# Markups and Public Procurement: Evidence from Czech Construction Tenders

**Marek Chadim**[1]

Writing Sample for Ph.D. Applications — Short Version (10 pages)

July 2026

**Abstract**

This paper analyzes the effect of public procurement on firm markups, using a panel dataset of Czech construction firms from 2006 to 2021. Markups are estimated through a structural framework, while biases are addressed using a selection on observables design and models for unobserved factors. Propensity score-based estimations indicate that firm markups increase by approximately 15% during contract years. A temporal analysis, employing synthetic control and matrix completion methods, reveals that treatment effects decline from around 30% in 2006 to 10% in 2021. These patterns are consistent with institutional improvements in the Czech Republic and offer empirical evidence of increasing efficiency in public spending.

**Keywords:** Models with Panel Data; Firm Behavior: Empirical Analysis; Procurement; Construction

**JEL:** C23, D22, H57, L74

[1]Predoctoral Fellow, Tobin Center for Economic Policy, Yale University (`marek.chadim@yale.edu`; `https://marek-chadim.github.io`). This paper is my MSc thesis, submitted to the Department of Economics, Stockholm School of Economics (defended December 2024, awarded the highest grade) and awarded First Place in the 32nd Young Economist of the Year competition of the Czech Economic Society, 2025. I am grateful to my thesis supervisor, Jaakko Meriläinen, for invaluable guidance; to Matěj Bajgar, Mitch Downey, and David Schönholzer for their helpful feedback; and to Jiří Skuhrovec of Datlab s.r.o. for providing administrative data on government procurement. This thesis builds on my earlier work at the Institute of Economic Studies, Charles University Prague (`https://dspace.cuni.cz/handle/20.500.11956/184831`), where I accessed firm financial statement data via MagnusWeb. Replication files: `https://github.com/marek-chadim/thesis`. A substantially revised and extended working-paper version—developed during Ph.D. coursework at Yale and presented in Graduate Industrial Organization II (ECON 601)—is available on request. At Yale, I am grateful to my Tobin Center supervisors, Seth Zimmerman and Barbara Biasi, for their mentorship, and to Joseph Altonji, Steven Berry, Xiaohong Chen, Timothy Christensen, Paul Goldsmith-Pinkham, Philip Haile, Charles Hodgson, Christopher Neilson, Katja Seim, and Edward Vytlacil, whose graduate courses shaped the revision; I especially thank Steven Berry and Katja Seim for detailed comments. This short version contains the introduction, institutional background, and the opening of the empirical framework; the complete 66-page sample is at `https://marek-chadim.github.io/chadim_writing_sample.pdf`.

# 1 Introduction

Public procurement constitutes approximately 12% of GDP in OECD countries, representing a substantial share of government expenditure (OECD, 2021). Ensuring its efficiency is vital for reducing resource waste and maximizing taxpayer value. Existing literature has examined key issues in public procurement, including discretion, political favoritism, and limited competition. However, much of this research has focused on comparing tenders, with relatively little attention devoted to differences between public procurement and private market outcomes.

This paper fills this gap by constructing a novel dataset and proposing methodological solutions to the identification challenges inherent in studying the dynamics of firm markups as firms enter and exit public procurement markets. To the best of my knowledge, this is the first study to provide credible econometric evidence on the relationship between markups and participation in public procurement. The findings contribute to ongoing debates on public procurement efficiency, market competition, and the implications of government contracts for firm-level market power.

Using the methodology of De Loecker and Warzynski (2012), I estimate markups under the assumption that cost-minimizing producers optimize variable input expenditures. This approach provides plant-level markup estimates without relying on assumptions about firms' competitive behavior in product markets. The results document the evolution of aggregate markups, showing a decline from 40% above marginal cost in 2006 to 30% in 2021, driven primarily by firms at the upper end of the markup distribution.

Next, I estimate the public procurement premium using non-experimental methods under the assumption of unconfoundedness, conditioning on information available prior to firms' engagement in public procurement. To mitigate biases arising from observational characteristics, I apply methodological advancements from Imbens and Xu (2024). Placebo tests, which compare the average markups of treated and control firms before the treated firms receive contracts, and sensitivity analyses benchmarking the omitted variable bias further reinforce the robustness of the findings.

Finally, I employ panel data models for binary interventions in longitudinal settings to estimate the causal effect of public procurement on the markups of government contractors, drawing on Arkhangelsky and Imbens (2024). These methods extend traditional difference-in-differences and two-way fixed-effects models by incorporating factor models, interactive fixed-effects models, and synthetic control approaches. The analysis emphasizes heterogeneity in causal effects and uses valid inference techniques suited to cases with a small number of treated units.

Drawing causal inferences from observational data, even with advanced methods, requires strong assumptions. The lack of natural experiments, instrumental variables, or regression discontinuity designs limits the availability of exogenous variation in this study. To address these challenges, I use a selection-on-observables approach, framed within the potential outcomes framework, to model counterfactual markups for government contractors in the absence of public procurement. The stable unit treatment value assumption (SUTVA) excludes interference between units and the dynamic effects of past treatments and outcomes. However, in competitive public procurement markets, violations of SUTVA can complicate the interpretation of causal effects. Despite this, the study prioritizes internal validity through a design-based balanced panel approach while enhancing generalizability by leveraging the full dataset. Balanced panel estimates, based on selection on observables, indicate that procurement raises contracting firms' markups by an average of 15%, with the full panel data model showing a decline in this effect from 30% in 2006 to 10% in 2021.

Industrial organization models with heterogeneous producers and firm specific markups provide potential explanations for these findings. More productive firms may absorb the costs associated with entering public procurement markets, which could explain why government contractors typically exhibit higher markups. Additionally, contractors may charge higher markups if they supply higher-quality goods produced using higher-quality inputs. Another plausible explanation is that procurement contracts often provide firms with long-term clients, leading to less elastic realized demand and enabling higher markups.

The literature on public procurement also highlights other mechanisms, such as discretion and political favoritism. My findings suggest that the public procurement markup premium has declined over time, consistent with reforms targeting single-bidding practices Titl (2023) and increased monitoring associated with European Union co-financing (Baránek and Titl, 2024).

The remainder of the paper is structured as follows. Section 1.1 reviews related work. Section 2 introduces the dataset and the construction of the public procurement indicator. Section 3.1 documents the distribution of markups and their descriptive relationship with public procurement. Section 3.2 presents the causal results under the unconfoundedness assumption, while Section 3.3 reports findings from panel data analyses. Methodological details supporting these sections are outlined in Appendices A, B, and C. Finally, Section 4 concludes the paper.

## 1.1 Related Work

**Public Procurement** A central challenge in public procurement is navigating the trade-off between the benefits of discretion and the risks of rent-seeking. When the interests of procuring officials diverge from those of the public, discretion can heighten the risk of corruption and lead to the misallocation of contracts to politically connected firms. Conversely, auctions, while more transparent, are often more costly and time-consuming to organize compared to less formal procedures like direct negotiations. In the Czech Republic, Titl (2023) reports that approximately 23% of public contracts are awarded through single-bid procedures and that a 2012 reform aimed at reducing single-bidding led to a decrease in procurement project prices by about 6% of their estimated costs. Similarly, Kang and Miller (2022) find that single-bidding is prevalent in the United States but argue that information asymmetries and noncontractible quality dimensions can necessitate some level of discretion to incentivize effective contract execution, leaving the net impact of discretion on procurement outcomes ambiguous. In the Czech context, Palguta and Pertold (2017) document that officials frequently manipulate procurement thresholds by adjusting contract values to circumvent competitive bidding. However, identification challenges associated with bunching complicate empirical assessments. Using a policy reform in Hungary, Szucs (2024) demonstrates that increasing discretion for public agencies results in higher prices, less productive contractors, and a greater prevalence of politically connected winners. Similar findings emerge in Italy, where Decarolis et al. (2020) show that discretionary procedures in government contracts, particularly those involving fewer bidders, heighten corruption risks and disproportionately benefit politically connected firms. In the Czech Republic, Baránek and Titl (2024) reveal that politically connected firms secure contracts priced 6% higher than competitively awarded contracts, with no corresponding improvement in quality. This misallocation aligns with evidence from Italy, where Bandiera et al. (2009) find that weaker governance leads public entities to overpay for comparable goods.

**Markups and Market Power** Markups, defined as the ratio of price to marginal cost, $\mu \equiv \frac{P}{c}$, are a natural measure of market power, emerging when a firm can influence the price at which it sells its product. Profit-maximizing firms leverage this ability to set prices above marginal cost, resulting in markups greater than one. Syverson (2024) highlights the dual role of markups. On the positive side, markups signal both the presence and magnitude of market power. On the normative side, markups contribute to deadweight loss, as firms with market power deliberately restrict output, forgoing socially beneficial transactions to maintain higher prices. Additionally, market power redistributes surplus, shifting it toward dominant market participants. As a result, market power influences total welfare under non-egalitarian social welfare functions, even beyond the inefficiencies caused by deadweight loss. Empirical evidence shows a significant rise in U.S. markups since 1980. De Loecker et al. (2020) document that average markups increased from 21% above marginal cost to 61% by 2016, with this growth concentrated among large firms, signaling heightened market concentration. Hall (2018) corroborates these findings, while Autor et al. (2020) examine the labor market consequences, observing that "superstar firms" increasingly dominate profits while employing fewer workers. This trend has contributed to a decline in labor's share of GDP and exacerbated income inequality. The implications of rising markups, however, remain a topic of debate. Berry et al. (2019) argues that higher markups often reflect reduced competition and rent extraction, particularly in industries with significant barriers to entry. In contrast, Shapiro and Yurukoglu (2024) contend that rising markups may instead signify competitive dynamics, where efficient firms expand by offering superior products. Miller (2024) attributes the increase in markups to technological advancements, in that productivity improvements and higher product quality enable firms to charge higher markups without necessarily harming consumer welfare.

**Markup Econometrics** Two methods are commonly used to estimate markups: one demand-based and one production-based. For measuring markups and conducting retrospective studies, the two approaches are arguably substitutes (De Loecker and Scott, 2016). However, the production approach cannot be used for counterfactual analysis, which requires a detailed understanding of demand. Recall the formula for monopoly pricing: $\frac{P}{c} = \frac{1}{1+E_D^{-1}}$, where $E_D^{-1}$ is the inverse elasticity of demand. In more complex settings, such as those involving differentiated products, one can still solve for markups as a function of demand elasticities. The demand-based approach has been the standard method for analyzing pricing behavior, but it depends on several restrictive assumptions. First, it typically assumes static Nash-Bertrand competition or some other form of imperfect competition that allows for a tractable equilibrium solution. Second, it requires instruments to identify demand elasticities accurately. Finally, the approach involves functional form assumptions on the demand system as well as a model of consumer heterogeneity, potentially limiting its applicability in complex or dynamic markets. De Loecker and Warzynski (2012) (DLW) introduced a method for estimating firm-level markups directly from production data. By combining input-output elasticities with input revenue shares, this approach captures market power in both product and factor markets. It overcomes many limitations of demand-based models by relying on more readily available production data. However, the DLW method has its own limitations. One critique is its assumption of Hicks-neutral productivity, where productivity shifts equally across all inputs. Raval (2023) shows that this assumption, combined with labor market frictions like hiring costs or monopsony power, can introduce bias by treating inputs such as labor and materials similarly. Bond et al. (2021) highlight that the DLW approach suffers from "omitted price bias" when using revenue data rather than actual output quantities, as revenue-based estimates fail to account for firm-specific price variations. Despite these criticisms, De Ridder et al. (2024) demonstrate that revenue-based markups still reveal important trends over time and across firms, making the DLW approach valuable for studying market power patterns, even if markup level estimates require careful interpretation. The DLW method relies on a first-order condition for the cost minimization of a variable input in production to express markups based on the output elasticity of the variable input and its revenue share. The assumption is that, in each period, producers minimize costs by choosing inputs optimally, free from frictions. The sketch of the main idea follows. Let $P^v_{it}$ denote the price of input $v$, and $P_{it}$ the price of output. The production function is given by $Q_{it} = Q_{it}(X^1_{it}, \dots, X^V_{it}, K_{it}, \omega_{it})$, where $v = 1, 2, \dots, V$ indexes variable inputs. Assuming that variable inputs are set each period to minimize costs, the Lagrangian for the cost minimization problem is as follows: $\mathcal{L}(X^1_{it}, \dots, X^V_{it}, K_{it}, \lambda_{it}) = \sum_{v=1}^{V} P^v_{it} X^v_{it} + r_{it}K_{it} + \lambda_{it}(Q_{it} - Q_{it}(\cdot))$. The first-order condition for input $v$ is $P^v_{it} - \lambda_{it}\frac{\partial Q_{it}(\cdot)}{\partial X^v_{it}} = 0$, where $\lambda_{it}$ is the marginal cost of production at the production level $Q_{it}$. Multiplying this condition by $X^v_{it}/Q_{it}$ leads to the following equation: $\frac{\partial Q_{it}(\cdot)}{\partial X^v_{it}}\frac{X^v_{it}}{Q_{it}} = \frac{1}{\lambda_{it}}\frac{P^v_{it}X^v_{it}}{Q_{it}}$. With $\mu_{it} \equiv P_{it}/\lambda_{it}$, this becomes

$$\frac{\partial Q_{it}(\cdot)}{\partial X^v_{it}}\frac{X^v_{it}}{Q_{it}} = \mu_{it}\frac{P^v_{it}X^v_{it}}{P_{it}Q_{it}}, \tag{1}$$

where we have multiplied and divided by $P_{it}$. This derivation produces a simple expression for the markup: $\mu_{it} = \theta^v_{it}(\alpha^v_{it})^{-1}$, where $\theta^v_{it}$ is the output elasticity with respect to input $v$, which is generally specific to the producer and time period, and $\alpha^v_{it}$ represents expenditures on input $v$ as a share of revenues. On its own, this formula is nothing new. What's new about DLW is how flexible they are about estimating $\theta^V_{it}$ and how they base their inferences about markups on careful production function estimation.[1]

[1]DLW adopt the control function approach, pioneered by Olley and Pakes (1996) and later refined by Levinsohn and Petrin (2003), to link the production function to the economic model describing firm behavior and the competitive environment in which firms operate. Ackerberg et al. (2015) consolidates these by introducing modified assumptions on the timing of input decisions, moving the identification of all coefficients of the production function to the second stage of the estimation, which I follow and outline in Methods Section A.

## 2 Data

I utilize a dataset obtained through licensed[2] access to the MagnusWeb database[3], which provides full company accounts for an unbalanced panel of 1,297 firms active in the Czech construction sector. The dataset contains 7,261 observations covering the period from 2006 to 2021. The sample size accounts for the markup estimation requirement of having at least one lag of data. For robustness of the production function estimates, the top and bottom percentiles of firms are trimmed based on the sales-to-cost-of-goods-sold ratio.

In the Czech Republic, public procurement contracts are awarded under nationwide regulations requiring the publication of contract details by procurers. The dataset includes contracts awarded by central, regional, and municipal governments, as well as government-owned enterprises. A private company[4] corrected and cleaned the raw data to ensure its accuracy. I link firm financial statements with public procurement records to construct a binary indicator of government tender sales. This classification identifies firms that operate exclusively in the private sector, those that enter public procurement, those that exit it, and those that remain as government contractors.

| **Year** | No. Firms | Share Contractors |
|---|---|---|
| 2006 | 227 | 0.38 |
| 2007 | 290 | 0.33 |
| 2008 | 348 | 0.32 |
| 2009 | 412 | 0.32 |
| 2010 | 497 | 0.28 |
| 2011 | 506 | 0.26 |
| 2012 | 457 | 0.33 |
| 2013 | 235 | 0.41 |
| 2014 | 243 | 0.40 |
| 2015 | 245 | 0.48 |
| 2016 | 338 | 0.51 |
| 2017 | 660 | 0.57 |
| 2018 | 708 | 0.55 |
| 2019 | 764 | 0.53 |
| 2020 | 769 | 0.55 |
| 2021 | 562 | 0.53 |
| Total | 7,261 | 0.44 |

[2]Charles University Prague, Faculty of Social Sciences, Institute of Economic Studies
[3]Dun & Bradstreet Czech Republic, a.s.
[4]Datlab s.r.o.

# 3 Results

Section 3.1 documents the main patterns of markups in the Czech construction sector and the relationship between markups and public procurement. I validate these results using two treatment effect identification strategies.

Section 3.2 employs a design-based approach relying on unconfoundedness, or selection based on observable financial statement information and pre-procurement markups. In contrast, Section 3.3 focuses on imputing the counterfactual markup evolution for government contracting firms, addressing selection on latent unobservable factors and incorporating synthetic control methods.

## 3.1 Markups and Public Procurement in the Czech Construction Sector

I use the empirical framework of De Loecker and Warzynski (2012)[5] to correlate markups with a firm's public procurement status and evaluate whether they change upon entering public procurement, controlling for input usage and aggregate trends. While I remain agnostic about a specific theoretical model, I draw on various economic mechanisms to interpret the findings. To estimate markups in the context of public procurement, I incorporate a firm's public procurement status into input demand equations, allowing it to directly affect the law of motion for productivity. After estimating the output elasticity of the variable inputs, I compute the implied markups using the first-order conditions leading to Equation 1.

**The Evolution of Markups in Czech Construction** Equation 1 defines the markup as the product of the output elasticity, $\theta$, and the inverse of the variable input's revenue share, $\frac{PQ}{P^v X^v}$. The revenue share is directly observed in firms' income statements, while the output elasticities are estimated. These elasticities are both firm- and time-specific, which reflect technological differences across firms and changes over time. The average markup is calculated as:

$$\mu_t = \sum_i m_{it}\mu_{it},$$

where $m_{it}$ represents the weight of each firm. I use the share of sales as the weight and compare it with total costs as the input weight. Figure 1 shows the evolution of average sales-weighted and input-weighted aggregate markups in the Czech construction sector.

[5] See Methods A for markup and production function estimation details.

Figure 1: Average Markups. Estimated output elasticities $\hat{\theta}^V_{it}$ are time- and firm-specific.

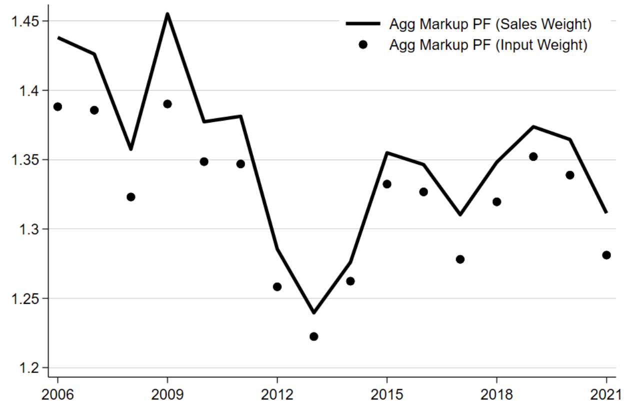


Early in the sample period, markups were stable at around 1.45, declined to approximately 1.25 during 2012–2014, and rose slightly to just above 1.35 by the early 2020s. In 2006, the average markup stood at 44% above marginal cost, compared to 31% in 2021. However, averages alone fail to capture the changes in the markup distribution. The strength of the De Loecker and Warzynski (2012) method lies in its ability to estimate firm-specific markups, which enables an analysis of the entire distribution. A key finding is that a small subset of firms drive the overall decline in markups, while most firms see only modest decreases.

Figure 2a illustrates the shifts in the distribution by showing the kernel density of the unweighted markups for 2006 and 2021. Over time, the variance of markups has narrowed, driven primarily by a thinning and shortening of the upper tail. Appendix Table 5 provides the unweighted markup distribution by year.

Figure 2: The Distribution of Markups $\hat{\mu}_{it}$

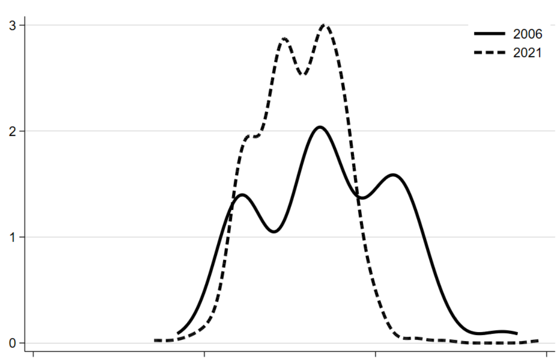


(a) Kernel Density (unweighted)

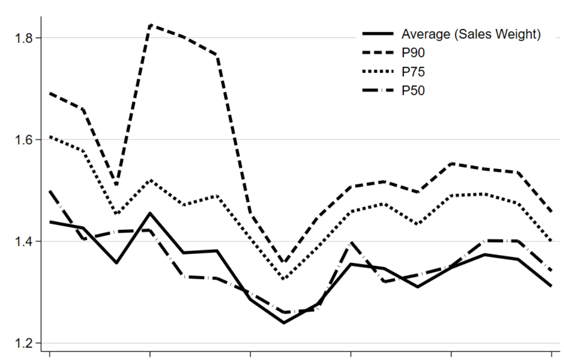


(b) Percentiles (revenue-weighted)

Because the kernel density does not account for firm weights, Figure 2b shows the moments of the sales-weighted markup distribution over time. Firms are ranked by markup, and percentiles are weighted by each firm's market share. This weighting ensures comparability between the percentiles and the share-weighted average. The rankings are updated annually, meaning that the top firms may vary from year to year. Firms in the upper half of the markup distribution primarily drive the decline in average markups. The median (P50) and lower percentiles remain mostly stable over time. The sharpest drop occurs at the 90th percentile, particularly before 2012, when markups fell significantly from 1.8 to 1.4. This indicates that the overall reduction in average markups is due to a few firms experiencing substantial declines compared to earlier periods. This finding provides preliminary evidence of the evolving relationship between markups and public procurement. It is particularly relevant when viewed alongside Titl (2023), who demonstrated that the 2012 Czech public procurement reform, which eliminated single-bid contracts, reduced prices relative to estimated costs for these contracts.

**Do Government Contractors Have Different Markups?** I examine whether government contractors systematically have different markups compared to private-sector firms by using firm-specific markups in a regression framework. Specifically, I estimate the percentage difference in markups between government contractors and private-sector firms.

This approach ensures robustness even if the variable inputs used to compute markups are subject to adjustment costs, provided that such costs do not disproportionately affect government contractors. After estimating the regression, I convert the percentage differences into absolute markup differences.

The regression specification is as follows:

$$\ln \mu_{it} = \delta_0 + \delta_1 pp_{it} + \mathbf{b}'_{it}\sigma + \nu_{it}, \tag{2}$$

where $pp_{it}$ is a dummy variable indicating a firm's public procurement status, and $\delta_1$ measures the percentage markup premium for government contractors. The vector $\mathbf{b}_{it}$ includes control variables[6], with $\sigma$ as their corresponding coefficients.

It is important to note that $\delta_1$ is not interpreted as a causal parameter. Instead, this specification identifies whether government contractors, on average, have different markups.

[6]Control variables include variable input and capital use to account for differences in size and factor intensity, along with 47 year–subindustry interaction dummies to capture aggregate markup trends.