{jt} \Omega_{jt} \beta_K K_{jt}^{\beta_K - 1} L_{jt}^{\beta_L}, \quad (32)$$

$$(1 + \theta) U_{jt} L_{jt}^\theta = \lambda_{jt} \Omega_{jt} \beta_L K_{jt}^{\beta_K} L_{jt}^{\beta_L - 1}. \quad (33)$$

These equations lead to the optimal choice of capital as a function of labor:

$$K_{jt} = \frac{\beta_K (1 + \theta)}{\beta_L p_K} U_{jt} L_{jt}^{1 + \theta}. \quad (34)$$

Substituting equation (34) into equation (29), the inverse Hicksian demand is

$$Q_{jt} = \Omega_{jt} \left(\frac{\beta_K (1 + \theta)}{\beta_L p_K} U_{jt} L_{jt}^{1 + \theta} \right)^{\beta_K} L_{jt}^{\beta_L} = \Phi_{jt} L_{jt}^\rho, \quad (35)$$

where we define $\rho \equiv (1 + \theta)\beta_K + \beta_L$ and $\Phi_{jt} \equiv \Omega_{jt} \left(\frac{\beta_K (1 + \theta)}{\beta_L p_K} U_{jt} \right)^{\beta_K}$. Thus, $L_{jt} = (Q_{jt}/\Phi_{jt})^{1/\rho}$. Substituting into the first-order condition for intermediate inputs (equation 10), the Hicksian expenditure on intermediate inputs is

$$X_{jt} \equiv p_M M_{jt} = p_M Q_{jt} / \beta_M = \frac{p_M}{\beta_M} \Phi_{jt} L_{jt}^\rho. \quad (36)$$

Lastly, letting $\kappa_U \equiv \frac{\beta_K}{\beta_L} (1 + \theta) + 1$, total costs can be expressed purely in terms of labor as

$$W_{jt} L_{jt} + p_K K_{jt} + p_M M_{jt} = \kappa_U U_{jt} L_{jt}^{1 + \theta} + \frac{p_M}{\beta_M} \Phi_{jt} L_{jt}^\rho. \quad (37)$$

A.2 Firm's Behavior in the Private Product Market

In this appendix, we derive equation (13) and several related results on firm behavior in the private market. We assume a downward-sloping private product demand curve ($\epsilon > 0$) and increasing composite returns to labor ($\rho > 1$), consistent with the empirical evidence.

If $d = 0$, the firm's profit maximization problem is,

$$\max_{L_{0jt}} p_H (\Phi_{jt} L_{0jt}^\rho)^{1 - \epsilon} - \kappa_U U_{jt} L_{0jt}^{1 + \theta} - \frac{p_M}{\beta_M} \Phi_{jt} L_{0jt}^\rho, \quad (38)$$

where we substituted equations (9) and (37) into equation (8) for the case with $d = 0$.

The profit-maximizing first-order condition is,

$$\frac{\partial \pi_{0jt}}{\partial L_{0jt}} \equiv \underbrace{p_H \Phi_{jt}^{1-\epsilon} (1-\epsilon) \rho L_{0jt}^{(1-\rho)\epsilon - (1-\rho) - \epsilon}}_{\text{MRP}} - \underbrace{\left(\kappa_U U_{jt} (1+\theta) L_{0jt}^\theta + \frac{p_M}{\beta_M} \Phi_{jt} \rho L_{0jt}^{-(1-\rho)} \right)}_{\text{MCL}} = 0. \quad (39)$$

This expression shows that L_{0jt} only varies across firms due to Φ_{jt} and U_{jt} . We now derive equation (13) and several implications. Equation (39) can be arranged as

$$\underbrace{\overbrace{(1-\epsilon)}^{\text{markup}^{-1}} \underbrace{p_H \Phi_{jt}^{-\epsilon} L_{0jt}^{-\rho\epsilon}}_{P_{jt}} \underbrace{\Phi_{jt} \rho L_{0jt}^{\rho-1} / \kappa_U}_{\text{MP}_{jt}}}_{\text{MRP}_{jt}} = \underbrace{\overbrace{(1+\theta)}^{\text{markdown}^{-1}} \underbrace{U_{jt} L_{0jt}^\theta}_{W_{jt}}}_{\text{MCL}_{jt}} + \underbrace{\frac{p_M}{\beta_M} \Phi_{jt} \rho L_{0jt}^{\rho-1} / \kappa_U}_{\text{marginal intermed. costs}}, \quad (40)$$

which is the same as equation (19), where we use that $\beta_M = Q_{jt}/M_{jt}$ implies $\frac{p_M}{\beta_M} = \frac{p_M M_{jt}}{Q_{jt}} = \frac{p_M M_{jt}}{Q_{jt} P_{jt}} P_{jt} = \frac{X_{jt}}{R_{jt}} P_{jt}$.

We will now show that MRP is greater than MCL as L_{0jt} approaches zero. Multiplying marginal profits in (39) by $L_{0jt}^{\rho\epsilon + (1-\rho)}$, which is strictly positive, we have

$$\frac{\partial \pi_{0jt}}{\partial L_{0jt}} L_{0jt}^{\rho\epsilon + (1-\rho)} = \underbrace{p_H \Phi_{jt}^{1-\epsilon} (1-\epsilon) \rho}_{\text{MRP} \times L_{0jt}^{\rho\epsilon + (1-\rho)}} - \underbrace{\left(\kappa_U U_{jt} (1+\theta) L_{0jt}^{\theta + \rho\epsilon + (1-\rho)} + \frac{p_M}{\beta_M} \Phi_{jt} \rho L_{0jt}^{\rho\epsilon} \right)}_{\text{MCL} \times L_{0jt}^{\rho\epsilon + (1-\rho)}}. \quad (41)$$

Note that $\text{MRP} \times L_{0jt}^{\rho\epsilon + (1-\rho)}$ is constant with respect to L_{0jt} and positive. By contrast, given $\theta + \rho\epsilon + (1-\rho) > 0$, then $\text{MCL} \times L_{0jt}^{\rho\epsilon + (1-\rho)}$ converges to zero as L_{0jt} approaches zero. Thus, we have shown that $\lim_{L_{0jt} \rightarrow 0^+} \frac{\partial \pi_{0jt}}{\partial L_{0jt}} > 0$. As a result, it is always optimal to choose $L_{0jt} > 0$ if $\theta + \rho\epsilon + (1-\rho) > 0$.

Furthermore, multiplying both sides of equation (39) by L_{0jt} , we have

$$\frac{\partial \pi_{0jt}}{\partial L_{0jt}} \equiv p_H \Phi_{jt}^{1-\epsilon} (1-\epsilon) \rho L_{0jt}^{\rho(1-\epsilon)} - \kappa_U U_{jt} (1+\theta) L_{0jt}^{1+\theta} - \frac{p_M}{\beta_M} \Phi_{jt} \rho L_{0jt}^\rho = 0. \quad (42)$$

Recall that $\kappa_U U_{jt} L_{0jt}^{1+\theta} = \frac{\rho}{\beta_L} B_{0jt}$, $R_{0jt}^H = p_H \Phi_{jt}^{1-\epsilon} L_{0jt}^{\rho(1-\epsilon)}$, and $X_{0jt} = \frac{p_M}{\beta_M} \Phi_{jt} L_{0jt}^\rho$. Substituting, we have

$$(1-\epsilon) R_{0jt}^H = \frac{1+\theta}{\beta_L} B_{0jt} + X_{0jt}. \quad (43)$$

Similarly, if $d = 1$, the firm's profit maximization problem is,

$$\max_{L_{1jt}: \Phi_{jt}L_{1jt}^\rho \geq \bar{Q}^G} p_H \left(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G \right)^{1-\epsilon} - \kappa_U U_{jt} L_{1jt}^{1+\theta} - \frac{p_M}{\beta_M} \Phi_{jt} L_{1jt}^\rho. \quad (44)$$

The first-order condition is,

$$\frac{\partial \pi_{1jt}}{\partial L_{1jt}} \equiv p_H \Phi_{jt} (1-\epsilon) \rho \left(\Phi_{jt} L_{1jt}^\rho - \bar{Q}^G \right)^{-\epsilon} L_{1jt}^{-(1-\rho)} - \kappa_U U_{jt} (1+\theta) L_{1jt}^\theta - \frac{p_M}{\beta_M} \Phi_{jt} \rho L_{1jt}^{-(1-\rho)} = 0. \quad (45)$$

As $\Phi_{jt}L_{1jt}^\rho$ approaches \bar{Q}^G , $\left(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G \right)^{-\epsilon}$ approaches infinity while all other terms involving L_{1jt} approach constants. Thus, $\Phi_{jt}L_{1jt}^\rho > \bar{Q}^G$ is necessary to satisfy the equation. Since $Q_{1jt} = \Phi_{jt}L_{1jt}^\rho$, it follows that $Q_{1jt}^H = Q_{1jt} - \bar{Q}^G > 0$, so the winning firm always produces for the private market. Furthermore, it is always true that $\frac{\partial \pi_{1jt}}{\partial L_{1jt}}|_{L_{1jt}=L_{0jt}} > 0$. Thus, Q_{djt} is larger if $d = 1$ than $d = 0$.

Multiplying both sides of equation (45) by L_{1jt} and replacing $R_{1jt}^H = p_H \left(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G \right)^{1-\epsilon}$,

$$R_{1jt}^H \Phi_{jt} (1-\epsilon) \left(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G \right)^{-1} L_{1jt}^\rho - \frac{1+\theta}{\beta_L} B_{1jt} - X_{1jt} = 0. \quad (46)$$

Since $Q_{1jt}^H = \Phi_{jt}L_{1jt}^\rho - \bar{Q}^G$ and $Q_{1jt} = \Phi_{jt}L_{1jt}^\rho$, it follows that

$$R_{1jt}^H (1-\epsilon) \frac{Q_{1jt}}{Q_{1jt}^H} - \frac{1+\theta}{\beta_L} B_{1jt} - X_{1jt} = 0. \quad (47)$$

Thus, combining equations (43) and (47), we have equation (13).

Lastly, it is interesting to consider if winning a procurement project will lead a firm to produce more for the private market (crowd-in) or less (crowd-out). To determine this, we evaluate the marginal profits of the winner when the total output is $\hat{Q}_{1jt} \equiv \bar{Q}^G + Q_{0jt}^H$; that is, \hat{Q}_{1jt} is the hypothetical output of the firm in the $d = 1$ case such that there is neither crowd-in nor crowd-out. The winner would prefer to produce more (less) than \hat{Q}_{1jt} if the marginal profit is positive (negative, respectively). Let the corresponding labor choice be \hat{L}_{1jt} such that $\Phi_{jt}\hat{L}_{1jt}^\rho - \bar{Q}^G = Q_{0jt}^H = \Phi_{jt}L_{0jt}^\rho$. Note that, since $\rho > 1$ and $\hat{L}_{1jt}^\rho = L_{0jt}^\rho + \bar{Q}^G/\Phi_{jt}$, then $\hat{L}_{1jt} > L_{0jt}$. Evaluating equation (45) at \hat{L}_{1jt} , marginal profits for the firm if it wins and produces hypothetical output \hat{Q}_{1jt} are,

$$\frac{\partial \pi_{1jt}}{\partial L_{1jt}}|_{L_{1jt}=\hat{L}_{1jt}} = p_H \Phi_{jt} (1-\epsilon) \rho (Q_{0jt}^H)^{-\epsilon} \hat{L}_{1jt}^{\rho-1} - \kappa_U U_{jt} (1+\theta) \hat{L}_{1jt}^\theta - \frac{p_M}{\beta_M} \Phi_{jt} \rho \hat{L}_{1jt}^{\rho-1}. \quad (48)$$

Multiplying by $L_{1jt}^{1-\rho}$ and substituting $Q_{0jt}^H = \Phi_{jt}L_{0jt}^\rho$, we have,

$$L_{1jt}^{1-\rho} \frac{\partial \pi_{1jt}}{\partial L_{1jt}} \Big|_{L_{1jt}=\hat{L}_{1jt}} = p_H \Phi_{jt}^{1-\epsilon} (1-\epsilon) \rho L_{0jt}^{-\rho\epsilon} - \kappa_U U_{jt} (1+\theta) \hat{L}_{1jt}^{\theta+1-\rho} - \frac{p_M}{\beta_M} \Phi_{jt} \rho. \quad (49)$$

Finally, substituting equation (42), this simplifies to,

$$L_{1jt}^{1-\rho} \frac{\partial \pi_{1jt}}{\partial L_{1jt}} \Big|_{L_{1jt}=\hat{L}_{1jt}} = \kappa_U U_{jt} (1+\theta) (L_{0jt}^{\theta+1-\rho} - \hat{L}_{1jt}^{\theta+1-\rho}). \quad (50)$$

Since $\hat{L}_{1jt} > L_{0jt}$, we have that $\frac{\partial \pi_{1jt}}{\partial L_{1jt}} \Big|_{L_{1jt}=\hat{L}_{1jt}} < 0$ if $\theta+1-\rho > 0$, and $\frac{\partial \pi_{1jt}}{\partial L_{1jt}} \Big|_{L_{1jt}=\hat{L}_{1jt}} > 0$ otherwise. Therefore, winning a government project crowds-out private projects when $1+\theta > \rho$ and crowds-in if $1+\theta < \rho$.

A.3 Worker Rents Expressions

We derive the key expressions for worker rents in Section 2.5. First, following Lamadon et al. (2022), total worker rents at firm j in year t are,

$$V_{jt}(W_{jt}) = \int_0^{W_{jt}} (W_{jt} - W) \frac{dL_{jt}(W)}{dW} dW. \quad (51)$$

Define $\omega \equiv \frac{W}{W_{jt}}$ so that $\frac{d\omega}{dW} = \frac{1}{W_{jt}}$, and note that labor supply can be expressed in terms ω as $\tilde{L}(\omega W_{jt}) = \omega^{1/\theta} L(W_{jt})$. Thus, $\frac{dL_{jt}(W)}{dW} dW = \frac{d\tilde{L}_{jt}(\omega W_{jt})}{d\omega} d\omega$. Moreover, $\frac{d\tilde{L}_{jt}(\omega W_{jt})}{d\omega} = L_{jt}(W_{jt}) \frac{\partial \omega^{1/\theta}}{\partial \omega}$, which implies

$$V_{jt}(W_{jt}) = W_{jt} \int_0^1 (1 - \omega) \frac{d\tilde{L}(\omega W_{jt})}{d\omega} d\omega = W_{jt} L_{jt}(W_{jt}) \int_0^1 (1 - \omega) \frac{\partial \omega^{1/\theta}}{\partial \omega} d\omega = \frac{W_{jt} L_{jt}(W_{jt})}{1 + 1/\theta}$$

Second, we derive the decomposition of incidence for incumbents and new hires. Let \mathcal{I}_{jt} denote the set of incumbent workers at firm j . For two potential wages W_{1jt} and W_{0jt} , the corresponding rents V_{1ijt} and V_{0ijt} for any $i \in \mathcal{I}_{jt}$ must satisfy,

$$\mathcal{U}_{it}(j, W_{1jt} - V_{1ijt}) = \max_{j' \neq j} \mathcal{U}_{it}(j', W_{j',t}) \quad \text{and} \quad \mathcal{U}_{it}(j, W_{0jt} - V_{0ijt}) = \max_{j' \neq j} \mathcal{U}_{it}(j', W_{j',t}).$$

Since the right-hand side is the same in both of these equations (that is, the outside option is unchanged by a wage increase at the incumbent employer), it follows that $\mathcal{U}_{it}(j, W_{1jt} - V_{1ijt}) = \mathcal{U}_{it}(j, W_{0jt} - V_{0ijt})$, $\forall i \in \mathcal{I}_{jt}$, which can only be satisfied by

$V_{1ijt} - V_{0ijt} = W_{1jt} - W_{0jt}$, $\forall i \in \mathcal{I}_{jt}$. Thus, the incidence for incumbents is

$$\underbrace{\sum_{\mathcal{I}_{jt}} (V_{1ijt} - V_{0ijt})}_{\text{Incidence for incumbents}} = \underbrace{L_{0jt}}_{\text{Number of incumbents}} \times \underbrace{(W_{1jt} - W_{0jt})}_{\text{Incidence for each incumbent}}.$$

The incidence for new hires is then,

$$\underbrace{V_{1jt} - V_{0jt}}_{\text{Incidence}} - \underbrace{L_{0jt} (W_{1jt} - W_{0jt})}_{\text{Incidence for incumbents}} = \underbrace{\frac{W_{1jt}L_{1jt}}{1 + 1/\theta} - \frac{W_{0jt}L_{0jt}}{1 + 1/\theta} - L_{0jt} (W_{1jt} - W_{0jt})}_{\text{Incidence for new hires}}$$

which can be rearranged as the decomposition in Section 2.5.

A.4 Over-identifying Restriction

In this appendix, we derive equation (25). Taking the log of both sides of equation (13) for the $d = 1$ case, we have,

$$\log(1-\epsilon) + r_{1jt}^H + q_{1jt} - q_{1jt}^H = \log\left(\frac{1+\theta}{\beta_L} B_{1jt} + X_{1jt}\right).$$

From equation (7), $r_{1jt}^H = \log p_H + (1-\epsilon)q_{1jt}^H$, so $q_{1jt}^H = \frac{1}{1-\epsilon}r_{1jt}^H - \frac{1}{1-\epsilon}\log p_H$. From equation (10), $q_{1jt} = \rho\ell_{1jt} + \phi_{jt} + e_{jt}$. Substituting, we have

$$\log(1-\epsilon) + r_{1jt}^H + (\rho\ell_{1jt} + \phi_{jt} + e_{jt}) - \left(\frac{1}{1-\epsilon}r_{1jt}^H - \frac{1}{1-\epsilon}\log p_H\right) = \log\left(\frac{1+\theta}{\beta_L} B_{1jt} + X_{1jt}\right),$$

which can be rearranged as equation (25).

B Product Market with Perfect Competition

This section solves the firm's problem in the private product market assuming the firm is a price-taker ($\epsilon = 0$). Denote the competitive price as p_H . In terms of the composite production function $Q_{jt} = \Phi_{jt}L_{jt}^\rho$, the firm's problem is

$$\max_{L_{jt}: \Phi_{jt}L_{jt}^\rho \geq d\bar{Q}^G} p_H \left(\Phi_{jt}L_{jt}^\rho - d\bar{Q}^G \right) - \kappa_U U_{jt} L_{jt}^{1+\theta} - \frac{p_M}{\beta_M} \Phi_{jt} L_{jt}^\rho$$

where the government's output must be produced if the firm receives a procurement contract ($\Phi_{jt}L_{jt}^\rho \geq d\bar{Q}^G$). We consider three cases:

Suppose $d = 0$. The government constraint is always satisfied, so we can ignore

this constraint. The profit-maximizing solution is simply $Q_{0jt} = \Phi_{jt}L_{0jt}^\rho$ and

$$L_{0jt} = \left(\left(p_H - \frac{p_M}{\beta_M} \right) \frac{\Phi_{jt}^\rho}{\kappa_U U_{jt} (1 + \theta)} \right)^{\frac{1}{\theta+1-\rho}}.$$

Suppose $d = 1$ and $Q_{0jt} > \bar{Q}^G$. Then, the solution $L_{1jt}^{interior} = L_{0jt}$ and $Q_{1jt}^{interior} = Q_{0jt}$ satisfies the government constraint and otherwise solves the profit-maximization problem, so this is the optimal solution. An implication is that $Q_{1jt}^{interior}$ is invariant to marginal changes in the size of the government contract, i.e., government projects crowd-out the firm's private market production one-for-one. Since input costs are not affected by receiving a procurement contract, the opportunity cost of receiving a procurement contract is simply the loss in revenues in the private product market, $\sigma_{jt}^{interior} = p_H \left(Q_{0jt} - \left(Q_{1jt}^{interior} - \bar{Q}^G \right) \right) = p_H \bar{Q}^G$.

Suppose $d = 1$ and $Q_{0jt} \leq \bar{Q}^G$. Then, the firm is at the corner solution in which it only produces for the government market, i.e., $Q_{1jt}^{corner} = \bar{Q}^G$ and $L_{1jt}^{corner} = \left(\bar{Q}^G / \Phi_{jt} \right)^{1/\rho}$. The opportunity cost is $\sigma_{jt}^{corner} = p_H Q_{0jt} - \{ T_{jt}(L_{0jt}) - T_{jt}(L_{1jt}^{corner}) \}$, where $T_{jt}(L) \equiv \kappa_U U_{jt} L^{1+\theta} + \frac{p_M}{\beta_M} \Phi_{jt} L^\rho$ is the total cost of production using labor L .

C Cobb-Douglas Production Function

C.1 Cobb-Douglas Model: Composite Production Function

Consider a Cobb-Douglas production function (in physical units)

$$Q_{jt} = \Omega_{jt} L_{jt}^{\beta_L} K_{jt}^{\beta_K} M_{jt}^{\beta_M}. \quad (52)$$

Given any production level Q_{jt} , the firm can find the most cost efficient combination (L_{jt}, K_{jt}, M_{jt}) by solving the cost-minimization problem,

$$\min_{L_{jt}, K_{jt}, M_{jt}} C_{jt} \quad \text{s.t.} \quad Q_{jt} = \Omega_{jt} L_{jt}^{\beta_L} K_{jt}^{\beta_K} M_{jt}^{\beta_M}. \quad (53)$$

where $C_{jt} \equiv U_{jt} L_{jt}^{1+\theta} + p_K K_{jt} + p_M M_{jt}$ denotes the total cost. This leads to the Lagrangian,

$$\mathcal{L}_{jt} \equiv U_{jt} L_{jt}^{1+\theta} + p_K K_{jt} + p_M M_{jt} + \lambda_{jt} (Q_{jt} - \Omega_{jt} K_{jt}^{\beta_K} L_{jt}^{\beta_L} M_{jt}^{\beta_M}) \quad (54)$$

where λ_{jt} is the Lagrange multiplier. The first-order conditions are:

$$\begin{aligned} p_K &= \lambda_{jt} \Omega_{jt} \beta_K K_{jt}^{\beta_K - 1} L_{jt}^{\beta_L} M_{jt}^{\beta_M}, \\ (1 + \theta) U_{jt} L_{jt}^\theta &= \lambda_{jt} \Omega_{jt} \beta_L K_{jt}^{\beta_K} L_{jt}^{\beta_L - 1} M_{jt}^{\beta_M}, \\ p_M &= \lambda_{jt} \Omega_{jt} \beta_M K_{jt}^{\beta_K} L_{jt}^{\beta_L} M_{jt}^{\beta_M - 1}. \end{aligned} \quad (55)$$

We can use these first-order conditions to write the optimal choices of capital and intermediate inputs as a function of labor

$$K_{jt} = \frac{\beta_K (1 + \theta) U_{jt}}{\beta_L p_K} L_{jt}^{1+\theta} = \chi^{(K)} U_{jt} L_{jt}^{1+\theta} \text{ and } M_{jt} = \frac{\beta_M (1 + \theta) U_{jt}}{\beta_L p_M} L_{jt}^{1+\theta} = \chi^{(M)} U_{jt} L_{jt}^{1+\theta} \quad (56)$$

where $\chi^{(K)} \equiv \frac{\beta_K (1+\theta)}{\beta_L p_K}$ and $\chi^{(M)} \equiv \frac{\beta_M (1+\theta)}{\beta_L p_M}$. We can substitute these expressions into $Q_{jt} = \Omega_{jt} L_{jt}^{\beta_L} K_{jt}^{\beta_K} M_{jt}^{\beta_M}$ and obtain

$$Q_{jt} = \Omega_{jt} \left[\chi_j^{(K)} U_{jt} L_{jt}^{1+\theta} \right]^{\beta_K} L_j^{\beta_L} \left[\chi_j^{(M)} U_{jt} L_{jt}^{1+\theta} \right]^{\beta_M} = \Phi_{jt} L_{jt}^\rho \quad (57)$$

where $\Phi_{jt} \equiv \Omega_{jt} \left[\chi^{(K)} U_{jt} \right]^{\beta_K} \left[\chi^{(M)} U_{jt} \right]^{\beta_M}$ and $\rho \equiv \beta_L + (1 + \theta)(\beta_K + \beta_M)$. We can also use equations (57) and equation (56) to rewrite the firm's problem in the private product market as

$$\max_{L_{djt}} \pi_{djt} = p_H (\Phi_{jt} L_{djt}^\rho - \bar{Q}^G d)^{1-\varepsilon} - \chi^{(W)} U_{jt} L_{djt}^{1+\theta}, \quad (58)$$

where cost-minimization implies $C_{djt} = \chi^{(W)} U_{jt} L_{djt}^{\theta+1}$, $\chi^{(W)} \equiv \frac{\rho}{\beta_L} \equiv \left(1 + \frac{(\beta_M + \beta_K)(1+\theta)}{\beta_L} \right)$.

C.2 Cobb-Douglas Model: First-order Conditions

We now derive the profit-maximizing first-order conditions in the model with Cobb-Douglas production. These derivations assume $\rho \equiv \beta_L + (1 + \theta)(\beta_K + \beta_M) > 1$ and $\varepsilon > 0$.

If the firm loses the auction, its profit maximization problem is

$$\max_{L_{0jt}} p_H (\Phi_{jt} L_{0jt}^\rho)^{1-\varepsilon} - \chi^{(W)} U_{jt} L_{0jt}^{1+\theta}. \quad (59)$$

The first-order condition is,

$$\rho(1 - \varepsilon)p_H\Phi_{jt}^{1-\varepsilon}L_{0jt}^{\rho(1-\varepsilon)-1} = \chi^{(W)}U_{jt}L_{0jt}^\theta(1 + \theta), \quad (60)$$

which implies,

$$L_{0jt} = \left[\frac{\rho(1 - \varepsilon)p_H\Phi_{jt}^{1-\varepsilon}}{\chi^{(W)}U_{jt}(1 + \theta)} \right]^{\frac{1}{\theta+1-\rho(1-\varepsilon)}}. \quad (61)$$

Thus $0 < L_{0jt} < \infty$.

Similarly, if the firm wins the auction, the profit maximization problem is:

$$\max_{L_{1jt}: \Phi_{jt}L_{1jt}^\rho \geq \bar{Q}^G} p_H(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G)^{1-\varepsilon} - \chi^{(W)}U_{jt}L_{1jt}^{1+\theta}. \quad (62)$$

The first-order condition is

$$\frac{\partial \pi_{1jt}}{\partial L_{1jt}} \equiv \rho(1 - \varepsilon)\Phi_{jt}p_H(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G)^{-\varepsilon}L_{1jt}^{\rho-1} - \chi^{(W)}U_{jt}L_{1jt}^\theta(1 + \theta) = 0, \quad (63)$$

which implies,

$$\rho(1 - \varepsilon)\Phi_{jt}p_H(\Phi_{jt}L_{1jt}^\rho - \bar{Q}^G)^{-\varepsilon} = \chi^{(W)}U_{jt}L_{1jt}^{1+\theta-\rho}(1 + \theta). \quad (64)$$

As $\Phi_{jt}L_{1jt}^\rho$ approaches \bar{Q}^G , the left-hand side of equation (64) approaches infinity while the RHS approaches a constant. Thus, $\Phi_{jt}L_{1jt}^\rho > \bar{Q}^G$ is necessary to satisfy the equation. Since $Q_{1jt} = \Phi_{jt}L_{1jt}^\rho$, it follows that $Q_{1jt}^H = Q_{1jt} - \bar{Q}^G > 0$, so the winning firm always produces for the private market.

Furthermore, since the solution is interior (i.e., $L_{0jt} \geq \bar{Q}^G$) due to $\varepsilon > 0$, equation (60) implies $\rho(1 - \varepsilon)\Phi_{jt}p_H(\Phi_{jt}L_{0jt}^\rho)^{-\varepsilon}L_{0jt}^{\rho-1} - \chi^{(W)}U_{jt}L_{0jt}^\theta(1 + \theta) = 0$ and therefore $\rho(1 - \varepsilon)\Phi_{jt}p_H(\Phi_{jt}L_{0jt}^\rho - \bar{Q}^G)^{-\varepsilon}L_{0jt}^{\rho-1} - \chi^{(W)}U_{jt}L_{0jt}^\theta(1 + \theta) \equiv \frac{\partial \pi_{1jt}}{\partial L_{1jt}}|_{L_{1jt}=L_{0jt}} > 0$. Thus, $\frac{\partial \pi_{1jt}}{\partial L_{1jt}}|_{L_{1jt}=L_{0jt}} > 0$, so total production will be larger if the firm receives a procurement contract than if it does not.

C.3 Cobb-Douglas Model: Identification

We now show identification of $(1-\varepsilon, \rho, \beta_L)$ in the model with a Cobb-Douglas production function.

In the $d = 0$ case, revenues are related to labor by

$$r_{jt} = \log p_H + (1-\varepsilon)\phi_{jt} + \rho(1-\varepsilon)\ell_{jt} \quad (65)$$

From this, we can identify $\rho(1-\epsilon)$ by regressing r_{jt} on ℓ_{jt} controlling for ϕ_{jt} among $D_{jt} = 0$ firms. In practice, we can control for $(\widehat{u}_{jt}, Z_{jt})$ in place of ϕ_{jt} as in equation (22) due to the invertibility of bids with respect to TFP, conditional on amenities. Thus, $\rho(1-\epsilon)$ is recovered by the estimator

$$\widehat{\rho(1-\epsilon)} \equiv \frac{\text{Cov}[r_{jt}, \ell_{jt} | \widehat{u}_{jt}, Z_{jt}, D_{jt} = 0]}{\text{Var}[\ell_{jt} | \widehat{u}_{jt}, Z_{jt}, D_{jt} = 0]}. \quad (66)$$

In the $d = 0$ case, equation (60) implies

$$\rho(1-\epsilon) \frac{R_{0jt}^H}{L_{0jt}} = \chi^{(W)} U_{jt} L_{0jt}^\theta (1 + \theta).$$

Since we showed above that cost-minimization requires $C_{dj} = \chi^{(W)} U_{jt} L_{dj}^{\theta+1}$, it follows that

$$\rho(1-\epsilon) = (1 + \theta) \frac{C_{0jt}}{R_{0jt}^H}. \quad (67)$$

Taking expectations in logs and rearranging, this yields another estimator that over-identifies $\rho(1-\epsilon)$:

$$\widetilde{\rho(1-\epsilon)} \equiv \exp(\log(1 + \theta) + \mathbb{E}[c_{jt} - r_{jt}^H | D_{jt} = 0]). \quad (68)$$

In the $d = 1$ case, multiplying both sides of equation (63) by L_{1jt} implies

$$\rho(1 - \epsilon) \Phi_{jt} p_H (\Phi_{jt} L_{1jt}^\rho - \bar{Q}^G)^{-\epsilon} L_{1jt}^\rho = \chi^{(W)} U_{jt} L_{1jt}^{\theta+1} (1 + \theta) = (1 + \theta) C_{jt}$$

Furthermore, since $(\Phi_{jt} L_{1jt}^\rho - \bar{Q}^G)^{-\epsilon} = (Q_{1jt}^H)^{-\epsilon} = (R_{1jt}^H / p_H)^{\frac{-\epsilon}{1-\epsilon}}$, we can rewrite this expression as

$$\rho(1 - \epsilon) p_H (R_{1jt}^H / p_H)^{\frac{-\epsilon}{1-\epsilon}} \Phi_{jt} L_{1jt}^\rho = (1 + \theta) C_{jt}$$

Taking logs,

$$\log \rho + \log(1 - \epsilon) + \log p_H + \frac{-\epsilon}{1 - \epsilon} r_{1jt}^H - \frac{-\epsilon}{1 - \epsilon} \log p_H + \phi_{jt} + \rho \ell_{1jt} = \log(1 + \theta) + c_{jt}$$

Rearranging, this gives,

$$\underbrace{c_{jt} + \frac{\epsilon}{1-\epsilon} r_{jt}^H}_{\Lambda_{jt}^{\text{CD}}(\epsilon)} = \text{constant} + \phi_{jt} + \rho \ell_{jt}, \quad (69)$$

where $\text{constant} \equiv \log \rho + \log(1 - \epsilon) + \frac{1}{1-\epsilon} \log p_H - \log(1 + \theta)$. Thus, for any candidate

value of ϵ , a regression of $\Lambda_{jt}^{\text{CD}}(\epsilon)$ on ℓ_{jt} controlling for ϕ_{jt} for the winners identifies ρ . Since $\rho(1-\epsilon)$ is identified above, this implies $(1-\epsilon)$ is uniquely determined by this implicit system of equations.

Furthermore, since we showed above that cost-minimization requires $C_{jt} = \frac{\rho}{\beta_L} B_{jt}$, the expected labor share of costs is

$$\frac{\beta_L}{\rho} = \mathbb{E} \left[\frac{B_{jt}}{C_{jt}} \right], \quad (70)$$

so we identify β_L given ρ .

In practice, we simultaneously estimate $(1-\epsilon, \rho, \beta_L)$ by applying equally-weighted GMM to equations (66), (68), (69), and (70).

For the remaining parameters, note that $X_{jt} = \frac{(1+\theta)\beta_M}{\beta_L} B_{jt}$ and $p_K K_{jt} = \frac{(1+\theta)\beta_K}{\beta_L} B_{jt}$, which implies the following expressions:

$$\beta_M = \exp \left(\mathbb{E} [x_{jt} - b_{jt}] - \log \frac{(1+\theta)}{\beta_L} \right), \quad (71)$$

$$\beta_K = \exp \left(\mathbb{E} [\log(p_K K_{jt}) - b_{jt}] - \log \frac{(1+\theta)}{\beta_L} \right), \quad (72)$$

$$\mathbb{E} [u_{jt}] = \mathbb{E} [b_{jt}] - (1+\theta) \mathbb{E} [\ell_{jt}], \quad (73)$$

$$\log p_H = \mathbb{E} [r_{jt}] - \rho(1-\epsilon) \mathbb{E} [\ell_{jt}], \quad (74)$$

where we normalize $\mathbb{E}[\phi_{jt}] = 0$ without loss of generality.

D Expected Impact of an Increase in Market Power

Set up: For simplicity, we consider a production function in which labor is the only input, returns to scale are constant ($\rho = 1$), and firms can only sell output to the private market when deriving theoretical predictions. We focus on firm j at time t , omitting these subscripts without loss of generality, and normalize TFP as $\Phi = 1$. The production function is then $Q = L$. This implies that revenue can be expressed in terms of labor as $R = p_H L^{1-\epsilon}$, so marginal revenue is $\text{MRP} = p_H(1-\epsilon)L^{-\epsilon}$. Since labor is the only input, the marginal cost of production is given by the marginal cost of labor, which is $\text{MCL} = U(1+\theta)L^\theta$. We solve for the baseline equilibrium by equating

MRP and MCL. The baseline equilibrium is characterized by $\bar{L} = \bar{Q} = \left(\frac{p_H}{U} \frac{1-\epsilon}{1+\theta}\right)^{\frac{1}{\theta+\epsilon}}$, $\bar{P} = p_H \bar{L}^{-\epsilon}$, $\bar{R} = p_H \bar{L}^{1-\epsilon}$, and $\bar{W}_{jt} = U \bar{L}^\theta$.

Rotation of labor supply curve: We now consider a compensated rotation of the labor supply curve. In particular, consider an (inverse) labor supply curve $W(L|U', \theta') = U' L^{\theta'}$ for some $\theta' \neq \theta$. This labor supply curve is a “rotation” around the initial equilibrium only if $W(\bar{L}|U', \theta') = \bar{W}$; that is, the baseline labor quantity receives the same wage after the rotation as it did in the baseline equilibrium. This rotation $W(\bar{L}|U', \theta') = \bar{W}$ is solved by $U' = \bar{W} \bar{L}^{-\theta'}$; that is, there is a unique “compensation” $U' - U$ to the location parameter of the labor supply curve such that $W(L|U', \theta')$ is a “rotation” around the initial equilibrium and $\theta' \neq \theta$.

Suppose labor supply is rotated to become more inelastic; that is, $\theta' > \theta$, which also implies $U' < U$. The new equilibrium satisfies $L' = Q' = \left(\frac{p_H}{U'} \frac{1-\epsilon}{1+\theta'}\right)^{\frac{1}{\theta'+\epsilon}} = \bar{L} \left(\frac{1+\theta}{1+\theta'}\right)^{\frac{1}{\theta'+\epsilon}}$. Since $\left(\frac{1+\theta}{1+\theta'}\right)^{\frac{1}{\theta'+\epsilon}} < 1$, then $L' < \bar{L}$ and thereby $Q' < \bar{Q}$. An implication is that $p_H (Q')^{-\epsilon} > p_H (Q)^{-\epsilon}$, so $P' > \bar{P}$. Another implication is that $W' = U' (L')^{\theta'} = U' \left(\bar{L} \left(\frac{1+\theta}{1+\theta'}\right)^{\frac{1}{\theta'+\epsilon}}\right)^{\theta'} = \bar{W} \frac{U'}{U} \left(\frac{1+\theta}{1+\theta'}\right)^{\frac{\theta'}{\theta'+\epsilon}}$. Since $\frac{U'}{U} \left(\frac{1+\theta}{1+\theta'}\right)^{\frac{\theta'}{\theta'+\epsilon}} < 1$, it follows that $W' < \bar{W}$. Therefore, a compensated rotation of the labor supply curve to become less elastic results in reductions in the firm’s employment, wage, and output, as well as an increase in its price.

Rotation of product demand curve: We now consider a compensated rotation of the product demand curve. In particular, consider an (inverse) product demand curve $P(Q|p'_H, \epsilon') = p'_H Q^{-\epsilon'}$ for some $\epsilon' \neq \epsilon$. This product demand curve is a “rotation” around the initial equilibrium only if $P(\bar{Q}|p'_H, \epsilon') = \bar{P}$; that is, the baseline output quantity receives the same price after the rotation as it did in the baseline equilibrium. This rotation $P(\bar{Q}|p'_H, \epsilon') = \bar{P}$ is solved by $p'_H = \bar{P} \bar{Q}^{\epsilon'}$; that is, there is a unique “compensation” $p'_H - p_H$ to the location parameter of the product demand curve such that $P(Q|p'_H, \epsilon')$ is a “rotation” around the initial equilibrium and $\epsilon' \neq \epsilon$.

Suppose product demand is rotated to become more inelastic; that is, $\epsilon' > \epsilon$, which also implies $p'_H > p_H$. The new equilibrium satisfies $L' = Q' = \left(\frac{p'_H}{U} \frac{1-\epsilon'}{1+\theta}\right)^{\frac{1}{\theta+\epsilon'}}$ =

$\bar{L} \left(\frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{1}{\theta+\epsilon'}}$. Since $\left(\frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{1}{\theta+\epsilon'}} < 1$, then $L' < \bar{L}$ and thereby $Q' < \bar{Q}$. An implication is that $U(L')^\theta < U(\bar{L})^\theta$, so $W' < \bar{W}$. Another implication is that $P' = p'_H(Q')^{-\epsilon'} = p'_H \left(\bar{Q} \left(\frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{1}{\theta+\epsilon'}}\right)^{-\epsilon'} = \bar{P} \frac{p'_H}{p_H} \left(\frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{-\epsilon'}{\theta+\epsilon'}}$. Since $\frac{p'_H}{p_H} \left(\frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{-\epsilon'}{\theta+\epsilon'}} > 1$, it follows that $P' > \bar{P}$. Therefore, a compensated rotation of the product demand curve to become less elastic results in reductions in the firm's employment, wage, and output, as well as an increase in its price.

Rotation of both labor supply and product demand curves: Lastly, we consider rotating both the labor supply and product demand curves to become more inelastic; that is, $\epsilon' > \epsilon$ and $\theta' > \theta$. Following the same logic as above, $L' = \bar{L} \left(\frac{1+\theta}{1+\theta'} \frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{1}{\theta'+\epsilon'}}$. Since $\left(\frac{1+\theta}{1+\theta'} \frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{1}{\theta'+\epsilon'}} < 1$, then $L' < \bar{L}$ and thereby $Q' < \bar{Q}$. Since $W' = \bar{W} \frac{U'}{U} \left(\frac{1+\theta}{1+\theta'} \frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{\theta'}{\theta'+\epsilon'}}$ and $\frac{U'}{U} \left(\frac{1+\theta}{1+\theta'} \frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{\theta'}{\theta'+\epsilon'}} < 1$, it follows that $W' < \bar{W}$. Since $P' = \bar{P} \frac{p'_H}{p_H} \left(\frac{1+\theta}{1+\theta'} \frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{-\epsilon'}{\theta'+\epsilon'}}$ and $\frac{p'_H}{p_H} \left(\frac{1+\theta}{1+\theta'} \frac{1-\epsilon'}{1-\epsilon}\right)^{\frac{-\epsilon'}{\theta'+\epsilon'}} > 1$, it follows that $P' > \bar{P}$. Therefore, a simultaneous compensated rotation of both the labor supply and product demand curves to become less elastic results in reductions in the firm's employment, wage, and output, as well as an increase in its price.

Lastly, we show that the impacts of increased market power in one market are attenuated by the existence of market power in the other market. In particular, we show that

$$\frac{\partial^2 L}{\partial \theta' \partial \epsilon'} \Bigg|_{\{P(\bar{Q}|p'_H, \epsilon') = \bar{P}, W(\bar{L}|U', \theta') = \bar{W}\}} = \bar{L} \frac{1}{(\theta + \epsilon)^2} \left[\frac{1}{1 + \theta} + \frac{1}{1 - \epsilon} + \frac{1}{(1 + \theta)(1 - \epsilon)} \right] > 0.$$

Proof. We start with $L = \bar{L} \left[\frac{(1+\theta)(1-\epsilon')}{(1-\epsilon)(1+\theta')}\right]^{\frac{1}{\theta'+\epsilon'}}$, which implies

$$\log L - \log \bar{L} = \frac{1}{\theta' + \epsilon'} \left[\log \frac{(1 + \theta)(1 - \epsilon')}{(1 - \epsilon)(1 + \theta')} \right]$$

Setting $\theta = \theta'$ and $\epsilon = \epsilon'$ delivers $\log L - \log \bar{L} = 0$. We can calculate the following derivatives:

$$\frac{d \log L}{d \theta'} = \frac{1}{L} \frac{dL}{d \theta'} = -\frac{1}{(\theta' + \epsilon')^2} \left[\log \frac{(1 + \theta)(1 - \epsilon')}{(1 - \epsilon)(1 + \theta')} \right] - \frac{1}{\theta' + \epsilon'} \frac{1}{1 + \theta'}$$

$$\frac{d \log L}{d \epsilon'} = \frac{1}{L} \frac{dL}{d \epsilon'} = -\frac{1}{(\theta' + \epsilon')^2} \left[\log \frac{(1 + \theta)(1 - \epsilon')}{(1 - \epsilon)(1 + \theta')} \right] - \frac{1}{\theta' + \epsilon'} \frac{1}{1 - \epsilon'}$$

Substituting,

$$\frac{dL}{d \theta'} = - \left[\frac{1}{(\theta' + \epsilon')^2} \left[\log \frac{(1 + \theta)(1 - \epsilon')}{(1 - \epsilon)(1 + \theta')} \right] + \frac{1}{\theta' + \epsilon'} \frac{1}{1 + \theta'} \right] L$$

$$\frac{dL}{d \theta'} = -\frac{1}{\theta' + \epsilon'} \left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] L$$

$$\frac{dL}{d \epsilon'} = -\frac{1}{\theta' + \epsilon'} \left[\log L - \log \bar{L} + \frac{1}{1 - \epsilon'} \right] L$$

Thus,

$$\begin{aligned} \frac{d^2 L}{d \theta' d \epsilon'} &= \frac{1}{(\theta' + \epsilon')^2} \left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] L - \frac{1}{\theta' + \epsilon'} L \frac{d \log L}{d \epsilon'} - \frac{dL}{d \epsilon'} \frac{1}{\theta' + \epsilon'} \left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] \\ &= \frac{1}{(\theta' + \epsilon')^2} \left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] L + \frac{1}{\theta' + \epsilon'} \left[\frac{1}{\theta' + \epsilon'} \left[\log L - \log \bar{L} + \frac{1}{1 - \epsilon'} \right] \right] L \\ &\quad + \frac{1}{\theta' + \epsilon'} \left[\log L - \log \bar{L} + \frac{1}{1 - \epsilon'} \right] \frac{1}{\theta' + \epsilon'} \left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] L \\ &= \frac{1}{(\theta' + \epsilon')^2} \left(\left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] L + \left[\log L - \log \bar{L} + \frac{1}{1 - \epsilon'} \right] L \right. \\ &\quad \left. + \left[\log L - \log \bar{L} + \frac{1}{1 - \epsilon'} \right] \left[\log L - \log \bar{L} + \frac{1}{1 + \theta'} \right] L \right) \end{aligned}$$

Finally, evaluating at $L = \bar{L}$, $\theta = \theta'$, and $\epsilon = \epsilon'$ delivers:

$$\frac{d^2 L}{d \theta' d \epsilon'} = \bar{L} \frac{1}{(\theta' + \epsilon')^2} \left[\frac{1}{1 + \theta'} + \frac{1}{1 - \epsilon'} + \frac{1}{1 - \epsilon'} \frac{1}{1 + \theta'} \right] > 0.$$

□

E Additional Institutional Details on the Construction Industry and Procurement Auctions

E.1 Prevalence of Non-wage Compensation

The Economic Census of the Construction Sector (EC), which is collected every five years by the US Census Bureau, provides informative descriptive statistics on the wage and non-wage compensation paid by the US construction industry. The EC provides measures of non-wage compensation separately for legally-required fringe benefits (social security, unemployment insurance, worker injury-compensation insurance, etc.) and voluntary fringe benefits (health insurance, pensions, training, etc.), as well as total payroll. We analyze the EC data for the construction industry, aggregated to the state-level at five-year frequency, for the period from 1977 to 2017.

In Online Appendix Figure A.10, we present the share of total compensation from wages, legally-required fringe benefits, and voluntary fringe benefits over time. In 2012, which is near the end of the time frame considered in our main analysis, we find that about 10% of total compensation is due to legally-required benefits and about 10% is due to voluntary benefits, with wages accounting for the remaining 80%. Thus, voluntary benefits – which is the component of non-wage compensation that the firm can in principle adjust – accounts for only one-tenth of total compensation in the construction industry.

It is perhaps not surprising that the construction industry primarily compensates workers through wages rather than voluntary benefits, as these benefits may be costly to provide. Industry experts summarize the role of wages versus non-wage benefits in recruiting workers to construction jobs by noting that “base pay has been the single most important piece of compensation” and “benefits aren’t usually a driving factor in recruiting.”¹ However, some workers may be offered non-wage benefits that are proportional to wages, such as bonuses and incentives which are offered as a

¹See “What you need to know about compensation in construction” by Carpenter-Beck, <https://www.sage.com/en-us/blog/need-know-compensation-construction/>.

percentage of base salary and thus adjust when wages increase.² We show in Online Appendix H that proportional adjustments in non-wage benefits do not introduce bias in our estimation of the labor supply curve.

E.2 Relevance of Prevailing Wage Laws

Prevailing wage laws require that workers employed by private construction firms on government-funded construction projects be paid at least the wages and benefits paid to similar workers in the same location where the project is located. The Davis-Bacon Act was passed by Congress in 1931 to require that private construction firms pay the prevailing wage to their employees on all federally-funded construction projects. Subsequently, most state governments passed so-called “little Davis-Bacon” laws to extend prevailing wage requirements to state-funded construction projects. However, 15 of those states have since repealed their prevailing wage laws.³

One potential concern in our estimation of the labor supply elasticity is that first-time procurement auction winners may become subject to the prevailing wage for the first time, which could force them to increase the wages of incumbent workers, independently of whether or not they hire new workers. In order for such an effect to occur, three conditions would need to be satisfied. First, the firm must have initial wages below the prevailing wage. This is unlikely to be true for many firms in our sample, as we find that procurement auction participants (both winners and losers) have higher than average wages in the pre-period. Second, even if the winning firms had initial wages below the prevailing wage, prevailing wage laws would only bind if the new wage after winning the procurement contract would not have met the prevailing wage in the absence of a prevailing wage law. In the presence of an upward-sloping labor supply curve, the wage increase required to recruit new workers may reach the prevailing wage even among winners that did not initially pay the prevailing

²See “Competitive pay to recruit and retain employees” by Robinson, <https://www.sage.com/en-us/blog/competitive-pay-to-recruit-and-retain-employees/>.

³The list of states that currently have prevailing wage laws as well as the history of repeals is provided by the US Department of Labor here: <https://www.dol.gov/agencies/whd/state/prevailing-wages>.

wage. Third, the procurement contract must be funded by a state government that currently has a prevailing wage law, or be funded by the federal government.

To investigate if there are actually effects of prevailing wage laws, we use repeals of state prevailing wage laws in a difference-in-differences analysis at the state-level to examine how wage and non-wage compensation are impacted by prevailing wages. For outcome measures, we use the EC data described above on wages and non-wage fringe benefits in the construction sector. In Online Appendix Table A.8, the difference-in-differences estimates for repeals suggest that prevailing wage laws have little to no effect on total compensation, wages, non-wage fringe benefits, or the share of total compensation from non-wage fringe benefits.⁴ For example, the effect of a repeal on the log wage in the construction industry is 0.009 with standard error 0.029, and the effect on log non-wage fringe benefits in the construction industry is 0.015 with standard error 0.031.

E.3 Safety Regulations and Procurement Auctions

The construction industry is governed by extensive safety regulations. For example, construction employers in California must comply with Cal/OSHA regulations found in the following subchapters of California Code of Regulations, title 8, chapter 4: subchapter 4 (Construction Safety Orders); subchapter 5 (Electrical Safety Orders); and subchapter 7 (General Industry Safety Orders).⁵ These regulations are typically task-specific, e.g., California requires certification and various safety procedures to be followed by crane operators.

It is important to observe that these safety regulations apply to all firms in the construction industry. In particular, government procurement projects are governed by the same safety regulations as private projects. Thus, receiving a procurement

⁴Prior work, reviewed by [Duncan and Ormiston \(2019\)](#), has found little evidence that prevailing wage laws increase wages in the construction industry. In their study of the first 9 repeals of prevailing wage laws and outcomes only through 1993, [Kessler and Katz \(2001\)](#) find economically small impacts on wages that become statistically insignificant when controlling for pre-trends. Our analysis includes more recent repeals, effects on non-wage outcomes, and uses modern difference-in-differences estimators for staggered treatment contexts ([Callaway and Sant'Anna, 2020](#)).

⁵See <https://www.dir.ca.gov/dosh/construction-guide-summary.html>.

contract does not change the safety regulations governing the construction firm. Similarly, we have read the annual reports of public construction firms to examine the language they use to describe participation in government projects.⁶ While they discuss the costs and opportunities associated with government contracts, we find no mention of changing safety policies in consideration of procurements.⁷

While we find no legal requirement that safety must be improved in response to winning a procurement auction, we may still worry that winning firms make safety improvements. In order to check for such effects, we download publicly-available data from OSHA safety inspections and link it to our procurement auction records. Since each state provides OSHA data in a different format, we focused on the largest state in our sample, California. Given our linked dataset between procurement auction and OSHA records, we used these data to run the same regression specification as our baseline research design in equation (21), but now using safety investigations and violations as the outcome variables.

Reassuringly, we find fairly precisely estimated zero effects on the probability of a safety violation and on the probability of a safety investigation: As shown in Online Appendix Table A.7, the point estimates are 0.000 and 0.009, with standard errors 0.006 and 0.008 respectively, for the safety violation and investigation probabilities. These estimates suggest little if any impacts. By comparison, the probability of a safety violation is 4.1% and the probability of an investigation is 7.5% in the year that a firm is a bidder (both winners and losers) in an auction.

E.4 Prevalence of Auctions with a Quality Dimension

Lewis and Bajari (2011) study a special type of auction, called an A+B auction, in which firms bid both a price and a time-to-completion. A bid submission contains

⁶We accessed annual reports using <https://www.annualreports.com>. E.g., the annual report for Granite Construction Inc, one of the the largest auction winners in the procurement auction records for California, is available at https://www.annualreports.com/HostedData/AnnualReports/PDF/NYSE_GVA_2022.pdf.

⁷Although receiving a procurement contract does not lead to different safety regulations, a history of serious workplace safety violations may be a reason to deny a firm the chance to participate in a procurement auction at the pre-qualification phase.

a dollar amount, the “A” component, and a total number of days to complete the project, the “B” component. The score is a weighted sum of these two components, and the bidder with the lowest score wins. The A+B design provides an incentive for the bidder to disrupt traffic for a shorter period of time, and it is rarely used outside a few categories of construction projects that require road lane or shoulder closure.

The empirical context considered by [Lewis and Bajari \(2011\)](#) is California DOT procurement auctions during 2003-2008. They restrict the sample of auctions to the small set of project types that are more likely to use the A+B format. They also restrict the analysis to 5 districts (4, 6, 8, 11, and 12) that use the A+B format more frequently, with most auctions coming from the San Francisco Bay Area. Even with such selective sampling, among 708 auctions that fit the criteria, only 80 have the A+B format. The contracts auctioned through the A+B format are exceptionally large: the average engineer’s estimate is \$21.9 million for A+B auctions versus \$4.6 million for other auctions.

To examine the prevalence of the A+B design in our analysis sample, we identify the list of A+B auctions in our 2001-2015 California DOT procurement auction data. In this sample, only 4.2% of auctions use the A+B format. We also check how these auctions enter our regressions. Recall that we only use auctions with a first-time winner to implement our estimation of the labor supply curve. In all of California during our sample period, only 12 auctions use the A+B format and are won by a first-time winner. Thus, A+B auctions comprise a very small share of our estimation sample for California.

Among the few auctions that use the A+B format in our sample, one may worry that the ultimate payment differs significantly from the total bid. According to [Lewis and Bajari \(2011\)](#), “standard contracts typically finish 7% early, whereas A+B contracts finish exactly on time.” More specifically, in their sample, 52% of the A+B contracts are completed exactly on time. Completing on time means that the payment equals the part A bid. Using the A+B auctions that satisfy their sample time frame and districts, we find that about 87% of the total bids are determined by the A price component rather than the B time-to-completion component, suggesting that the price-only auction might remain a good approximation. We find similar statistics

using the A+B auctions that fit our analysis sample criteria.

Another paper discussing auctions with a quality dimension is [Takahashi \(2018\)](#). He considers so-called design-bid auctions, in which each bidder submits a design and a price bid, and the price-per-quality score ratio determines the winner. In his sample from the Florida DOT between 2000 and 2011, only 152 auctions are design-bid auctions. In this case, quality-and-price-based auctions are again exceptions.

E.5 Prevalence of Subcontracting

Along with the project announcement, DOT publishes a detailed project description, including an engineer’s estimate of the size and cost of each item. The winners of procurement auctions may subcontract some of those project items to be completed by other firms. However, data on subcontracting is limited, so there has been little empirical research on subcontracting.

The California DOT requires that bidders settle all subcontracting agreements prior to bid submission. Bidders must list all project items that will be subcontracted, along with the prices that will be paid to associated subcontractors. While subcontracting agreements are not available in a readily-usable data form, [Balat et al. \(2017\)](#) digitized this information from bidding documents submitted to the California DOT for projects auctioned between 2002 and 2016. Because their sample overlaps well with ours, their summary statistics represent ours well.

Using their data on subcontracting in California procurement auctions, [Balat et al. \(2017\)](#) find that subcontracting is common, with about 95% of auction winners subcontracting at least one item in road or bridge projects. However, the amount of subcontracting relative to the total value of the project is small. They find that, for the average (median) winner of a procurement contract, subcontractors only account for 8% (4%) of the bid value of the project. Thus, for the largest state in our data, subcontracting accounts for a relatively small share of the bid value.

F Details on Labor Supply Elasticity Estimators

F.1 Identification using the LMS Estimator

Following [Lamadon et al. \(2022, LMS\)](#), we consider instrumenting for long-differences in log labor using short-differences in log value added (VA). Denoting the short-difference in log VA by $\Delta va_{jt} \equiv \log VA_{jt} - \log VA_{jt-1}$, the estimator of LMS is,

$$\hat{\theta}_{\Delta va} \equiv \frac{\text{Cov}[w_{jt+e} - w_{jt-e'}, \Delta va_{jt}]}{\text{Cov}[\ell_{jt+e} - \ell_{jt-e'}, \Delta va_{jt}]}.$$

Unlike the estimators of Propositions 3–4, this estimator does not make use of information on procurement auctions and thus can be applied to the entire construction industry rather than only construction firms that bid for procurement projects.

Consider the following assumption, which compares short-run changes in VA to longer-run changes in TFP and firm-specific amenity shocks:

Assumption 1. *Suppose $\exists e, e' > 0$ sufficiently large such that (i) $\phi_{jt+e} - \phi_{jt-e'}$ is correlated with Δva_{jt} , and (ii) Δva_{jt} is orthogonal to $\nu_{jt+e} - \nu_{jt-e'}$.*

We then have the following result:

Proposition 6. *Under Assumption 1 and the rank condition $\text{Cov}[(\ell_{jt+e} - \ell_{jt-e'}), \Delta va_{jt}] \neq 0$, $\hat{\theta}_{\Delta va}$ recovers θ .*

Proof. By equation (4),

$$\hat{\theta}_{\Delta va} = \frac{\text{Cov}[\theta(\ell_{jt+e} - \ell_{jt-e'}), \Delta va_{jt}]}{\underbrace{\text{Cov}[(\ell_{jt+e} - \ell_{jt-e'}), \Delta va_{jt}]}_{= \theta}} + \frac{\text{Cov}[(\nu_{jt+e} - \nu_{jt-e'}), \Delta va_{jt}]}{\underbrace{\text{Cov}[(\ell_{jt+e} - \ell_{jt-e'}), \Delta va_{jt}]}_{= 0}} = \theta,$$

where the denominator of each term is non-zero (i.e., the rank condition is satisfied) by Assumption 1(i) and the second term is zero by Assumption 1(ii). \square

The key result, the exclusion condition $\text{Cov}[(\nu_{jt+e} - \nu_{jt-e'}), \Delta va_{jt}] = 0$, relies on the assumption that firm-specific amenity shocks are transitory while TFP shocks are persistent. Thus, short-run VA growth is correlated with long-run employment growth

(satisfying the rank condition due to persistence in TFP shocks), but orthogonal to long-run firm-specific amenity shocks. In practice, we must take a stand on the persistence of the transitory shocks. LMS argue that these transitory shocks are well-approximated as a moving average of order one, in which case, Assumption 1 holds as long as $e \geq 2, e' \geq 3$ and TFP shocks persist for at least e periods. We use the same choices of e and e' as LMS in our empirical implementation.

F.2 Implementation of RDD Estimators

We now describe the implementation of the RDD estimation defined in Proposition 4. Recall that, for a firm j that bids in auction ι at time t , we define the loss margin as $\tau_{jt} \equiv \frac{Z_{jt} - Z_{\iota}^*}{Z_{\iota}^*}$, where Z_{ι}^* is the winning bid in auction ι . We can interpret τ_{jt} as a measure of proximity to the discontinuity, satisfying $\tau_{jt} = 0$ for auction winners ($D_{jt} = 1$) and $\tau_{jt} > 0$ for auction losers ($D_{jt} = 0$). The regression specification is,

$$\begin{aligned}
 1 \{ \tau_{jt} \leq \bar{\tau} \} w_{jt+e} = & \underbrace{\sum_{e' \neq \bar{e}} 1 \{ \tau_{jt} \leq \bar{\tau} \} 1 \{ e' = e \} \mu_{te'}^{\bar{\tau}}}_{\text{event time fixed effect}} + \underbrace{\sum_{j'} \sum_{\iota'} 1 \{ \tau_{jt} \leq \bar{\tau} \} 1 \{ j' = j \text{ and } \iota' = \iota \} \psi_{j'\iota't}^{\bar{\tau}}}_{\text{firm-auction fixed effect}} \\
 & + \underbrace{\sum_{e' \neq \bar{e}} 1 \{ \tau_{jt} \leq \bar{\tau} \} 1 \{ e' = e \} D_{jt} \lambda_{te'}^{\bar{\tau}}}_{\text{treatment status by event time}} + \underbrace{1 \{ \tau_{jt} \leq \bar{\tau} \} \epsilon_{jte}}_{\text{residual}}
 \end{aligned} \tag{75}$$

where we have fully interacted the regression with the indicator $1 \{ \tau_{jt} \leq \bar{\tau} \}$ to remove any control units that do not place close bids to the winning firms. The parameter $\lambda_{te}^{\bar{\tau}}$ recovers the numerator of $\theta_{\text{RDD}}(\bar{\tau})$ for a particular choice of $\bar{\tau}$, pair (e, t) , and $\bar{e} = -2$ is the omitted event time. As in the baseline implementation, we estimate $\lambda_{te}^{\bar{\tau}}$ for all t and e and then average across t , using the delta method to compute standard errors. The analogous regression in which ℓ_{jt+e} is the outcome recovers the denominator of $\theta_{\text{RDD}}(\bar{\tau})$. We average across event times e to form the main estimate.

As an alternative implementation of the RDD estimator defined in Proposition 4, consider a regression that controls for the loss margin, rather than restricting the

sample based on the loss margin, as,

$$\begin{aligned}
w_{jt+e} = & \underbrace{\sum_{e' \neq \bar{e}} 1 \{e' = e\} \mu_{te'}}_{\text{event time fixed effect}} + \underbrace{\sum_{j'} \sum_{\iota'} 1 \{j' = j \text{ and } \iota' = \iota\} \psi_{j'\iota't}}_{\text{firm-auction fixed effect}} & (76) \\
& + \underbrace{\sum_{e' \neq \bar{e}} 1 \{e' = e\} D_{jt} \lambda_{te'}}_{\text{treatment status by event time}} + \underbrace{\sum_{e' \neq \bar{e}} 1 \{e' = e\} \varrho_e(\tau_{jt})}_{\text{polynomial in loss margin}} + \underbrace{\epsilon_{jte}}_{\text{residual}},
\end{aligned}$$

where $\varrho_e(\tau_{jt})$ is an event time-specific polynomial in the loss margin τ_{jt} . In practice, we consider a linear specification, $\varrho_e(\tau_{jt}) = \varphi_e \tau_{jt}$, and a third-order polynomial specification, $\varrho_e(\tau_{jt}) = \varphi_{1e} \tau_{jt} + \varphi_{2e} \tau_{jt}^2 + \varphi_{3e} \tau_{jt}^3$.

F.3 Implementation of Local Labor Market Estimators

Let m denote the market in which the firm participates, and let \mathcal{J}_m denote the set of firms that participate in market m . Possible markets include the auction in which the firm bids or the commuting zone in which the firm employs workers. Our goal is to estimate θ_{DiD} while controlling for market-specific shocks. To do so, we will implement the Proposition 3 estimator separately by market, allowing each market to experience its own sequence of event time effects.

Consider the cohort of firms that receive a procurement contract in year t ($D_{jt} = 1$) and the set of comparison firms that bid for a procurement in year t but lose ($D_{jt} = 0$). Let e denote an event time relative to t . For each event time $e = -4, \dots, 4$, our DiD estimation for market m is implemented as

$$\begin{aligned}
w_{jt+e} = & \underbrace{\sum_{e' \neq \bar{e}} 1 \{e' = e\} \mu_{te'}^m}_{\text{market-specific event time fixed effect}} + \underbrace{\sum_{j'} 1 \{j' = j\} \psi_{j't}}_{\text{firm fixed effect}} & (77) \\
& + \underbrace{\sum_{e' \neq \bar{e}} 1 \{e' = e\} D_{jt} \lambda_{te'}^m}_{\text{treatment status by event time}} + \underbrace{\epsilon_{jte}}_{\text{residual}} \quad \text{among } j \in \mathcal{J}_m
\end{aligned}$$

where, for market m , λ_{te}^m recovers the numerator of θ_{DiD}^m for a particular pair (e, t) and $\bar{e} = -2$ is the omitted event time. This estimator differs from the baseline specification in that it only considers comparison firms that participate in the same market as the firm that receives a procurement contract, ensuring that market-specific shocks are controlled. In particular, market-specific shocks are captured by the μ_{te}^m parameters. The analogous regression in which ℓ_{jt+e} is the outcome recovers the denominator of θ_{DiD}^m . We average across event times e to form the main estimate. Finally, we average θ_{DiD}^m across all of the markets m to form the overall θ_{DiD} estimate that controls flexibly for market-specific shocks.

G Additional Tables and Figures

State	DOT Auction Records		Final Sample: Matched Auction-Tax Data		
	Data Source	Includes EIN	Bidders in 2010 (Num. Firms)	Share of 2010 Construction Sector: Value Added	FTE Workers
AL	State Website	✗	196	15.7%	17.4%
AR	State Website	✗	149	7.9%	12.8%
AZ	No	✗	*	*	*
CA	State Website	✗	1,041	8.3%	11.2%
CO	FOIA Request	✓	241	12.6%	14.7%
CT	FOIA Request	✗	126	9.4%	15.5%
FL	State Website	✓	344	30.7%	10.6%
GA	BidX Website	✗	137	4.3%	7.0%
IA	BidX Website	✗	256	15.4%	20.7%
ID	BidX Website	✗	112	17.2%	13.6%
IL	No	✗	*	*	*
IN	State Website	✓	213	10.6%	16.6%
KS	BidX Website	✓	130	13.7%	21.6%
KY	No	✗	*	*	*
LA	BidX Website	✗	167	11.5%	10.8%
MA	No	✗	*	*	*
MD	No	✗	*	*	*
ME	BidX Website	✗	141	13.7%	16.9%
MI	BidX Website	✗	391	9.5%	16.3%
MN	BidX Website	✗	262	13.5%	19.8%
MO	BidX Website	✗	179	14.9%	13.3%
MS	No	✗	*	*	*
MT	FOIA Request	✗	122	15.0%	23.6%
NC	BidX Website	✗	135	5.2%	9.8%
ND	FOIA Request	✗	*	*	*
NE	No	✗	*	*	*
NH	No	✗	*	*	*
NJ	No	✗	*	*	*
NM	BidX Website	✗	*	*	*
NV	No	✗	*	*	*
NY	No	✗	*	*	*
OH	BidX Website	✗	320	43.7%	17.5%
OK	No	✗	*	*	*
OR	No	✗	*	*	*
PA	No	✗	*	*	*
SC	No	✗	*	*	*
SD	No	✗	*	*	*
TN	BidX Website	✗	140	5.3%	11.5%
TX	FOIA Request	✓	551	4.9%	9.6%
UT	No	✗	*	*	*
VA	BidX Website	✗	241	14.2%	12.0%
VT	BidX Website	✗	*	*	*
WA	BidX Website	✗	200	7.5%	14.0%
WI	BidX Website	✗	194	12.1%	14.6%
WV	BidX and State Websites	✓	103	13.7%	19.0%
National			6,792	10.7%	9.9%

Table A.1: Summary of Auction Data by State

Notes: The first two columns provide information on in-state DOT data sources by state, where “state” refers to the state in which the auction occurred. The first column indicates the source from which we obtained data on that state’s DOT auctions, and the second column indicates whether or not EINs were included in the auction records. The final three columns provide information on the final sample of firms in the matched auction-tax data, where “state” refers to the state in which the firm filed taxes. Among firms in the construction industry in 2010, the last two columns consider the share of value added and FTE workers due to the firms that participated in auctions in our sample. We drop from these calculations firms that have missing values on the variables displayed, so the total sample size must be smaller than in Online Appendix Table A.2. An asterisk (*) denotes that number of bidders is non-zero but below the disclosure threshold.

	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 A.2: Sample Characteristics

Notes: This table displays descriptive statistics for the sample of firms that place bids in 2010. The third column compares aggregates for this sample to all firms in the construction industry in the 2010 tax records.

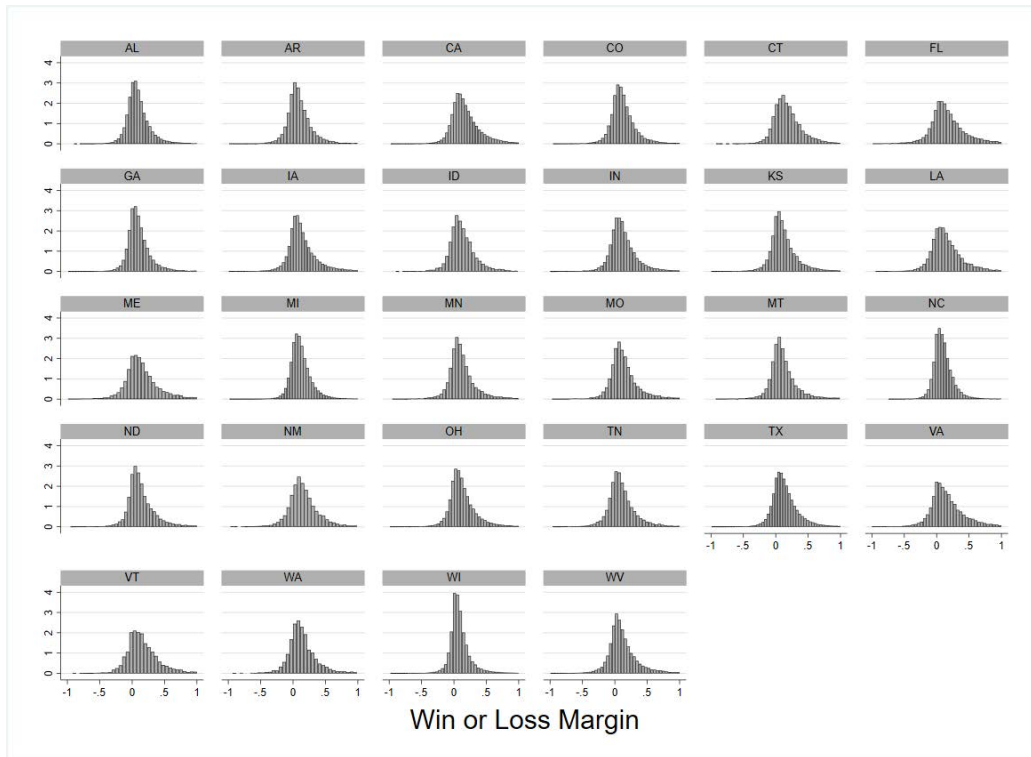


Figure A.1: Chassang et al. (2022) Visual Test for Collusion

Notes: This figure displays the histogram of bid competition for each of the 28 states in our sample. Negative values indicate the difference between the winner’s bid and the bid of the runner-up. Positive values indicate the difference between each loser’s bid and the winner’s bid. Differences are scaled by the winner’s bid in each case. Chassang et al. (2022) demonstrate that, under some assumptions on the auction environment, these differences should display discontinuities in the histogram near zero if there is collusion.

	Effect on Employment		Effect on Earnings		Implications for Labor Market		
	Estimate	Std. Err.	Estimate	Std. Err.	Parameter θ	Elasticity $1/\theta$	Markdown $(1 + \theta)^{-1}$
Panel A) By Proximity:							
Any Proximity	0.083	(0.019)	0.020	(0.008)	0.245	4.084	0.803
Proximity 1.0	0.083	(0.019)	0.021	(0.008)	0.251	3.991	0.800
Proximity 0.5	0.080	(0.019)	0.020	(0.008)	0.251	3.980	0.799
Proximity 0.4	0.079	(0.020)	0.022	(0.008)	0.277	3.608	0.783
Proximity 0.3	0.079	(0.020)	0.022	(0.008)	0.281	3.559	0.781
Proximity 0.2	0.079	(0.021)	0.020	(0.009)	0.257	3.892	0.796
Proximity 0.1	0.065	(0.025)	0.019	(0.010)	0.286	3.491	0.777
Panel B) By Proximity for Stayers:							
Any Proximity	0.083	(0.019)	0.023	(0.006)	0.278	3.600	0.783
Proximity 1.0	0.083	(0.019)	0.024	(0.006)	0.283	3.530	0.779
Proximity 0.5	0.080	(0.019)	0.022	(0.006)	0.277	3.605	0.783
Proximity 0.4	0.079	(0.020)	0.023	(0.006)	0.294	3.403	0.773
Proximity 0.3	0.079	(0.020)	0.023	(0.006)	0.288	3.467	0.776
Proximity 0.2	0.079	(0.021)	0.021	(0.007)	0.271	3.689	0.787
Proximity 0.1	0.065	(0.025)	0.019	(0.007)	0.286	3.499	0.778
Panel C) By Worker Incumbency:							
Stayer Spell: $(-1, \dots, 1)$	0.083	(0.019)	0.023	(0.005)	0.272	3.681	0.786
Stayer Spell: $(-2, \dots, 2)$	0.083	(0.019)	0.023	(0.006)	0.278	3.600	0.783
Stayer Spell: $(-3, \dots, 3)$	0.083	(0.019)	0.021	(0.007)	0.253	3.957	0.798
Tenure: 1 Year	0.083	(0.019)	0.023	(0.006)	0.277	3.615	0.783
Tenure: 2 Years	0.083	(0.019)	0.023	(0.006)	0.277	3.615	0.783
Tenure: 3 Years	0.083	(0.019)	0.025	(0.006)	0.301	3.326	0.769
Tenure: 4 Years	0.083	(0.019)	0.022	(0.006)	0.266	3.766	0.790
Panel D) By Employment Intensity:							
Add Indep. Contractors	0.092	(0.021)	0.020	(0.009)	0.220	4.548	0.820
110% of FTE Wage	0.083	(0.019)	0.023	(0.006)	0.276	3.627	0.784
120% of FTE Wage	0.083	(0.019)	0.022	(0.006)	0.266	3.755	0.790
130% of FTE Wage	0.083	(0.019)	0.022	(0.006)	0.264	3.788	0.791
140% of FTE Wage	0.083	(0.019)	0.021	(0.006)	0.256	3.909	0.796
150% of FTE Wage	0.083	(0.019)	0.019	(0.006)	0.230	4.346	0.813
Panel E) Interacted DiD Designs:							
Fully-interacted: CZ	0.074	(0.019)	0.017	(0.009)	0.224	4.467	0.817
Fully-interacted: Auction	0.088	(0.027)	0.018	(0.009)	0.206	4.846	0.829
Panel F) LMS Passthrough Regressions:							
Full Construction Industry	0.087		0.023		0.266	3.758	0.790
Full Construction Industry, Control CZ	0.092		0.027		0.299	3.349	0.770

Table A.3: Specifications for Estimating the Parameter θ

Notes: This table presents estimates of the impacts of winning a procurement contract on employment and earnings per worker in the post-treatment time period. Employment and earnings per worker are measured in log units. Proximity refers to the largest value of τ_{jt} permitted in the sample. Specification details and sample definitions are provided in the main text.

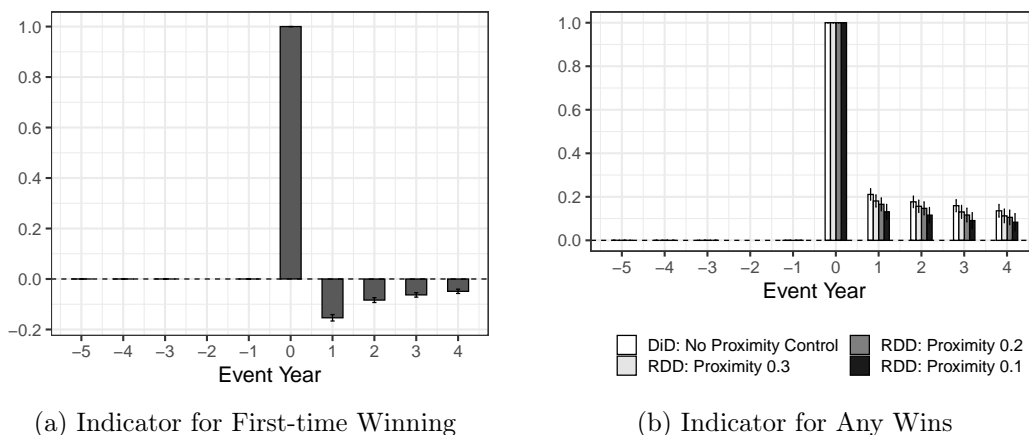


Figure A.2: Visualizing the Research Design

Notes: This figure presents estimates based on equation (21) of the impacts of winning a procurement contract. The omitted event time is -2 . The outcomes considered are the probability of winning an auction for the first time (subfigure a) and the probability of winning any auction in the current year (subfigure b). It provides these estimates separately by event year. Proximity refers to the largest value of τ_{jt} permitted in the sample. 95% confidence intervals are displayed in brackets. Specification details and sample definitions are provided in the main text.

Design:	DiD	RDD		
	Any	0.3	0.2	0.1
Impact: Before Treatment	-0.003 (0.013)	0.003 (0.019)	0.003 (0.020)	-0.007 (0.024)
Impact: After Treatment	0.010 (0.011)	0.017 (0.017)	0.016 (0.018)	0.009 (0.021)

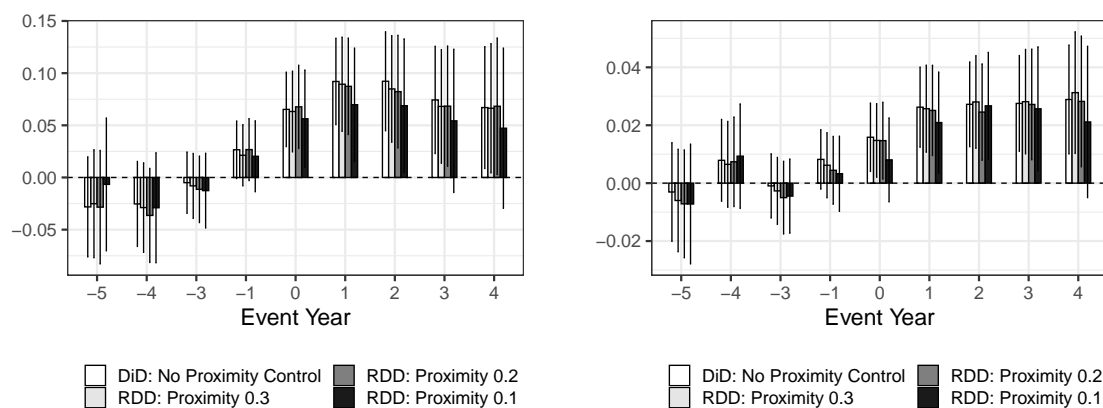
Table A.4: Impacts of Winning a Procurement Contract: Earnings of New Hires

Notes: This table presents estimates of the impacts of winning a procurement contract on the log earnings per worker at the new employer among new hires in the post-treatment and pre-treatment time periods. Proximity refers to the largest value of τ_{jt} permitted in the sample. Specification details and sample definitions are provided in the main text.

Design:	DiD		RDD	
	Any	0.3	0.2	0.1
Impact: Before Treatment	-0.002 (0.019)	-0.009 (0.027)	-0.003 (0.029)	0.004 (0.034)
Impact: After Treatment	0.009 (0.016)	0.012 (0.024)	0.014 (0.026)	0.023 (0.030)

Table A.5: Impacts of Winning a Procurement Contract: Quality of New Hires

Notes: This table presents estimates of the impacts of winning a procurement contract on worker quality. Worker quality is measured by log earnings per worker in the previous firm. Proximity refers to the largest value of τ_{jt} permitted in the sample. Specification details and sample definitions are provided in the main text.




(a) Outcome: Log Number of Employees


(b) Outcome: Log Earnings per Employee

Figure A.3: DiD and RDD Estimates of the Effects of Winning a Procurement

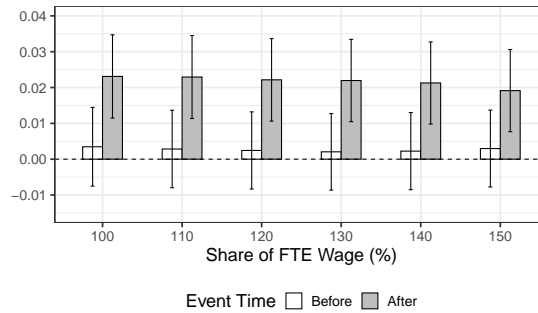
Notes: This figure displays estimates of the effects of winning a procurement contract on winners relative to losers across event times. The winner is announced during event time 0 and outcomes are normalized to zero for both winners and losers in event time -2. All specifications control for time-invariant auction and firm characteristics as well as time fixed effects. The outcomes are log employment in subfigure a and log earnings per employee in subfigure b. Proximity refers to the largest value of τ_{jt} permitted in the sample. 95% confidence intervals are displayed in brackets.



(a) Stayer Definitions



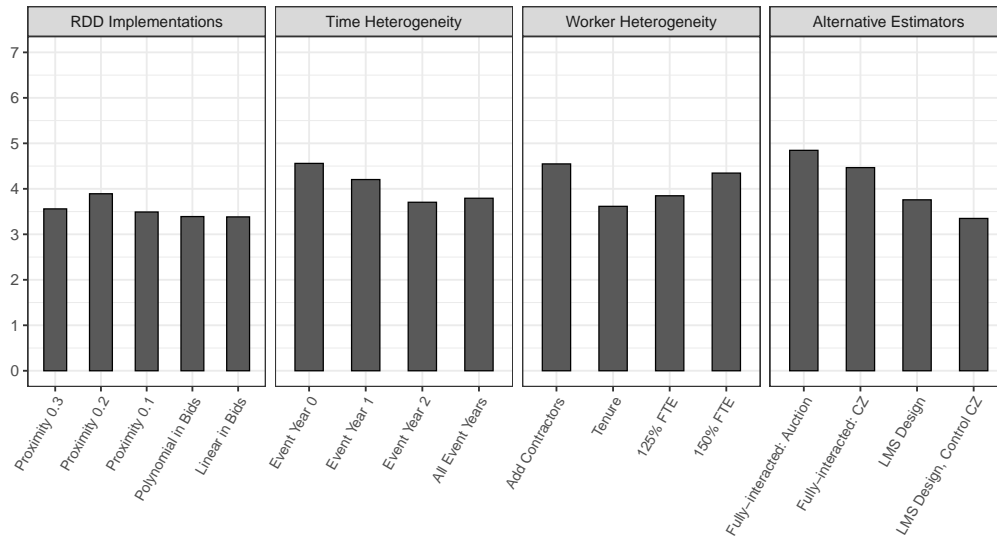
(b) Tenure Definitions




(c) FTE Definitions

Figure A.4: Impacts on log Earnings per Worker: Sensitivity to Worker Sample

Notes: This figure presents estimates based on equation (21) of the impacts of winning a procurement contract. The outcome considered is log earnings per worker. It provides these estimates for alternative sample definitions for stayers (subfigure a), tenured workers (subfigure b), and full-time equivalence (FTE) thresholds as a percentage of the annualized minimum wage (subfigure c). 95% confidence intervals are displayed in brackets. Specification details and sample definitions are provided in the main text.




(a) Labor Supply Elasticity, $1/\theta$



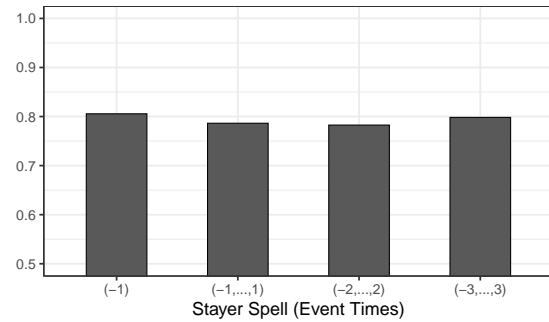
(b) Wage Markdown relative to MRPL, $(1 + \theta)^{-1}$

Figure A.5: Specification Checks: Estimates of the Labor Supply Elasticity and Wage Markdown

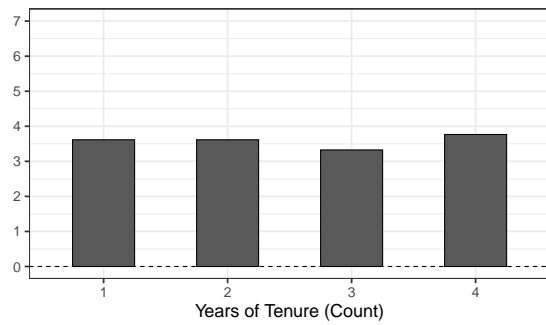
Notes: This figure presents the sensitivity checks for the estimates of the labor supply elasticity, $1/\theta$, and the wage markdown relative to MRPL, $(1 + \theta)^{-1}$. Specification details and sample definitions are provided in the main text.



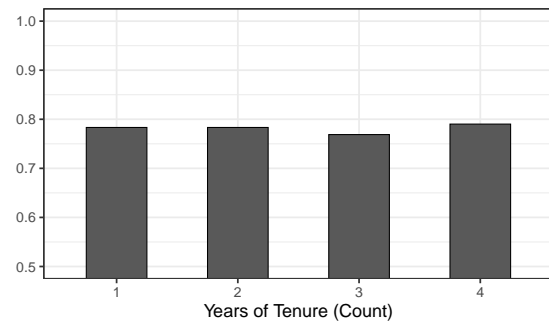
(a) Stayer Definitions: Labor Supply Elasticity



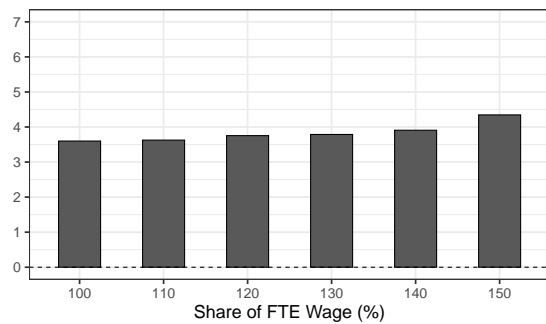
(b) Stayer Definitions: Wage Markdown



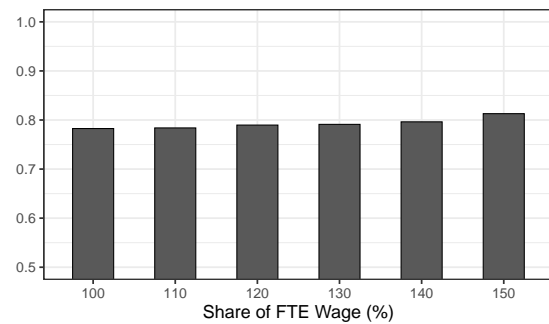
(c) Tenure Definitions: Labor Supply Elasticity



(d) Tenure Definitions: Wage Markdown



(e) FTE Definitions: Labor Supply Elasticity



(f) FTE Definitions: Wage Markdown

Figure A.6: Labor Supply Elasticity and Wage Markdown: Sensitivity to Worker Sample

Notes: This figure presents estimates of the labor supply elasticity, $1/\theta$, and the wage markdown relative to MRPL, $(1 + \theta)^{-1}$ for alternative sample definitions for stayers (subfigure a-b), tenured workers (subfigure c-d), and full-time equivalence (FTE) thresholds as a percentage of the annualized minimum wage (subfigure e-f). Specification details and sample definitions are provided in the main text.

	All States		Prevailing Wage States	
	All Workers	Stayers	All Workers	Stayers
Impacts of Winning an Auction:				
Log Employment:	0.083 (0.019)		0.081 (0.023)	
Log Earnings per Worker:	0.020 (0.008)	0.023 (0.006)	0.023 (0.010)	0.027 (0.007)
Implied Labor Parameters:				
Labor Supply Elasticity:	4.084	3.600	3.508	3.054
Markdown relative to MRPL:	0.803	0.783	0.778	0.753

Table A.6: Impacts of Winning a Procurement Contract: Prevailing Wage States

Notes: This table presents estimates of the impacts of winning a procurement contract on log employment and log earnings per worker in the post-treatment time period. It provides the impacts on log earnings per worker separately for all workers in the firm and stayers. It compares all states to states that have a state-specific prevailing wage requirement. Specification details and sample definitions are provided in the main text.

	OSHA Investigations		OSHA Violations	
	Probability	Count	Probability	Count
	Occurrence			
Observed Average:	0.075	0.139	0.041	0.110
	Impacts of Winning a Procurement Auction			
Impact: Before Treatment	0.000 (0.006)	-0.012 (0.016)	0.000 (0.004)	-0.009 (0.018)
Impact: After Treatment	0.009 (0.008)	0.004 (0.020)	0.000 (0.006)	-0.006 (0.023)

Table A.7: Impacts of Winning a Procurement Contract: OSHA Safety Outcomes in California

Notes: This table presents estimates of the impacts of winning a procurement contract on OSHA safety investigations and violations in California during 2001-2015. The observed average is reported for bidders (winners and losers) in the years of active bidding.




Figure A.7: Interquartile Range in TFP by Year

Notes: This figure presents the interquartile range of TFP estimated separately by calendar year. Specification details and sample definitions are provided in the main text.




Figure A.8: Heterogeneity across Census Regions in the Estimate of $1-\epsilon$

Notes: This figure presents heterogeneity across Census regions in the estimate of $1-\epsilon$ using the estimator in equation (24), as well as the implied price markup $(1-\epsilon)^{-1}$.

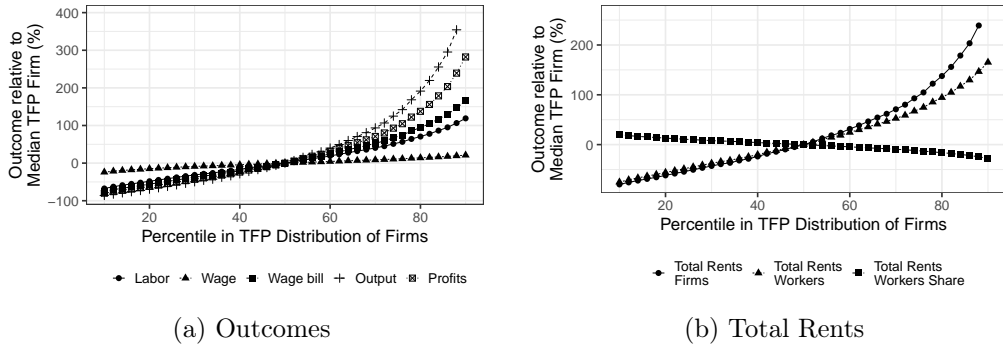


Figure A.9: Outcomes and Rents for Alternative TFP Percentiles, among Firms without Procurement Contracts ($D_{jt} = 0$)

Notes: In this figure, we assign alternative TFP quantiles to the median-TFP firm in the $D_{jt} = 0$ sample without changing any other primitives of the model, then re-solve the model to obtain this firm's alternative outcomes and rents. The x-axis displays the alternative TFP assigned to the firm (as a percentile in the population TFP distribution). Each y-axis value is expressed as a percent change relative to the actual value for the median-TFP firm.

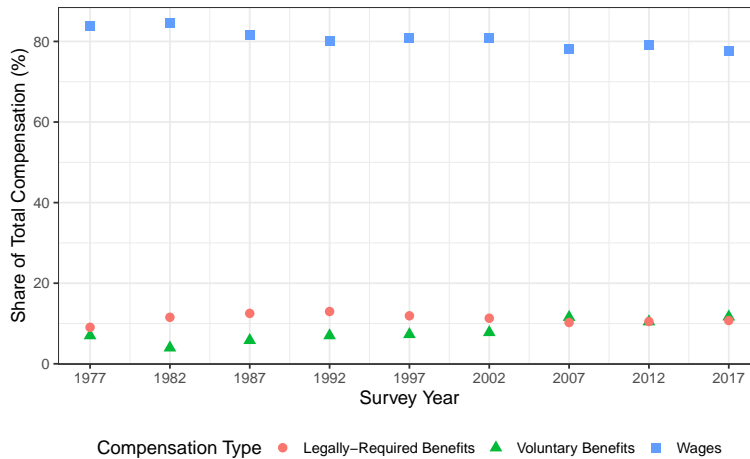


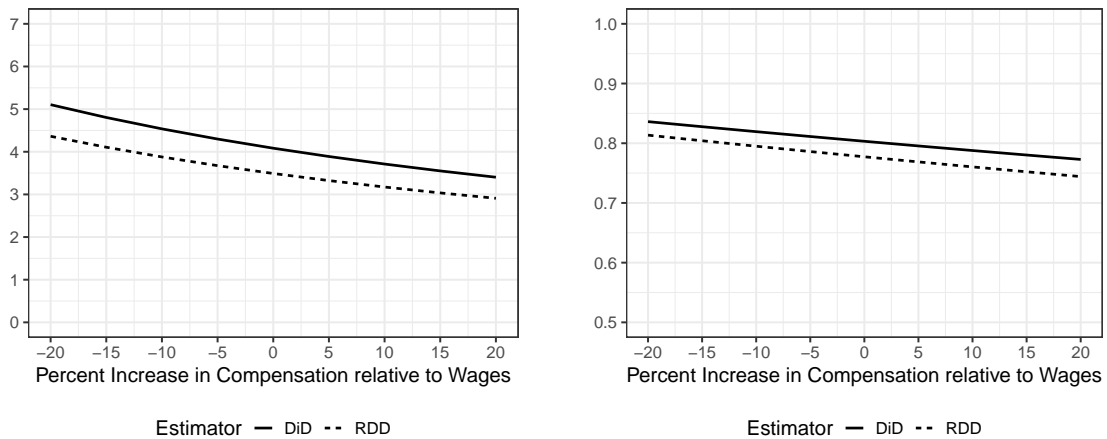
Figure A.10: Survey Evidence from the Economic Census of the Construction Sector: Sources of Compensation in the Construction Industry

Notes: This figure uses survey data from the Economic Census of the Construction Sector to estimate the share of total compensation from legally-required benefits, voluntary benefits, and wages. See Appendix E.1 for details.

Total Compensation (log)	Wage Compensation (log)	Non-wage Fringe Benefits (log)	Share Non-wage Fringe Benefits (fraction)
Difference-in-Differences for State Davis-Bacon Repeals			
0.009 (0.026)	0.009 (0.029)	0.015 (0.031)	0.000 (0.005)

Table A.8: Impacts of Repeals of State Prevailing Wage Laws

Notes: This table combines information on repeals of state prevailing wage laws with survey data from the Economic Census of the Construction Sector to estimate the impacts of prevailing wage law repeals on wage and non-wage compensation. See Appendix E.2 for details.



(a) Labor Supply Elasticity

(b) Wage Markdown relative to MRPL

Figure A.11: Sensitivity to Allowing Endogenous Amenity Responses to Winning a Procurement Auction

Notes: In this figure, we allow for amenity responses to winning a procurement auction, and adjust the θ_{DiD} and θ_{RDD} estimates to account for these amenity responses. See Appendix H for details.

H Sensitivity to Assuming there are Causal Effects of Winning an Auction on Amenities

Let Comp_{jt} denote the total compensation offered by the firm (inclusive of wages and amenities). If we observed Comp_{jt} , we could infer the (inverse) labor supply elasticity with respect compensation from the estimand $\tilde{\theta} = \frac{\mathbb{E}[\Delta \log \text{Comp}_{jt} | \tau_{jt}=0] - \mathbb{E}[\Delta \log \text{Comp}_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]}{\mathbb{E}[\Delta \ell_{jt} | \tau_{jt}=0] - \mathbb{E}[\Delta \ell_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]}$. In practice, however, we only observe wages W_{jt} , so we use $\Delta \log W_{jt}$ in place of $\Delta \log \text{Comp}_{jt}$.

We now examine how the key conclusions regarding the labor supply curve would change if log compensation in reality increased more (or less) than log earning. To do so, it is useful to define $\lambda \equiv \frac{\mathbb{E}[\Delta \log \text{Comp}_{jt} | \tau_{jt}=0] - \mathbb{E}[\Delta \log \text{Comp}_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]}{\mathbb{E}[\Delta \log W_{jt} | \tau_{jt}=0] - \mathbb{E}[\Delta \log W_{jt} | 0 < \tau_{jt} \leq \bar{\tau}]}$ - 1, so that $\lambda \times 100\%$ is the percent increase in log compensation relative to log wages. If $\lambda = 0$, then the change in compensation is proportional to the change in wages. In this case, there is no bias in using wages in place of total compensation when estimating θ . However, if $\lambda > 0$ (or $\lambda < 0$), then the change in log compensation is greater (or smaller) than the change in log wages, so the estimate of θ using wages may be downward-biased (or upward-biased).

In Online Appendix Figure A.11, we calibrate $\lambda \times 100\%$ and examine how our conclusions would change if winning a procurement auction had a causal effect on amenity provision. We find that, even in the rather extreme case in which log compensation increases by 20% more (or less) than log earnings due to amenity responses, the labor supply curve is upward sloping with an elasticity that would be close to our baseline estimates in Figure 3. Furthermore, there is still a significant wage markdown relative to MRPL, and it is close to our baseline estimates using log wages.

I Computational Details

Overview: Simulating model outcomes is computationally challenging. Since $1/\theta$ and $-1/\epsilon$ both appear in the firm's opportunity cost $\sigma(\phi_{jt}, u_{jt})$ (recall the definition associated with equation 8), it follows that changing these parameters also changes the optimal bid Z_{jt}^* (equation 14). In turn, the bid affects the additional rents captured

by firms from winning a procurement contract. To simulate from the model, we first solve the second stage problem for each ϕ_{jt} to find the distribution of opportunity costs. Next, we solve the first stage problem to obtain the distribution of optimal bids given the 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 all outcomes. To ease the computational burden in solving for these distributions we implement the quantile representation method of Luo (2020). Our main results focus on outcomes for the typical firm (the firm with the median value of ϕ_{jt}), which further reduces the computational burden.

Second stage: Denote the TFP quantile function as $\phi(\alpha)$ 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 ϕ 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 (d, α, \bar{Q}^G) , 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 2), the production function (equation 9), and optimal intermediate inputs (equation 10).

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 $\sigma(\alpha|\bar{Q}^G) = \pi_0^H(\alpha) - \pi_1^H(\alpha|\bar{Q}^G)$. Since π_{1jt}^H is the winning firm's revenue in the private market net of the total cost, it follows that $\pi_{0jt}^H > \pi_{1jt}^H$ and thus $\sigma > 0$. π_1^H is decreasing in \bar{Q}^G , and π_0^H does not depend on \bar{Q}^G . Moreover, σ is decreasing in α . In other words, a higher TFP firm has a 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 $f_0(\alpha, I) = \frac{I}{I-1}(1 - \alpha^{I-1})$.

Solution: Let $Y_d(\alpha|\bar{Q}^G)$ 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 $\mathbb{E}[Y_d|\bar{Q}^G, I] = \int_0^1 Y_d(\alpha|\bar{Q}^G) f_d(\alpha, I) d\alpha$. For example, the probability that a bidder with TFP ϕ_{jt} wins the project is the probability that its TFP is the highest among all participating bidders, i.e., $H(\phi_{jt})^I$, where H denotes the distribution of TFP. This implies that the density function of the winner's TFP is $IH(\phi_{jt})^{I-1}h(\phi_{jt})$. The profit function depends on who wins the auction, in particular, the TFP of the winner. The expected profit of the winner is then

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

Note that this expectation depends on the combinations (\bar{Q}^G, I) . One possibility is to solve the model for each possible combination of (\bar{Q}^G, I) , and then average across them. In our setting, this is computationally infeasible. An alternative is to evaluate (\bar{Q}^G, I) at representative values. In practice, we choose the values of (\bar{Q}^G, I) that provide the best fit to the additional rents from procurement projects, $(V_{jt\Delta}, \pi_{jt\Delta})$, for the typical firm. The best fit yields a model-simulated incidence on workers of about \$6,500, which is the same as the main estimate in Table 4, and incidence on firms of \$9,200, which is very close to the main estimate of about \$9,600 in Table 4. The implied incidence share on workers is about 41%, which is about the same as our main estimate. The best fit is achieved at $I = 5$ bidders per auction, which is in the right ballpark to the mean observed value in the data of around 8 bidders per auction.

Additional details: We now provide the derivation of the quantile representation of the optimal bidding strategy. Consider a standard first-price 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:

$$c = b - \frac{1}{I-1} \frac{1 - \mathfrak{H}(b)}{\mathfrak{h}(b)},$$

where $\mathfrak{H}(\cdot)$ and $\mathfrak{h}(\cdot)$ are the bid distribution and density, respectively. Since the bidding strategy is strictly increasing, we can further rewrite this expression 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 \(2020\)](#), we solve this ODE and obtain the mapping from the cost quantile function to the bid quantile function:

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

This representation is convenient for numerically solving the first-price procurement auction model.

S Online Data Supplement

S.1 Acquisition and Preparation of Auction Data

This supplement describes our data sources for auction bids and how we build the data set for our main application. Online Appendix Table A.1 provides a summary of the sources of DOT records by state.

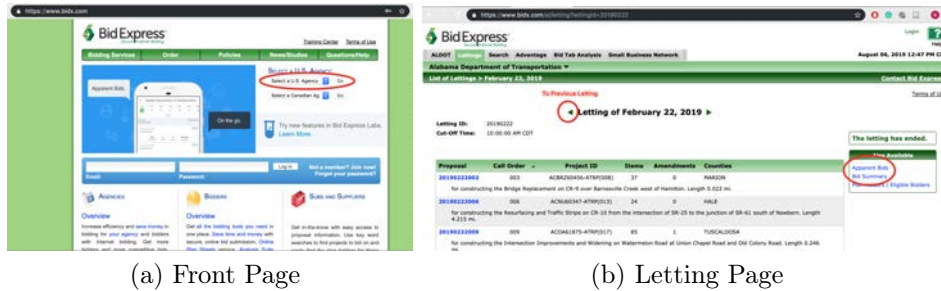
Bid Express Auction Records

The Bid Express website collects information on bids and bidders for procurement auctions held by Departments of Transportation of many US states. It can be freely accessed at www.bidx.com, although the access to information on the bidders requires a paid account registration. We obtained 17 states' DOT auction records from Bid Express. We performed the download 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”. An illustration is provided in Appendix Figure S.1a. 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. An illustration is provided in Appendix Figure S.1b.
2. There are two different sources of information - “Apparent Bids” and “Bid Summary” - on a letting page. More specifically, “Apparent Bids” and “Bid Summary” contain auction information but in different formats, and both of them have links to additional bidder information, which requires a paid account to access. Starting from the latest letting page, our function clicks the hyperlink “Apparent Bids” (Appendix Figure S.1c) then downloads a csv file for every

bidder by clicking on the bidder hyperlink (Appendix Figure S.1d) and “Export(csv)” on the refreshed page.



(a) Front Page

(b) Letting Page



(c) Apparent Bids Page

(d) Bid Summary Page

Figure S.1: Web Pages from BidX.com

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:

- auction information from “Apparent Bids”, which contains: bidder names, bidder ID, bidder ranks, bid amounts, bidder call orders, project description, counties, letting ID and letting date. We do note that a few states record two extra variables: DBE Percentage and DBE Manual.
- auction information from “Bid Summary”, which contains: bidder names, bidder ID, bidder ranks, bid amounts, bidder call orders, counties, proposal ID and letting date.

- 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, two 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 is 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 the 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 $X_apparentbid$.¹ 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: Date of the auction, unique auction identifier, name of the bidding firm, address of the bidding firm, unique firm identifier, total bid amount, and state in which the auction occurred.

State-specific Auction Records

We obtained auction records on 12 other states from two types of sources: downloading from state-specific bidding websites (7 states) and submitting Freedom of Information Act (FOIA) requests to state governments (5 states). Each dataset included different variables and were organized in different formats. For example, the data from Texas included 121 variables while the data from West Virginia included

¹Proposal in $X_bidsummary$ is treated as Letting ID.

only 11 variables. We harmonized these datasets focusing on the core set of variables: Date of the auction, unique auction identifier, name of the bidding firm, address of the bidding firm, unique firm identifier, total bid amount, and state in which the auction occurred.

Note that one state, West Virginia, transitioned from its own website to Bid Express in 2011, so we use combined records from both sources. Once harmonized, we combined the various state-specific DOT auction records with the records obtained from Bid Express.

EIN Availability

We were able to obtain the EINs for firms that bid in DOT auctions in six states:

- Florida, Indiana, and West Virginia: These states' DOT auction records were downloaded from state-specific websites. The EINs were available from these websites.
- Colorado and Kansas: These states' DOT auction records were obtained through FOIA requests. The requested data included EINs.
- Texas: This state's DOT auction records were obtained through a FOIA request. Although this request did not include EINs, we were able to look up EINs by firm name and address through a Texas state government website: <https://mycpa.cpa.state.tx.us/coa/>.

S.2 Matching Auction Data to Tax Records

This supplement describes the procedure adopted to match the bidders in our auction sample to the tax data. For a subset of bidders, the EIN is available in the auction 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.

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.² The distinctive part is named by the business owner and is usually required by governments to be “*substantially different*” from any other existing name. The descriptive part describes what the business does, or its sector.³ 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 corporation 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 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 Dictionary

In order to construct a uniform abbreviation in the legal part, we constructed a many-to-one dictionary using a subsample of our training data. We manually select abbreviations (including for misspelled words) by comparing mismatched names for the same firm in multiple databases. For example, “Incorporated” appears as “Inc.”

²Although there are no specific regulations on this naming structure, it is in alignment with naming convention and government guidelines. <https://www.ic.gc.ca/eic/site/cd-dgc.nsf/eng/cs01070.html>

³An 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>

Algorithm 1 Matching Algorithm Pseudocode

```
Input:  $S^a$ ,  $State^a$ ,  $Dict$ ,  $IRS.Firmlist.normalized$ 
Output:  $Match^a$ ,  $Score^a$ 
1 1. Name Normalization:
2  $S^a_{norm}$  = Remove non-alphanumeric characters and double spaces from  $S^a$  and set uppercase letters;
3  $W^a \equiv \{w_1^a, w_2^a, \dots, w_n^a\}$  = Split substrings at spaces ( $S^a_{norm}$ );
4 for  $i = 1, \dots, n$  do
5   if  $w_i \in Dict$  then  $W = W - \{w_i\}$ ;
6  $S^a_{norm}$  = Merge words in  $W$ ;
7 2. Shortlisting:  $Shortlist = IRS.Firmlist.normalized$ 
8 Candidate="Unmatched"
9 Out=0
10 i=0
11 repeat
12    $i=i+1$ 
13    $C = \{FirmName \in IRS.Firmlist.normalized \mid w_i^a \in FirmName\}$ 
14   Shortlist = Shortlist  $\cap C$ 
15   if Shortlist is singleton then
16     Candidate = Shortlist
17     Out = 1
18   if Shortlist is empty then
19     Out = 1
20   else
21     Candidate = Shortlist
22 until Out=1;
23 3. Scoring
24 for  $c \in Candidate$  do
25    $Score^c = \text{Levenshtein distance}(c, S^a_{norm})$ 
26 Best =  $\text{argmax}\{Score^c\}$ 
27 if  $\text{Levenshtein distance}(\text{Best}, S^a_{norm}) < 0.6$  then
28    $Match^a = \text{Unmatched}$ 
29 else
30    $Match^a = \text{Best}$ ;
31    $Score^a = \text{Levenshtein distance}(\text{Best}, S^a_{norm})$ 
```

“INC”, “Incorp” and so on in our data. Therefore, these abbreviations, when found, are mapped into “Incorporated” as described below. Our dictionary and matching algorithm are available upon request for replicability.

Matching Algorithm

We now describe the database matching algorithm (written in Python). A pseudocode representation of this procedure is provided in Appendix Algorithm 1. For each company name in the auction database, 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. In our main application, the auction 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 by using a “blocking” procedure that prioritizes firms from the same origin state, as explained below. Let a be the firm, S^a be the firm’s string name and $State^a$ be the firm’s 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 sub-string separated by one space is considered a “word”. Every word in the legal-parts 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 “Brothers” is recognized in our dictionary as a synonym for “BROS” and “Inc.” is recognized in our dictionary as a legal part and therefore removed. We refer to the normalized string as S_{norm}^a . The same normalization is applied to every company name in the tax database. If the normalized name is not unique in the tax database, we restrict to the ones that ever filed at least one of the three firm tax returns (1120,1120-S or 1065). If the same firm name filed multiple firm tax returns, we select the one with highest value added, as the firm with greater value added is participating in more economic activity and therefore more likely to be the firm that participated in the auction.

2. SHORTLISTING

Let S_{norm}^a 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_{norm}^a and consider only the candidate matches that also contain 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 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 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 $T = 0.6$.

Appendix Table S.1 illustrates how the algorithm works with an example search, using “Hannaford Bros. Distribution Co.” as the search query. In our example, *strict* and *state* are disabled.

In-Sample Algorithm Validation

In order to validate the algorithm, we apply it to the subset of firms for which we were provided the EIN by the state DOT, thus allowing us to link records exactly rather than using the algorithm (the “Known EIN” sample). The results are displayed in Appendix Table S.2. In Column (1), we provide results from using a simple string matching algorithm, in which a firm in the auction database is only matched to a firm in the tax database if they have identical names. In Columns (2-5), we apply our approach presented in Appendix Algorithm 1. Overall, the algorithm outperforms string matching in both accuracy and number of matches achieved. In our preferred specification in column (5), the algorithm correctly matches 84.5 percent of the bidders whose EIN are known and could be found in tax database. The use of

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

Table S.1: Example Search

the *State* option proves effective in increasing the number of true matches, while the *Strict* option with $T = 0.6$ improves accuracy by reducing the false matches.

Out-of-Sample Algorithm Validation

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). Our test data sets, *PensionData* and *PensionTest*, are constructed using Form 5500 data sets that are published by the Employee Benefits Security Administration (ESBA)⁴. 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 EIN, then remove duplicate observations. For every unique EIN, we find all names that are associated with it, then we discard any duplicate names. Most of the EINs are associated with multiple company names, which reproduces a challenge in the tax database. For each EIN, if multiple names are associated with it, we select the first name and put

⁴<https://www.dol.gov/agencies/ebsa/researchers/data>

	Simple Search	Fuzzy Match			
	(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	✓	✗	✗	✓	✓

Table S.2: In-Sample Algorithm Validation

Notes: This table provides summary statistics on the in-sample performance of the matching algorithm when applied to the six states that provided EINs. For these six states, we observe the true match between auction and tax records. Since some contractors are individuals rather than firms or are otherwise not required to file one of the three firm tax forms, not all contractors in auction data have a true match in firm tax records. First row provides share of contractors in the auction data that the algorithm matches to a firm tax record. Second row provides share of contractors in the auction data that the algorithm matches to a firm tax record and the match is true. Third row provides share of contractors in the auction data that the algorithm matches to a firm tax record and the match is true, among contractors in the auction data for which the true match exists in the firm tax data.

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.

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%

(a) Performance for Values of T

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

(b) Quantiles of Shortlist Lengths

Table S.3: Out-of-Sample Algorithm Validation using Pension Data

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 sets, we can evaluate the matching performance. The program achieved an average speed of 152 queries per second and an average accuracy of 73.39 percent among matched queries for a $T = 0.6$ using the *strict* option. Appendix Table S.3a presents the percentage of correctly matched firms and false matches for different values of T . We note 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 leads to a loss of correct matches that outweighs the gains in precision. Therefore, we do not recommend setting T above 0.9. In Appendix Table S.3b, instead, we provide a closer look at the effectiveness of the shortlisting step. Looking at the distribution of the shortlists' length, we see that over 50% of the sample is matched at the shortlisting step and 70 percent of the candidate matches requires the scoring of at most 2 candidates. Furthermore, the 99th percentile of the longest shortlist amounts to 2,733 candidates. This is only 0.2 percent of the potential matches that a standard matching algorithm would have to consider for each query and, therefore, much more efficient.

S.3 Description of the Tax Data

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 annualized full-time minimum wage and where

the EIN is this worker's highest-paying employer.

- **Wage bill:** Total earnings among employees in year t .
- **Independent contractors:** Number of workers matched to an EIN in year t from Form 1099-MISC with annual compensation above the annualized full-time minimum wage.
- **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.
- **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 Input Expenditures:** 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.

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 Form 1120, 1120-S, or 1065 is considered part of the baseline sample centered around auction cohort t if it is observed bidding in an auction in year t .
- **Sample of non-winners:** A firm in the baseline sample at t that does not win an auction before or during t is called a non-winner if it continues to not win any auctions until at least relative time $e \geq 4$. For example, if $t = 2005$, then a non-winner must not win its first auction until at least 2009.
- **Sample of first-timers:** A firm in the baseline sample at t that does not bid in an auction before t and bids in an auction at t .
- **Sample in the same location:** Firm j and j' are in the same location at t if their business address zip codes reported on the business tax filings correspond to the same commuting zone at t .
- **Known EIN sample:** Firms from the six states in which the auction records included the EIN, thus allowing us to link records exactly rather than using a fuzzy matching algorithm.

Worker Sample Definitions:

- **Main sample:** A worker is considered part of the main sample at t if the worker’s highest-paying firm at t 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.

- **Add Contractors:** Add to the main sample any independent contractor whose highest-paying firm at t on Form 1099 is in the baseline sample of firms and the 1099 wage payments from that firm are greater than \$15,000 in 2015 USD. We also restrict to contractors aged 25-60.
- **Stayers:** A worker is a stayer for $2e + 1$ years at firm j in the baseline sample of firms at t if the worker's highest-paying W-2 firm is the same firm during each time period in $(t - e, \dots, t + e)$ and the W-2 wage payments from that firm in each year are greater than \$15,000 in 2015 USD.
- **Tenure:** A worker has e years of tenure at firm j in the baseline sample of firms at t if the worker's highest-paying W-2 firm is the same firm during each time period in $(t - e, \dots, t)$ and the W-2 wage payments from that firm in each year are greater than \$15,000 in 2015 USD.
- **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 received W-2 wage payments greater than \$15,000 in 2015 USD from j' in $t - 1$ as well as from j in t .

A potential drawback of tax data is limited coverage of undocumented immigrants. As a result, we primarily interpret our paper as providing an analysis of the legal labor market. However, there is substantial coverage of undocumented immigrants in our W-2 returns. Since 1996, the IRS has assigned a tax identification number, called the ITIN, to undocumented immigrants in order to facilitate filing. By law, the IRS cannot share undocumented immigrant status with other agencies for purposes of immigration enforcement, so filing does not pose deportation risks. The IRS imposes penalties on employers for failing to file W-2 tax forms on behalf of undocumented employees, while tax refunds (e.g., child tax credits) and other benefits (e.g., evidence of consistently filing taxes can be used in support of citizenship applications) provide

substantial incentives for undocumented immigrants to file. For example, the [CBO \(2007\)](#) estimated that up to 75 percent of all undocumented immigrants filed during the earlier part of our sample, and this rate may have risen due to DACA and other reforms instituted during the latter part of our sample.

S.4 Description of the Norwegian Data

The Norwegian data comes from the State Register of Employers and Employees, which covers the universe of workers and firms. Our sample spans 2009-2014. For each job, it includes information on start and end dates, annual earnings, and contracted hours. We construct annual earnings at the primary employer as our main outcome of interest. Because the Norwegian data also provides hours worked per day, we can construct the average hourly wage. We supplement the employer-employee data with a measure of value added, which we define as the difference in sales and non-wage operating costs as reported to the Norwegian tax authority by the firm.

To harmonize the Norwegian data with our sample from the US, we follow [Bonhomme et al. \(2023\)](#) by applying the following steps. First, as is common in the literature, whenever a worker is employed by multiple employers in the same year, we focus on the employer associated with the greatest annual earnings. Second, we restrict attention to workers employed in the construction industry. Third, we restrict attention to workers who are between 25 and 60 years of age. Lastly, we restrict attention to full-time equivalent (FTE) workers. Recall that, since we do not observe hours worked in US data, or a formal measure of full-time employment, we defined a worker as FTE if his or her annual earnings exceed \$15,000, which is approximately the annualized minimum wage and corresponds to 32.5% of the national average. To harmonize the sample selection across countries, we similarly restrict the Norwegian sample to workers with annual earnings above 32.5% of the national average.