Individuals and Organizations as Sources of State Effectiveness*

Michael Carlos Best† Jonas Hjort‡ David Szakonyi§

August 2022

Abstract

Bureaucrats implement policy. How important are they for a state's productivity? And do the tradeoffs between policies depend on their effectiveness? Using data on 16 million public purchases in Russia, we show that 39 percent of the variation in prices paid for narrowly defined items is due to the individual bureaucrats and organizations who manage procurement. Low-price buyers also display higher spending quality. Theory suggests that such differences in effectiveness can be pivotal for policy design. To illustrate, we show that a common one—bid preferences for domestic suppliers—substantially improves procurement performance, but only when implemented by *ineffective* bureaucrats.

JEL codes: O1, H1

*We are grateful to editor Esther Duflo and two anonymous referees for comments that substantially improved the paper. We also thank Ben Olken for his very helpful discussion of the paper at the AEA meetings, and John M. Abowd, Daron Acemoglu, Hunt Allcott, Tim Besley, Nick Bloom, Richard Blundell, Raj Chetty, Allan Collard-Wexler, Francesco Decarolis, Andres Drenik, Pascaline Dupas, Ben Faber, Fred Finan, Matthew Gentzkow, Josh Gottlieb, Caroline Hoxby, Amit Khandelwal, Brad Larsen, David Margolis, Torsten Persson, Andrea Prat, Imran Rasul, Jimmy Roberts, Jonah Rockoff, Tobias Salz, Orie Shelef, Andrei Shleifer, David Silver, Eric Verhoogen, Guo Xu, and Danny Yagan for valuable comments; Andrei Yakovlev for institutional guidance; Andrey Druzhaev, Len Goff, Vinayak Iyer, and Georgiy Syunyaev for outstanding research assistance; and seminar participants at the AEA meetings, U.C. Berkeley, University of British Columbia, Central European University, Columbia, University of Copenhagen, Duke, the Econometric Society, Harvard Business School, Higher School of Economics, McGill, University of Michigan, the NBER Development Economics Program meeting, NEUDC, NYU, Northwestern, University of Oslo, Paris School of Economics, Santa Clara, Stanford, UCL, UCLA, University of Toronto, University of Virginia, and the World Bank for comments. Best thanks the Stanford Institute for Economic Policy Research, where he was a postdoctoral fellow while much of the work on the paper was done. The study has been funded within the framework of the Basic Research Program at the National Research University Higher School of Economics (HSE) and by the Russian Academic Excellence Project "5-100". All remaining errors are ours alone.

†michael.best@columbia.edu. Columbia University, BREAD, CEPR, & NBER

‡j.hjort@ucl.ac.uk. UCL, BREAD, CEPR, & NBER

§dszakonyi@gwu.edu. Harvard University, George Washington University, & Higher School of Economics

1 Introduction

A successful state is the foundation economic development is built on (Besley & Persson, 2009; Page & Pande, 2018). States delegate policy implementation to their middle management tier, the bureaucracy. Historically, the dominant view in social science was that states could and should strive for a Weberian bureaucracy—"machines" mindlessly translating policy into output, ensuring uniform provision of public services (Weber, 1921). In reality, the skills, organizational capacity, and priorities of bureaucrats differ. But by how much? And what are the implications for policy design?

This paper aims to advance our understanding of the state's production function, an object that remains almost entirely unknown.[1] Our goals are two-fold. First, to quantify the importance of the bureaucracy for the productivity of the state. Second, to explore how the tradeoffs between different policies depend on the effectiveness of the bureaucracy in charge of implementation. The second goal is of particular importance in the public sector, where policy design may be relatively malleable compared to modifying hiring, training, and incentive practices to directly improve bureaucratic effectiveness. Both goals are challenging, as bureaucracies produce a wide array of outputs that are difficult to measure. However, one task—the procurement of off-the-shelf goods—is performed throughout the state enterprise, and has a well-defined and quantifiable primary output: prices paid.

We use a simple conceptual framework of procurement with endogenous supplier entry to guide our analysis of administrative data covering the universe of public procurement in Russia. With an empirical specification derived from the model, we estimate that 39 percent of the variation in performance—quality-adjusted prices paid—is attributable to the bureaucrats who manage procurement, roughly half to individual procurement officers and half to the end-user public organizations. The procurers also explain 24 percent of the variation in spending quality, and price- and quality-effectiveness are positively correlated. Differences in effectiveness of such magnitude have far-reaching implications for policy design. To illustrate, we examine the introduction of a bid preference regime common throughout the world. Under Russia's bid preferences, contract-winners offering goods manufactured abroad are paid 85 percent of their bid. Consistent with our model's predictions, we find that preferences substantially reduce costs and increase competitiveness, but only when the policy is implemented by *ineffective* bureaucrats.

Public procurement in Russia is an ideal setting to study micro-level state effectiveness. First, procurement makes up roughly 8 percent of worldwide GDP (Schapper *et al.*, 2009). Second, for purchases of items that are precisely defined ("off-the-shelf" goods), procurers' mandate is simply to pay the lowest possible price while following the government's policy rules (see also Bandiera *et al.*, 2009; Ferraz *et al.*, 2015).[2] This makes performance measurable and comparable across the entire state

[1] This is despite a growing literature on front-line public sector workers (see e.g. Finan *et al.*, 2017, for an overview).

[2] Russia spends over half of its total public procurement budget on such goods.

enterprise. Third, Russia's massive and diverse bureaucracy spans a wide range of state effectiveness. Finally, the labor market for Russian procurement officers is decentralized and the resulting private-sector-like churn makes it possible to identify individuals' and their employers' effectiveness.

In our stylized model, bureaucratic effectiveness affects procurement outcomes in two ways. First, ineffective bureaucracies impose costs (e.g. unusual product specifications) that raise the cost to suppliers of fulfilling the contract. Second, ineffective bureaucracies impose higher participation costs (e.g. required deposits, or bribes to enter the auction) on bidders. As a result, they attract fewer participants and pay higher quality-adjusted prices.

To compare the performance of bureaucrats (procurement officers) and organizations (e.g. ministries, schools or hospitals) across the country empirically, we need to ensure that they are performing the same task—buying the same type and quality of good. To do this, we adapt tools from machine learning to develop a methodology that uses the text of procurement contracts to classify purchases into homogeneous bins.[3] We also confirm that our results are very similar in a subsample of goods that are by nature homogeneous—pharmaceuticals—and alongside price-effectiveness consider spending quality outcomes (such as delays, contract renegotiation, and cost over-runs). For identification we exploit the fact that many organizations are observed working with multiple bureaucrats and vice versa. This provides us with thousands of quasi-experiments that capture the causal impact of individual bureaucrats and organizations on prices paid under weak assumptions on bureaucrat–organization matching. Event studies reveal large and sharp decreases in quality-adjusted prices paid when organizations switch to more effective bureaucrats, and vice versa, supporting a causal interpretation of these effects.[4]

To aggregate the impacts of individual bureaucrats and organizations on prices paid into an estimate of the share of the total variation explained by the bureaucratic apparatus as a whole, we extend the variance decomposition approach pioneered by Abowd *et al.* (1999, 2002) (hereafter AKM) in two ways. First, we correct the fixed-effect estimates for sampling error using split-sample methods (Finkelstein *et al.*, 2016; Silver, 2016), and by extending shrinkage methods (Kane & Staiger, 2008; Chetty *et al.*, 2014) to a two-dimensional context to explicitly account for the covariance between the estimation error in the bureaucrat and the organization effects (Andrews *et al.*, 2008).[5] Second, we show how to estimate lower bounds on the variation explained by bureaucrats and organizations in a setting—like ours—where bureaucrat switching does not link *all* organizations and how the combined productivity effect of bureaucrats and organizations can nevertheless be identified.

[3]Our methodology ensures that within-category quality differences are minimal, while maintaining generality by not restricting to very specific types of goods. In foregoing conventional methods for categorizing comparable goods and instead using text analysis, we follow Hoberg & Phillips (2016). They classify *firm* similarity based on the goods produced.

[4]Importantly, our estimates can be interpreted causally even if bureaucrats sort across organizations based on the effectiveness of the bureaucrat and/or the organization. Instead, the assumptions needed for causal interpretation are that they do not sort based on unmodelled *match* effects, and that drift in effectiveness and switches are uncorrelated. The event studies provide compelling support for these assumptions, as does a battery of additional tests. Studies of the wages of workers and firms in the private sector tend to find the same (see Card *et al.* (2018); Bloom *et al.* (2019) for overviews of the literature).

[5]To our knowledge, two-dimensional shrinkage estimators like the ones we develop have not been used before.

We find that the bureaucracy jointly accounts for 39 percent of the variation in quality-adjusted prices paid, of which individuals and organizations account for roughly equal shares. Moving the worst-performing quartile of procurers to 75th percentile-effectiveness would reduce procurement expenditures by around 12 percent, or USD 10 billion each year—roughly 15 percent of the total amount the Russian state spends on health care. This would likely entail *better* spending quality too: the buyers in charge of procurement explain a quarter of the variation in spending quality, and price- and spending quality-effectiveness are positive correlated. Procurers' "type" thus appears to influence performance more than any multitasking incentives pulling price and spending quality apart.

We exploit our rich set of indicators on each procurer's auctions—measures of entry barriers chosen, how the auction was executed, procurer experience, etc—to select and explore the 30 most predictive correlates of estimated effectiveness (see also Lacetera *et al.*, 2016). Consistent with our model, we find that effective procurers set lower reservation prices, and attract and admit more applicants to their auctions. While some other measures of bureaucrat behavior also predict effectiveness—for example, low-price procurers attract a somewhat wider variety of bidders—they generally do so to a lesser extent, and a wide range of auction characteristics, including measures of corruption, do not.

The second part of the paper focuses on the implications of heterogeneity in policy *implementer* effectiveness for policy *design*—often a more feasible path to better performance than attempting to directly increase bureaucrats' productivity. We focus on the particular example of industrial policy implemented through public procurement: bid preferences benefitting domestically manufactured goods.

In our model, introducing bid preferences makes participation less attractive to foreign bidders and more attractive to locals. When state effectiveness is high, so is baseline participation and so preferences induce a modest decrease in participation. However, when state effectiveness is low, baseline participation is low and so is the likelihood that a local bidder who enters has to face a more efficient, foreign, bidder. Bid preferences then have a large impact on the likelihood that a local bidder can win the contract, leading to a significant increase in participation. Additionally, foreign bidders shade their bids upward to offset the bid penalty. The overall impact on prices paid combines these participation and bidding responses with the mechanical effect of paying less to foreign winners. We show that the ultimate price effect depends negatively on baseline state effectiveness: effective buyers see performance worsen and vice versa.

We identify the impact of the bid preference regime using a generalized difference-in-differences approach that takes advantage of the fact that preferences apply to an evolving set of goods and are in effect for only parts of each year. Our results reveal that, *on average*, bid preferences achieve the Russian government's goal of channeling demand to domestic manufacturers, and do so at no cost to the government. If anything, average prices paid decrease slightly.[6]

[6]The contrast between this average treatment effect and that of similar policies in more effective states (see e.g. Marion, 2007; Krasnokutskaya & Seim, 2011) suggests that industrial policies in public procurement may be more successful in countries with *low* average bureaucratic effectiveness. This foreshadows our next findings.

To test our model's heterogeneous treatment effect predictions, we interact the bid preference regime with our estimates of the effectiveness of the implementing bureaucrats. The small negative average effect on prices masks considerable heterogeneity. Our estimates imply savings of 12 vs. 0.7 percent when the policy is implemented by the least vs. the most effective quartile of bureaucrats, and that prices increase for some of the *most* effective bureaucrats (as has been shown for similar policies implemented in the U.S.).[7] We also find that characteristics of the procurers and how they initiate purchases that predict effectiveness in a constant policy regime—part one of our analysis—also predict how the policy affects prices paid when implemented by a given procurer. This suggests that policy changes can markedly affect state productivity even absent significant changes in policy implementer behavior.[8]

Overall, this paper demonstrates that state effectiveness is to a large extent embedded in the individuals and organizations of the bureaucratic apparatus, and that tailoring the design of policy to the implementing bureaucracy can partly offset the costs of bureaucratic ineffectiveness.

The first of two strands of literature on state effectiveness we contribute to focuses on individuals and the incentives they face as sources of productivity (see, among many others, Dal Bo *et al.*, 2013; Duflo *et al.*, 2013, 2018; Bertrand *et al.*, 2020; Khan *et al.*, 2016, 2018; Rasul & Rogger, 2018).[9] We quantify, for the first time, the "macro" importance of the bureaucracy for public sector output—the share of overall variation in performance explained by bureaucrats *relative to (all) other contributors*. We sidestep concerns about multitasking and unobserved dimensions of performance by developing a new approach to measuring task-specific productivity and in parallel analyzing spending quality as a separate outcome.[10]

The second strand focuses on how public policy design should be tailored to context (see e.g. Laffont, 2005; Best *et al.*, 2015; Duflo *et al.*, 2018; Hansman *et al.*, 2019). The fact that *implementation* of policy is delegated to bureaucracies is often overlooked. Bureaucracies differ in effectiveness

[7]In the pharmaceuticals sample, where we observe goods' origin, we also find that purchases administered by ineffective bureaucrats see a bigger increase in the probability that an auction is won by a supplier selling locally manufactured goods when bid preferences apply, consistent with our theoretical framework.

[8]We also find that additional features downstream in procurement processes—characteristics of the auction itself and the supplier—become important under bid preferences.

[9]Jones & Olken (2005); Xu (2018) study how public sector leaders and politicians matter for aggregate economic outcomes. In addition to Bandiera *et al.* (2009); Ferraz *et al.* (2015)—who, like us, focus on purchases of off-the-shelf goods—Lewis-Faupel *et al.* (2016); Coviello *et al.* (2017, 2018); Decarolis *et al.* (2018) also study state effectiveness in the context of public procurement. The innovative study by Decarolis *et al.* (2018) is especially related to this paper. The authors investigate how bureaucratic competence affects procurement outcomes in a setting where there are multiple dimensions to both competence and procurement outcomes, and find large effects.

[10]We thus avoid the limitations that arise from comparing workers and/or organizations (e.g. firms) (i) engaging in different or competing activities and/or (ii) based on wages and profits. The seminal work of Abowd *et al.* (1999, 2002) spawned a large empirical literature using employer–employee matched datasets to address a range of important questions in labor economics (see, among many others, the papers cited in footnote 4, and also Bertrand & Schoar (2003) and the literature that followed on CEO effects). Wages do not necessarily reflect productivity (Card *et al.*, 2016), but are important objects in and of themselves. Existing applications of the AKM method have used samples that include workers performing many different tasks. Carneiro *et al.* (2012) show the potential importance of accounting for differences in tasks. On the organization/firm side, conventional methods estimate productivity from revenue or profits data and thus risk conflating productivity itself with mark-ups and quality differentiation (see e.g. Goldberg & De Loecker, 2014).

across contexts. We provide tools for measuring the performance of a bureaucracy and show that effectiveness affects the relative costs and benefits of different policies (see also Dehejia *et al.*, forthcoming).[11] We are not aware of prior studies that estimate treatment effects conditional on an unobserved characteristic such as effectiveness (see e.g. Heckman & Smith, 1997; Angrist, 2004, for discussion of the estimation of treatment effects conditional on observed characteristics).

2 Public Procurement in Russia

2.1 A decentralized system with centralized rules

Public procurement comprises 10 percent of Russia's non-resource GDP. In 1991, it created an extremely decentralized procurement system (see e.g. Enikolopov & Zhuravskaya, 2007). Each government entity has the legal authority to make its own purchases and there are no centralized purchases (such as framework contracts). Conversely, a federal law provides the legal framework for all procurement purchases above USD 35,000 for all levels of government (Yakovlev *et al.*, 2010).

We focus our analysis on electronic auctions—the most common vehicle, used for 53.5 percent of purchasing during our 2011–2016 data period—so as to study bureaucrats and organizations performing exactly the same task. Auctions are conducted through one of five independent web platforms. At the time of the auction, only the platform knows the identities of the bidders, making it possible to conduct auctions in which the bidding firms are anonymous to the procurers.

Appendix Figure H.1 traces out the steps involved together with the number of purchases that followed each path to contracting. The auction announcement is drawn up by a procurement officer. It contains technical details on the item(s) to be purchased (from clients), a maximum price, the required security deposit (between 0.5 and 5 percent of the maximum price), other participation requirements, and the auction date. Suppliers can then prepare a two-part application. The first part describes the good(s) that they are offering to fulfill the procurement order. The second part—which cannot be accessed by the procurers until the auction is concluded—contains information on the supplier itself (name etc.).

A five-member commission, including the purchasing bureaucrat and organization, oversees the purchase. They receive and evaluate the anonymized first part of each application before the auction. The purchasing bureaucrat directs the commission's review to deny applications from suppliers that are not accredited, cannot pay the deposit, or whose proposal does not comply with the requested specifications.[12] If only one supplier is approved, the auction is declared "not held" and its contract is drawn up at the maximum price. This is relatively common, occurring in 1.4 million (22 percent of) cases. If there are no approved applicants, the purchase is cancelled (13 percent of purchases).

If more than one supplier is approved, the auction is held. Approved suppliers are assigned a

[11] Our findings resonate with those of the first studies to compare experimentally identified program effects across branches of companies or private-versus-public status of the implementing agency (see Bold *et al.*, 2018; Allcott, 2015).

[12] The platform accredits suppliers that are not in a state of bankruptcy; do not have substantial unpaid taxes; and are not listed in a registry of suppliers who have violated procurement rules during the last two years.

number and remain anonymous. They log in to the platform and participate in a descending, open-outcry auction. Following the conclusion of the auction, the commission receives and reviews the second part of the applications. These contain the identifying information of the participants, but they cannot be linked to their bids. The commission checks the suppliers' accreditations, licenses, names, registration, and tax ID numbers. These are verifiable so there is little scope for subjective judgment at this stage. The contract is signed with the approved bidder who submitted the lowest bid.

2.2 The role of bureaucrats and organizations in procurement

The labor market for Russian procurement officers resembles that for private sector jobs. Interested individuals seek out educational and employment opportunities in decentralized markets as in the private sector, creating labor market churn from procurement officers' and their employers' job search.[13] The Russian government does not educate bureaucrats, nor does it operate a centralized civil service administration to recruit, train, or assign public servants to postings (Barabashev & Straussman, 2007). In all cases we are aware of, procurement bureaucrats are paid a flat salary.

Purchases are made for the public entity that pays for and uses the goods; an *organization*. It may, for example, be a school, hospital or ministry, at the municipal, regional or federal level. To make a purchase, the organization must work with a procurement officer—individual *bureaucrats*. Together, the organization and bureaucrat (the *procurers*) acquire the good according to the centrally set rules, and at the lowest possible price. Any policy goals the central government may have, such as influencing which types of goods or firms win contracts, manifest themselves in the rules followed by all procurers. Conditional on following those rules, procurers' only mandate is to pay the lowest possible price. For any given rules, the price paid is thus the primary measure of how effective procurers have been at implementing the government's procurement policy.

Bureaucrats can either be "in-house" (employees of the organization) or "external".[14] This means that we observe bureaucrats working with more than one organization (and vice versa) for two distinct reasons. The first is that bureaucrats change employers. The other is that external bureaucrats may conduct purchases with multiple organizations, and organizations may work with multiple external bureaucrats. On average, bureaucrats in our data are observed working with 5.2 organizations, and organizations with 4.8 bureaucrats. This high degree of churn is a powerful source of variation for this paper's empirical exercise.[15]

[13] Examples of private academies offering trainings on procurement include ArtAleks `http://artaleks.ru/` and the Granit Center `http://www.granit.ru/`. The primary prerequisites are a legal education, management experience, and knowledge of current procurement laws.

[14] Each regional authority sets rules dictating the type of bureaucrat used for each type of purchase, as defined by the maximum price and the nature of the item. External procurement agencies can be organized by a given authority (e.g. an education or health ministry), at the federal, regional, or municipal level. Part of the motivation for creating such agencies was to allow organizations purchasing similar goods to join forces and achieve lower prices. In practice, the decentralized management of procurement in Russia and coordination required means that joint purchases are very rare. Note that we control for the factors that authorities use to determine whether an in-house or external bureaucrat is used—the type of good and/or maximum allowable price of the contract—in our empirical analysis.

[15] Our setting features more turnover than would be observed in comparable private sector labor markets. German workers e.g. work at an average of 1.19 firms over the period 2002–2009 (authors' calculations based on Card *et al.*, 2013).

Since 2014, the division of labor between a procuring organization and an external procurement officer has been specified by law. The organization submits all technical documentation, and chooses and justifies the maximum price. The pair then together designate the commission to oversee the auction process. The bureaucrat manages consultations with specialists, collects information needed to design the tender, and works with the committee to conduct the first stage review, the auction, and the second stage review. The organization then signs the contract with the winner and verifies delivery. The same or a similar division of labor applies when in-house bureaucrats are used, and 2014.

2.3 Preferences for domestically manufactured goods

During our study period (2011–16), certain goods manufactured in Russia received a 15 percent bid preference for parts of each year. Where preferences are in place, if at least one bidder offers foreign-made goods and at least one offers locally manufactured goods, a bidder offering foreign-made goods only receives 85 percent of her final bid as the contract price.[16]

Each year from 2011 to 2014 a list of good categories for which a preference for domestic goods applied was drawn up.[17] The presidential order defining the list was passed in May or June and remained in effect until the end of the year, after which the preference ceased to operate until a new list had been created and approved the following year (except in 2015 and 2016, when the 2014 list was extended through 2016). Preferenced goods spanned many categories, including automobiles, clocks, various food products, medical equipment, pharmaceuticals, textiles and furs (see Table H.1). Procurers were required to publicly inform potential suppliers that the preference applied.

Our analysis of the bureaucratic apparatus's role in procurement performance restricts attention to the policy regime without preferences. In Section 6 we analyze impacts of the preferences.

3 Data and Measurement of Procurement Performance

Since 2011, a centralized procurement website (`http://zakupki.gov.ru/`) has provided public information about all purchases. We use data from this website on the universe of electronic auction requests, review protocols, auction protocols, and contracts from January 1, 2011 through December 31, 2016. The data cover 6.5 million auction announcements for the purchase of 21 million items. However, purchases of services and public works are idiosyncratic, and do not lend themselves to our approach to measuring performance, and so we remove them, resulting in a final sample of 16 million purchases of homogeneous goods. Table 1 describes the sample.

To use this data to evaluate procurement performance, we must overcome two challenges. First, we need very precise measures of the items being procured to use prices paid as our main measure of performance. Section 3.1 describes our text-based item measures. Second, prices are not the only

[16] When the law is active, preferences formally apply to goods for which "the cost of goods produced in the territory of Russia, Belarus, and Armenia exceeds 50% of the total cost". Incorrect reporting of origin country may occur, but we found no coverage to suggest that such manipulation is common.

[17] Preferences were first given to domestic manufacturers in 2008 to stimulate the economy during the financial crisis. The list of goods covered was slightly changed in 2009, before expiring completely on December 31, 2010. The government then adopted an annual approach to determining which goods were covered beginning in 2011.

outcome that matters in public procurement. Sub-section 3.2 describes the additional data we bring in to study bureaucrats' and organizations' impacts on spending quality. We round out this section by discussing how corruption affects our performance measures (Sub-section 3.3) and the additional purchase process data we use to study the correlates of procurement performance in Sub-section 3.4.

3.1 Prices as performance

Our main measure of performance is the price paid, holding constant the precise nature of the item being procured. Holding constant the item being procured is crucial to avoid conflating differences in prices paid with differences in the precise variety of item being procured. A great deal of previous research in economics has faced this challenge, but typically achieves within-category homogeneity at the cost of losing generality.[18] To avoid doing so, we use the text of the final contracts, in which the precise nature of the good purchased is laid out. We classify purchases into narrow product categories within which quality differences are likely to be negligible using text analysis methods (see also Hoberg & Phillips, 2016).

Our method proceeds in three steps. First, we transform the good descriptions in contracts into vectors of word tokens. Second, we use the universe of Russian customs declarations to train a classification algorithm to assign goods descriptions to a 10-digit Harmonized System product code, and apply it to the good descriptions in our procurement data. Third, for goods that are not reliably classified in the second step, either because the goods are non-traded, or because their description is insufficiently specific, we develop a clustering algorithm to group good descriptions into clusters of similar "width" to the categories from the second step. Details are in Online Appendix A.[19]

To complement this approach, we collect additional data on purchases of pharmaceuticals, an especially homogeneous category of goods (Bronnenberg *et al.*, 2015). Russia's government regulates the pharmaceutical market, compelling suppliers of certain drugs to register in a List of Vital and Essential Medicinal Drugs (LVEMD). This list includes information on each drug's active ingredient, i.e. international nonproprietary name (INN); the manufacturer's name and location; date of registration; and maximum price. Matching the LVEMD to our data, we can construct a barcode-level classification of pharmaceuticals.[20] The pharmaceuticals subsample is summarized in column (4) of Table 1.

[18] Broadly, three approaches have been taken: using hedonic regressions to estimate consumers' demand for and/or suppliers' costs of producing good attributes when rich attribute data is available (see e.g. Bandiera *et al.*, 2009); using product codes provided by e.g. customs agencies to partition goods (see e.g. Rauch, 1999); or restricting attention to products that are by nature especially homogeneous (Syverson, 2004).

[19] Online Appendix A also analyzes the sensitivity of our main findings to the choices made when developing our text analysis methodology. As Figure D.1 and Table E.4 show, the findings are robust.

[20] We use fuzzy string matching to combine the contract data on medicines with corresponding entries in LVEMD using each drug's international brand (trademark) name, active ingredient (INN), dosage, active units, concentration, volume, and units. We restrict the Pharmaceuticals Subsample to purchases of drugs we can match to LVEMD. Failure to match can arise if a medicine is not considered "essential" or because insufficient information is available in the procurement contract.

3.2 Spending quality

Sourcing inputs at low prices is the primary goal of public procurement,[21] but it is not the only outcome that matters. Contracts should not need to be unduly renegotiated or terminated, and goods should be delivered as specified, without delays. These outcomes reflect the quality of public spending and may conflict with the goal of achieving low prices, creating a multi-tasking problem for buyers. If this problem is severe, then we may misclassify bureaucrats and organizations as high-performing if they achieve low prices but this is offset by poor performance on spending quality.

To address this, we build direct measures of spending quality and use them as an additional outcome in our analysis. We combine six proxies for the quality of the non-price outcomes of a procurement purchase: the number of contract renegotiations, the size of any cost over-run, the length of any delays, whether the end user complained about the execution of the contract, whether the contract was contested and canceled, and whether the product delivered was deemed to be low quality or banned for use in Russia because it didn't meet official standards. The first five of these measures come from the zakupki data, while the last one is sourced from a civil society organization—clearspending.ru—that scrutinizes the government's spending and publishes infractions they detect.

We focus on these six measures as they capture outcomes of a procurement purchase as opposed to inputs into the process leading up to the award of a contract. These are events that happen after the contract is signed that may not be captured in the contract price, but which alter the benefit to the government of the purchase. As a result, they should be thought of as outcomes which bureaucrats and organizations may affect. To summarize them in a single number, we take the six and create an index of spending quality y_i as the average of the six quality proxies after standardizing each one to have mean zero and standard deviation one: $y_i = \frac{1}{6}\sum_{k=1}^{6}(y_i^k - \bar{y}^k)/\sigma^k$ (Kling *et al.*, 2007).[22]

3.3 Corruption

Both public procurement and Russia are associated with widespread corruption (Transparency International, 2016; Szakonyi, 2018). By its very nature, corruption is unobserved, and so we must take care to ensure that our measures of performance are not tainted by corruption. Corruption can lead to low quality goods being purchased at high prices. However, since our performance measure—the price paid conditional on the good—carefully controls for the precise nature of the good that is ultimately purchased, it captures both high prices and low item quality.[23] The reforms that introduced electronic procurement in Russia also imposed strict requirements on government customers whereby the final contract could only be ratified for the amount publicly disclosed on

[21] Article 1 of Federal Law 94 (FZ-94), which transformed the public procurement system in 2005, declares the aim of procurement as the "effective, efficient use of budget funds". The law also introduced minimum price as the key criterion for selecting winners for most types of selection mechanisms (Yakovlev *et al.*, 2011).

[22] We also use the first principal component of the five proxies and show that our results are similar. We prefer this index since it does not take a stand on placing higher weight on some proxies than on others (see Kling *et al.*, 2007).

[23] Note that it is important here that we use the item described in the final contract rather than the tender documents to capture leakage between what the tender documents specify and what is ultimately delivered.

the auction platform. Corruption therefore is hiding in plain sight, 'on-the-books' in the mounds of procurement data. Savvy journalists have built numerous investigations into the misuse of government funds by analyzing the publicly available zakupki dataset.[24]

The quality-adjusted price paid is an attractive measure of performance in the potential presence of unobserved corruption for a number of reasons. First, governments mandate that procurers target exactly this—the price paid for goods of specified quality. Second, quality-adjusted prices are the relevant metric when policy-makers decide which services can be offered given costs. Finally, both high prices stemming from a lack of effort or ability and high prices stemming from corruption represent transfers between taxpayers and bureaucrats and as such have similar welfare implications.

Of course, the underlying source of ineffectiveness may have welfare implications for higher-order efficiency or equity reasons.[25] However, the above arguments hold irrespective of whether quality-adjusted price differences are due to corruption or "intrinsic" ineffectiveness, and so in the model and empirical analysis below, we remain largely agnostic about their relative contributions. In Sub-section 5.5 we provide some evidence that corruption is probably not the primary driver of variation in bureaucratic effectiveness in Russia (see also Bandiera *et al.*, 2009).

3.4 Process measurement

In addition to measuring the performance outcomes described above, we also want to paint a detailed picture of the inputs bureaucrats and organizations provide in the procurement process. To do this, we exploit the richness of the zakupi procurement data, which contains details of the entire procurement process. This allows us to measure things such as the extent to which buyers rush at the end of the fiscal year (Liebman & Mahoney, 2017), the reservation prices buyers set, the number of items they bundle into purchases, the number of bidders who apply; are accepted; and bid in the auction, the competitiveness of the auction, the experience and types of products the buyers buy etc.

We supplement it with data from two sources. First, we use data from clearspending.ru on how well the process is run (these include whether identifier codes in the tender documents are correctly filled out, whether the names of the products in document headings are correct, whether sufficient time is provided to prepare a bid, whether the contract specifies the contractors correctly etc.). Second, we match firms in the procurement data to the Bureau Van Dijk's *Ruslana* database, which covers the vast majority of firms that file financial information. This allows us to measure the types of firms that bid on, and that win, contracts from different buyers. Table F.1 summarizes the large number of variables we use on procurers' purchasing processes.

[24] See for example Tom Bergin, and Stephen Grey:"Opaque Middlemen Exact High Price in Russia's Deals with the West." *Reuters*, December 19, 2014.

[25] Such consideration could for example arise if the source matters for whether ineffectiveness affects efficiency by changing which firms win government contracts, or if transfers to taxpayers and bureaucrats are valued differently for equity reasons. These possibilities present an important avenue for future research.

4 A Model of Procurement with Heterogeneous State Effectiveness

In this section we present a stylized model of public procurement. We model state effectiveness as costs imposed on potential sellers wishing to participate in public procurement and show how variation in these costs leads to variation in output—the prices paid, motivating our empirical analysis in Section 5. We also show how the introduction of bid preferences differentially affects procurement by bureaucracies with different levels of state effectiveness, patterns we test for in Sub-section 6.2.

4.1 Performance heterogeneity in a constant policy environment

Consider a pair of a bureaucrat and an end-user organization—jointly, a bureaucracy—wishing to purchase an item from a supplier through a second-price descending auction. State effectiveness affects the prices the government is able to achieve in two ways. First, by directly increasing suppliers' contract fulfillment costs $\bar{\theta}/\theta_i$. $\bar{\theta}$ is a common cost component with three parts: $\log(\bar{\theta})=\mathbf{X}'\boldsymbol{\beta}+\alpha_\theta+\psi_\theta$. $\mathbf{X}$ are observable attributes of the item and α_θ and ψ_θ are the costs of satisfying requirements stipulated by bureaucrats and organizations, respectively. These may include the date and place of delivery, the size of the order, and other requirements that directly affect the cost of fulfilling the contract. $\theta_i \geq 1$ is a firm-specific productivity term.

Second, bureaucrats and organizations indirectly affect prices by adding specifications α_c and ψ_c that affect the cost to firms of participating in the procurement process. These may include deposits required, the time granted to prepare bids, the clarity of the tender documents, bribes paid to enter the auction, and any other specifications affecting the cost of bidding, but not of fulfilling the contract.

In the first stage of the procurement process, two firms—one local and one foreign—observe the specifications $\{\mathbf{X},\alpha_\theta,\alpha_c,\psi_\theta,\psi_c\}$ and decide whether to pay a participation cost c_i to learn their productivity θ_i and enter the auction.[26] The foreign firm $i=F$ and the local firm $i=L$ differ in both their expected productivity and their participation costs. Productivities θ_i are independent and Pareto distributed with Pareto parameters δ_F and δ_L. Foreign firms have higher expected productivities ($\delta_F<\delta_L$)[27] but face higher participation costs: $c_i=\frac{\bar{\theta}}{1+\delta_i}-\frac{\bar{\theta}}{1+\delta_L}\sqrt{1-\alpha_c-\psi_c}$.[28] In the second stage, if only one supplier chose to enter the auction, she is awarded the contract at price $\bar{\theta}$. If neither supplier entered, the bureaucracy finds an outside supplier and awards her the contract at a price of $\bar{\theta}$.[29] Finally, if both suppliers enter, a descending, open-outcry auction takes place, which we approximate with a second-price sealed-bid auction (see e.g. Milgrom, 2004).

[26] We assume that firms do not know their productivity when they decide whether to enter the auction, as in Samuelson (1985). A more general approach would allow firms to have a signal of their productivity before deciding on entry as in Gentry & Li (2014). This significantly complicates the analysis, but the qualitative conclusions are the same. A sketch of such a model is available from the authors upon request.

[27] This fact is well established in the literature on international trade (see e.g. Bernard *et al.*, 2007)

[28] This functional form makes the expressions for profits and prices tractable. However, the qualitative results only require the participation costs to be increasing in α_c and ψ_c.

[29] A more realistic assumption might be that auctions in which no firms enter have to be re-run at some cost. Our assumption makes the model static, simplifying the exposition. The qualitative results are unlikely to depend on this choice since no firms entering is more likely for low-effectiveness buyers (since, as discussed below, firms weigh entry costs against expected profits from participation and low-effectiveness buyers impose higher entry costs), and so this channel only adds to the additional costs that low-effectiveness buyers create.

The suppliers choose their entry and bidding strategies to maximize expected profits. We outline the equilibrium here, relegating a detailed characterization and the proofs of propositions to Online Appendix B. Working backwards from the second stage, when both firms enter, it is a dominant strategy for bidders to bid their fulfillment cost since bidder valuations are independent (see e.g. Milgrom, 2004). The winner is the bidder with the lowest fulfillment cost and receives the contract at the other bidder's fulfillment cost. The participation decision depends on the size of the participation costs c_i. When participation costs are sufficiently small, both firms enter and the auction always takes place. For larger participation costs the equilibrium involves mixed strategies with entry probabilities q_i. We can summarize the equilibrium in the following proposition:

Proposition 1. *In the Nash equilibrium of the auction, the bidders, $i \in \{F,L\}$ enter with probabilities $q_i = \sqrt{\kappa(1-\alpha_c-\psi_c)}$, where $\kappa = \min\left\{[(1+\delta_F+\delta_L)/(1+\delta_L)]^2, 1/(1-\alpha_c-\psi_c)\right\}$. Expected log prices are*

$$\mathbb{E}[\log(p)] = \log(\bar{\theta}) - \tfrac{q_F q_L}{\delta_F+\delta_L} = \mathbf{X}'\boldsymbol{\beta} - \tfrac{\kappa}{\delta_F+\delta_L} + \tilde{\alpha} + \tilde{\psi}, \tag{1}$$

where $\tilde{\alpha} = \alpha_\theta + \frac{\kappa}{\delta_F+\delta_L}\alpha_c$, and $\tilde{\psi} = \psi_\theta + \frac{\kappa}{\delta_F+\delta_L}\psi_c$. In equilibrium

1. *Bureaucracies that impose higher contract fulfillment costs α_θ, ψ_θ pay higher prices for otherwise identical goods.*
2. *Bureaucracies that impose higher participation costs α_c, ψ_c pay higher prices for otherwise identical goods, and also attract fewer bidders to auctions they run.*

Equation (1) shows how prices vary with with the costs imposed by bureaucrats ($\tilde{\alpha}$) and organizations ($\tilde{\psi}$) managing the procurement process, and forms the basis of our empirical approach.

4.2 Policy change with heterogeneous state effectiveness: bid preferences

We now study the impact of introducing bid preferences favoring the locally producing bidder L. If the lowest-bid, winner of the auction is foreign, the contract price will only be $p = \gamma b_L$, where $\gamma < 1$, while a local winner receives the undiscounted $p = b_F$. Otherwise the auction protocol is unchanged. Preferences make it optimal for bidder F to shade so that her contract price should she win is equal to her fulfillment cost $b_F = \bar{\theta}/\gamma\theta_F$. However, when her shaded bid would have no chance of winning ($\theta_F < 1/\gamma$), she drops out and the contract is awarded to bidder L.

The effects on prices depend on the balance of four effects. First, the penalty mechanically lowers prices in auctions with foreign winners. Second, local bidders, who are less productive on average, are advantaged in the auction, raising prices.[30] Third, since foreign bidders are less likely to win auctions, they are less likely to participate. Fourth, local bidders are emboldened to enter by their higher chance of winning the contract. The interesting cases arise when the preferences are strong enough that the effect on L's entry decision is considerable, but not so large as to make

[30] There is extensive evidence that exporters are more productive than other firms, see e.g. Bernard *et al.* (2007).

it very unlikely F can win the auction. Formally, we focus on the case when $\gamma^{-\delta_F} > 1 - \log\left(\gamma^{\delta_L}\right)$.[31] In this case, introducing bid preferences has heterogeneous effects depending on the effectiveness of the bureaucracy that we summarize in the following proposition:

Proposition 2. *When $\gamma^{-\delta_F} > 1 - \log\left(\gamma^{\delta_L}\right)$, the introduction of bid preferences has different effects on three groups of bureaucracies differing in their effectiveness.*

1. *For bureaucracies with $\alpha_c + \psi_c \leq \underline{c}$, prices rise, the expected number of bidders is unchanged, and the probability that bidder L wins the contract at auction increases;*
2. *For bureaucracies with $\underline{c} < \alpha_c + \psi_c \leq \bar{c}$, prices rise, the expected number of bidders falls, and the probability that bidder L wins the contract at auction decreases;*
3. *For bureaucracies with $\bar{c} < \alpha_c + \psi_c$, prices fall, the expected number of bidders increases, and the probability that bidder L wins the contract at auction increases. The probability that bidder L wins the contract at auction increases by more than in case 1.*

The thresholds $\underline{c}$ and $\bar{c}$ are defined by

$$\underline{c} = 1 - \left(\frac{1+\delta_L}{1+\delta_F}\left(1 - \gamma^{1+\delta_F}\right) + \frac{1+\delta_L}{1+\delta_F+\delta_L}\gamma^{1+\delta_F}\right)^2 \qquad \bar{c} = 1 - \left(\frac{1+\delta_L}{1+\delta_F+\delta_L}\gamma^{\delta_F}\right)^2.$$

For effective bureaucracies that impose low participation costs on potential bidders ($\alpha_c + \psi_c \leq \underline{c}$), preferences do not deter foreign firms from entering the auction, but the local bidder is more likely to win, and the less aggressive bidding by the foreign bidder raises expected prices. For bureaucracies with intermediate effectiveness ($\underline{c} < \alpha_c + \psi_c \leq \bar{c}$), foreign bidders no longer find it profitable to enter. Since only the local bidder enters, the auction does not take place and the local firm gets the contract at the maximum price $\bar{\theta}$. Finally, when bureaucracies impose high participation costs ($\bar{c} < \alpha_c + \psi_c$), the increase in bidder L's willingness to enter is larger than the decrease in bidder F's willingness to enter, increasing the probability of both bidders entering and the auction taking place, lowering expected prices. Moreover, the entry effect is larger than the increase in prices caused by the changes in the bidding behavior in the auction, resulting in an overall decrease in expected prices.

Proposition 2 makes three predictions about heterogeneity in the impact of bid preferences. First, bureaucracies that pay higher prices when there are no bid preferences—which Proposition 1 shows is associated with higher participation costs—should experience price *decreases*, while bureaucracies that pay lower prices absent the bid preferences experience price *increases*. Second, the average number of participants in procurement processes should increase for bureaucracies that pay higher prices when there are no bid preferences. Third, we should see that the probability that an auction is won by a bidder offering to supply locally manufactured goods increases by more for bureaucracies that pay higher prices when there are no bid preferences. These are the patterns we test for in Sub-section 6.2

[31] Essentially, this condition requires that δ_L not be too much larger than δ_F. If this is violated, even with the preferences, the probability the local bidder wins is still very small and so when there is an auction the foreign bidder still wins but has her bid penalized lowering final prices even when the bureaucracy is very effective.

5 How Important is a Good Bureaucracy?

In this section we estimate the extent to which procurement effectiveness can be attributed to the individuals and organizations in the bureaucracy. We extend the method pioneered by Abowd *et al.* (1999) exploiting *switchers*—bureaucrats who make purchases with multiple organizations, and organizations who make purchases with multiple bureaucrats—for identification.

5.1 Identifying the effectiveness of individuals and organizations

We start by showing that bureaucrat-organization switches identify the causal impact of the individual in charge and the organization he or she works with on the purchase price. We use an event study analysis to study the dynamics of prices paid by organizations around the time that they switch the bureaucrat they work with. This happens frequently in Russia. As detailed in Table D.1, we observe 65,000 events in which organizations switch bureaucrats, with an average of 45 observations per event.

We define an event as chronological pairs of employment spells involving the same organization but two different bureaucrats. Figure 1 shows how prices change around such events. Each of the two employment spells is a sequence of at least two weeks less than 400 days apart in which a bureaucrat-organization pair makes purchases together. The horizontal axis displays event time, i.e. purchase weeks. The vertical axis displays the average quality-adjusted prices paid in a given week. The figure shows the evolution of prices paid by buyers starting with a bureaucrat in the top or bottom quartile of effectiveness, which we define using purchases made by the bureaucrats involved in the event, but that are not included in the event itself. Specifically, we use the average quality-adjusted price they pay in purchases made for *other* organizations they work with during the half-year that the spell ends (for the earlier spell) or starts (for the later one), akin to Card *et al.* (2013).[32]

Four key findings emerge from Figure 1. First, quality-adjusted prices paid change sharply, and in the expected direction, precisely when an organization switches to a less or more effective bureaucrat. The estimates suggest that an organization switching from a worst quartile-bureaucrat to a best quartile-bureaucrat on average experiences an 18 percent decrease in prices paid. Second, the figure shows no sign that performance is improving in organizations that subsequently switch to a better bureaucrat, and vice versa.[33] This suggests that drift in effectiveness and switches are uncorrelated. Third, we do not see a systematic dip or spike in performance before a bureaucrat switch, indicating that switches are not driven by temporary improvements or deteriorations in performance. Fourth, the price changes associated with switching bureaucrats appear symmetric:

[32] We quality-adjust prices by regressing them on log quantity, good fixed effects, month fixed effects, interactions between 2-digit HS product categories, years, regions, and lot size (as detailed in the next sub-section). Table D.1 highlights that the number of switches used to construct each quartile-to-quartile plot in Figure 1, and the average number of purchases observed for each bureaucrat-organization involved, are symmetric both around the events, and across quartile-to-quartile plots. The table also displays the average number of calendar weeks between each purchase week on the x-axis of Figure 1.

[33] More formally, of the sixteen groups formed by the possible trajectories between the four quartiles of bureaucrat effectiveness, we are unable to reject the null of no pre-trend in ten. Of the remaining six, five have pre-trends that point in the opposite direction of this concern.

organizations switching from a bureaucrat in the best quartile of average prices to one in the worst quartile experience a price *increase* of similar magnitude to those switching in the other direction. In Online Appendix D we show that these patterns are robust to changing a series of choices made in constructing the event studies.

Taken together, the evidence in this sub-section suggests that the thousands of quasi-experiments arising from organizations switching bureaucrats and vice versa in Russian public procurement can be used to estimate specific procurers' causal impact on performance, and that this impact is large.[34]

5.2 Variance decomposition method

We now aggregate the causal effects of specific bureaucrats and organizations from Sub-section 5.1 into estimates of the share of sample-wide variation in procurement performance bureaucrats and organizations as a whole explain. To do so we first extend the method pioneered by Abowd *et al.* (1999) to study wage dispersion in the private sector, and then show how to correct for sampling error to form predictions of the impact of each bureaucrat and organization on prices paid. We use these predictions to examine the mechanisms through which procurers affect prices in Sub-section 5.5 and how bureaucratic effectiveness impacts the way policies map into public sector output in Section 6.

We model the price paid for item i procured by organization j and bureaucrat $b(i,j)$ as a function of item attributes $\mathbf{X}_i$, a price premium due to the bureaucrat $\tilde{\alpha}_{b(i,j)}$, and a price premium due to the organization $\tilde{\psi}_j$. As the theoretical framework in Section 4 shows, these price premia can be thought of as a reduced form for the impact on prices of the participation costs that bureaucrats and organizations of different levels of effectiveness impose on suppliers. The log unit price paid for an item is

$$p_i = \mathbf{X}_i \boldsymbol{\beta} + \tilde{\alpha}_{b(i,j)} + \tilde{\psi}_j + \varepsilon_i \qquad (2)$$

To control flexibly for the item being purchased, $\mathbf{X}_i$ includes log quantity, good and month fixed effects, and interactions of 2-digit HS product categories, years, regions, and lot size.[35]

Identifying the bureaucrat and organization premia is made possible by the switches we documented in Sub-section 5.1. As Abowd *et al.* (2002) show, individual and organization effects are only identified *within* sets of organizations connected by individuals moving between them.[36]

[34] We also construct analogous event study figures for organizations and bureaucrats switching from purchasing one type of *good* to another. The results are in Figure D.4. Each event study shows the same general patterns as in Figure 1.

[35] By lot size we mean the maximum allowable price for all items to be purchased in the auction. We divide this price into bins to allow our estimates of effectiveness to capture the impact on prices of the procurers' choice of the exact maximum price posted. The interactions help address e.g. concerns that systematic spatial variation in the average prices of different types of goods—Russian regions are highly heterogeneous (Enikolopov & Zhuravskaya, 2007; Yakovlev & Zhuravskaya, 2014)—in combination with differences across procurers in items purchased, confound our estimates of effectiveness. Hereafter we refer to the goods categories constructed using the method described in Sub-section 3.1 as "goods".

[36] More precisely, within each connected set s containing $N_{b,s}$ bureaucrats and $N_{j,s}$ organizations, we can identify at most $N_{b,s}+N_{j,s}-1$ linear combinations of bureaucrat and organization fixed effects. In fact, we estimate models with three sets of high-dimensional fixed effects, for bureaucrats, organizations, and goods (the models also contain month dummies to control for common time trends, but there are few enough of these month effects such that "month-connectedness" is not an issue). To our knowledge, identification results for models with more than two sets of fixed effects are not yet available (Gaure, 2013), however our focus is on the estimates of only two of the three

However, such switches do not connect all bureaucrats and organizations that conduct procurement in Russia. Our data contain 616 connected sets. This relatively large number comes about for several reasons. First, focusing on bureaucrats performing a *single task*, rather than comparing many types of workers through their wages—the approach taken in existing related work—limits connectedness. Second, workers change employers less often in the public than in the private sector. Finally, the decentralized nature of Russian procurement means that some geographically remote organizations do not have bureaucrat links to other organizations.

To form our Analysis Sample, we focus on connected sets containing at least three bureaucrats and organizations after we make the following restrictions. We remove any procurer pairs that only ever occur together (as in this case it is not possible to distinguish bureaucrat and organization effects), and similarly for bureaucrat-good pairs and organization-good pairs as well as any levels of our control fixed effects that only appear once in the data. We also require that all bureaucrats and organizations in the Analysis Sample make at least five purchases. Table 1 compares the full sample and the Analysis Sample. The organizations in the Analysis Sample are less likely to be federal, but their purchases are of similar size and quantity to those in the full sample.[37] Overall the sample we use for analysis appears to be fairly representative.[38]

To proceed, we normalize the $\tilde{\alpha}_{b(i,j)}$ and $\tilde{\psi}_j$ to have mean zero in each connected set and augment (2) to include intercepts $\gamma_{s(b,j)}$ for each connected set:

$$p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i \tag{3}$$

In Online Appendix C, we show that while the $\tilde{\alpha}$s and $\tilde{\psi}$s in equation (2) are not identified, the αs, ψs and γs in equation (3) are. These are related to the underlying bureaucrat and organization effects as follows: $\alpha_b = \tilde{\alpha}_b - \overline{\alpha}_{s(b)}$, $\psi_j = \tilde{\psi}_j - \overline{\psi}_{s(j)}$, and $\gamma_{s(b,j)} = \overline{\alpha}_{s(b,j)} + \overline{\psi}_{s(b,j)}$, where $\overline{\alpha}_{s(b)}$ is the mean bureaucrat effect in the connected set containing bureaucrat b, and similarly $\overline{\psi}_{s(j)}$ is the mean organization effect in organization j's connected set.[39]

We can use equation (3) to decompose the variance of prices into its constituent parts using

$$\begin{aligned} \mathsf{Var}(p_i) = {} & \mathsf{Var}\Big(\alpha_{b(i,j)}\Big) + \mathsf{Var}(\psi_j) + 2\mathsf{Cov}\Big(\alpha_{b(i,j)}, \psi_j\Big) \\ & + 2\mathsf{Cov}\Big(\alpha_{b(i,j)} + \psi_j, \gamma_{s(b,j)} + \mathbf{X}_i\boldsymbol{\beta}\Big) + \mathsf{Var}\Big(\gamma_{s(b,j)} + \mathbf{X}_i\boldsymbol{\beta}\Big) + \mathsf{Var}(\varepsilon_i) \end{aligned} \tag{4}$$

all of which can be identified. Since $\mathsf{Var}\Big(\alpha_{b(i,j)}\Big)$ and $\mathsf{Var}(\psi_j)$ are variances within connected sets, they are lower bounds on the underlying variances of bureaucrat and organization effects.[40]

---

dimensions–the bureaucrat and the organization effects.

[37] We find that bureaucrats at federal agencies switch jobs less often, since there is more scope for both horizontal and vertical mobility within these larger organizations.

[38] In Table E.3 we show that our results are robust to using only the largest set of connected organizations. Table E.2 compares the Analysis Sample to its largest connected set.

[39] Faced with this issue, previous work on firms and workers has tended to restrict attention to the largest connected set, normalizing an arbitrary firm effect to 0, and estimating unconditional variances. An exception is Card *et al.* (2016).

[40] Formally, $\mathsf{Var}(\tilde{\alpha}_b) \equiv \mathbb{E}\left[\mathsf{Var}(\tilde{\alpha}_b|s(b))\right] + \mathsf{Var}\left(\mathbb{E}[\tilde{\alpha}_b|s(b)]\right) = \mathsf{Var}(\alpha_b) + \mathsf{Var}\left(\mathbb{E}[\tilde{\alpha}_b|s(b)]\right) \geq \mathsf{Var}(\alpha_b)$. Similarly, $\mathsf{Var}(\tilde{\psi}_j) = \mathsf{Var}(\psi_j) + \mathsf{Var}\left(\mathbb{E}\left[\tilde{\psi}_j|s(j)\right]\right) \geq \mathsf{Var}(\psi_j)$.

We can obtain unbiased estimates of procurer effects using OLS under the assumption that the residuals in (3) are uncorrelated with the identity of the bureaucrat or organization making a purchase (conditional on $\mathbf{X}_i$). There are two principal reasons this might not be the case. First, it could be that prices change around the time bureaucrats move across organizations or vice versa, for reasons unrelated to the switch. However, as shown in Sub-section 5.1, we do not see evidence of such pre-trends.

Second, equation (3) assumes that prices are log-linear in the procurer effects—an assumption about the degree of complementarity between the bureaucrat and the organization working on a purchase and associated sorting patterns. If the model is misspecified, then the omitted complementary terms are a component of the residuals in (3).[41] These complementarities may be correlated with the identity of the bureaucrat or organization making a purchase if, for example, organizations recruit bureaucrats who specialize in particular goods. Then estimates from (3) would recover a mixture of the true effect and the average complementarity of bureaucrat-organization matches.

Such sorting would imply that organizations switching from bureaucrats who pay high prices to bureaucrats who pay low prices enjoy larger decreases than the price increase suffered from moving in the opposite direction. Organizations hiring a low-price bureaucrat benefit from *both* a lower average price and an improved match effect, and organizations hiring a high-price bureaucrat lose from the lower average price but benefit from an offsetting improved match effect. We see no evidence of such patterns in Figure 1.[42] The symmetry of the event study evidence indicates that omitted complementarities are unlikely to bias our estimates. Online Appendix E.1 provides further tests for misspecification.

We use a large sample of public procurers, but nevertheless, our estimates need not be consistently estimated, even if they are unbiased. Consistency of the estimated fixed effects requires that the number of observations *on each group* tends to infinity (Lancaster, 2000). Our data contains 284,710 bureaucrat-organization pairs and an average of 40 observations per pair, so we cannot be confident a priori that the error in the bureaucrat and organization effect estimates has asymptoted to zero, particularly for the less frequently observed pairs. Moreover, since we are estimating two sets of fixed effects, the problem is compounded if the network features too few switches. Such *limited mobility bias* results in a spurious negative correlation between the two dimensions of estimated fixed effects (Andrews *et al.*, 2008). Each connected set in our data is densely connected—we observe bureaucrats working with 5.2 organizations on average, and organizations with 4.8 bureaucrats—but limited mobility bias may still be a concern.[43]

We address these sampling error issues in three ways. First, we bootstrap to estimate standard

[41]Our identifying assumption does not rule out effective bureaucrats and organizations matching with each other.

[42]If anything, the price decreases when organizations switch to lower average-price bureaucrats in Figure 1 appear slightly *smaller* than the corresponding increases when organizations switch to higher average-price bureaucrats.

[43]Moreover, in 76% of organizations, *all* the bureaucrats they work with are switchers (work with multiple organizations). Similarly, for 94% of bureaucrats, all the organizations they work with are switchers. This is reassuring since it is these switches that allow us to identify their effects.

errors for our variance decomposition.[44] Second, we take a non-parametric, split-sample approach to estimating the variance components in (4), akin to Finkelstein *et al.* (2016) and Silver (2016). We randomly split our sample in half, stratifying by bureaucrat-organization pair. We then estimate equation (3) separately on each sample, yielding two estimates ($k=1,2$) for each bureaucrat ($\hat{\alpha}_b^k$), organization ($\hat{\psi}_j^k$), and connected set ($\hat{\gamma}_s^k$) effect. Both estimates are estimated with error, but the errors in the two estimates should be uncorrelated, so we can create split-sample estimates of the variance decomposition terms as follows: $\widehat{\mathsf{Var}}^{SS}(\alpha_b)=\mathsf{Cov}(\hat{\alpha}_b^1,\hat{\alpha}_b^2)$, $\widehat{\mathsf{Var}}^{SS}(\psi_j)=\mathsf{Cov}\left(\hat{\psi}_j^1,\hat{\psi}_j^2\right)$, $\widehat{\mathsf{Var}}^{SS}(\gamma_s)=\mathsf{Cov}(\hat{\gamma}_s^1,\hat{\gamma}_s^2)$, and $\widehat{\mathsf{Var}}^{SS}(\alpha_b+\psi_j)=\mathsf{Cov}\left(\hat{\alpha}_b^1+\hat{\psi}_j^1,\hat{\alpha}_b^2+\hat{\psi}_j^2\right)$.

Third, we adopt two shrinkage approaches to create predictions of each bureaucrat and each organization effect. The variance in our estimated fixed effects comes from two sources: the true, signal variance in bureaucrats' and organizations' effects, σ_α^2 and σ_ψ^2 respectively, and sampling error with variances σ_μ^2 and σ_ω^2. Bootstrapping the estimation of equation (3) yields estimates of the variance of the sampling error which we use to perform a standard shrinkage procedure for the bureaucrat and organization estimates separately, as is common in studies of teacher value-added (see e.g. Kane & Staiger, 2008; Chetty *et al.*, 2014).[45] To address limited mobility bias, we extend the shrinkage approach used in existing work to explicitly account for the correlation between the estimation errors of the bureaucrat and organization effects. Our bootstrap also provides estimates of the covariance of all the estimation errors which we use to form minimum mean-squared error predictions of the full vector of bureaucrat and organization effects.[46] We label this method "covariance shrinkage". It yields our preferred estimates of the price variance decomposition in equation (4).[47]

5.3 Results

Table 2 shows results from our variance decomposition (4). The first column shows estimates of the standard deviations using the raw fixed effects estimates from equation (3), while estimates from the

[44] We construct partial residuals $\epsilon_i = p_i - \mathbf{X}_i\hat{\boldsymbol{\beta}}$ and randomly resample the residuals, stratifying by bureaucrat-organization pair to preserve the match structure of the observations. We then re-estimate the bureaucrat and organization effects. We repeat this procedure 100 times, and use the distribution of the estimates to compute standard errors. This procedure does not fully account for uncertainty arising from the data's match structure and finite sample correlations between bureaucrat and organization assignment and $\mathbf{X}$, but is computationally feasible.

[45] Formally, we find $\lambda_b=\operatorname{argmin}_{\tilde{\lambda}}\mathbb{E}\left[\alpha_b-\tilde{\lambda}\hat{\alpha}_b\right]=\sigma_\alpha^2/\left(\sigma_\alpha^2+\sigma_{\mu_b}^2\right)$, and analogously for λ_j. Our shrinkage estimators replace these terms with their sample analogues $\hat{\alpha}_b^{Sh}=\hat{\lambda}_b\hat{\alpha}_b$ and $\hat{\psi}_j^{Sh}=\hat{\lambda}_j\hat{\psi}_j$.

[46] Formally, we seek the linear combination of the full vector of fixed effects that minimizes the expected mean-squared error of the predictions. Denoting the vector of estimated bureaucrat and organization fixed effects by $\hat{\boldsymbol{\theta}}$ and the matrix of weights by $\boldsymbol{\Lambda}$, the objective is $\min_{\boldsymbol{\Lambda}}\mathbb{E}\left[\left(\boldsymbol{\theta}-\boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)'\left(\boldsymbol{\theta}-\boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)\right]$, which has solution $\boldsymbol{\Lambda}^*=\mathbb{E}\left[\boldsymbol{\theta}\hat{\boldsymbol{\theta}}'\right]\left(\mathbb{E}\left[\hat{\boldsymbol{\theta}}\hat{\boldsymbol{\theta}}'\right]\right)^{-1}$. Replacing the expectations with their sample analogues yields the shrinkage matrix $\hat{\boldsymbol{\Lambda}}^*=diag\left(\hat{\sigma}_\alpha^2,\hat{\sigma}_\psi^2\right)\left(diag\left(\hat{\sigma}_\alpha^2,\hat{\sigma}_\psi^2\right)+\boldsymbol{\Sigma}\right)^{-1}$, where $\boldsymbol{\Sigma}$ is the covariance matrix of the bootstrap estimates and $diag\left(\hat{\sigma}_\alpha^2,\hat{\sigma}_\psi^2\right)$ is the diagonal matrix with $\hat{\sigma}_\alpha^2$ in entries corresponding to entries for bureaucrats in $\boldsymbol{\theta}$ and $\hat{\sigma}_\psi^2$ in entries corresponding to organizations.

[47] We thus use "covariance shrunk" estimates in our analysis of the determinants of bureaucratic capacity in Sub-section 5.5 and the analysis of the effects of procurement policy changes in Section 6. For computational reasons, we perform covariance shrinking separately in each connected set. Since the estimated fixed effects are all normalized to be mean zero within each connected set and by definition the observations are unrelated across connected sets, this is without loss.

split-sample approach are in Column (3). The corresponding standard errors are in columns (2) and (4). The results from the shrinkage and covariance shrinkage methods are in columns (5) and (6).

Three key findings emerge. First, bureaucrats and organizations are each important determinants of policy performance. After controlling for the good being purchased and the month of the purchase, the standard deviation of log unit prices is 1.280. Compared to this, the bureaucrat fixed effects have a standard deviation of 0.795 and the organization fixed effects' standard deviation is 0.931. The split-sample estimates in Column (3) are similar. The shrinkage methods in columns (5) and (6) deliver slightly smaller estimates of the bureaucrat and organization variances, but even the covariance shrinkage estimates imply large effects of bureaucrats and organizations on policy performance.

Second, the covariance shrinkage method shown in Column (6) appears to best deal with the finite-sample inconsistency of our estimates. The fixed effects, split-sample, and shrunk estimates all yield a negative estimate of the correlation between bureaucrat and organization effects.[48] However, our covariance shrinkage approach yields a more plausible estimate of the correlation of 0.297.[49] As a result, the covariance shrunk estimates of share of the variation in performance explained by bureaucrats and organizations—20 and 28 percent percent respectively—are our preferred estimates.

Third, the combined importance of bureaucrats and organizations for policy performance is large. Our estimates of the within-connected-set standard deviation of the combined bureaucrat and organization effects are consistent across the four methods, ranging from 0.66 for the split-sample approach down to our preferred estimate of 0.50 for the covariance-shrunk estimates—39 percent of the standard deviation of log unit prices. Overall, our estimates imply that bureaucrats and organizations jointly explain a remarkably large share of the variation in procurement effectiveness in Russia, of which about half in turn is due to bureaucrats and half to organizations.

The large estimates in Table 2 have correspondingly dramatic implications for the scope of potential savings from improving the effectiveness of the bureaucracy. To illustrate the magnitude, we can consider simple counterfactual bureaucracies in which bureaucrats and/or organizations with low effectiveness are improved, for example through changes in recruiting, training of existing bureaucrats, or improved organizational management. Our estimates indicate that increasing the effectiveness of the lowest quartile of bureaucrats to the 75th percentile would save the Russian government 4.5 percent of annual procurement expenses. Moving all bureaucrats *and* organizations below 25th percentile-effectiveness to 75th percentile-effectiveness would save the government 12.1 percent of procurement expenditures. Annual procurement expenses are USD 86 billion, so this implies savings of USD 10 billion each year, or 0.7 percent of non-resource GDP (see Table H.2)—roughly one fifth, for example, of the total amount spent on health care in 2013 and 2014.[50]

[48]The same is found in many studies applying the AKM method to private sector wages. This led Andrews *et al.* (2008) to show that the AKM-estimated covariance term is downward biased (see Sub-section 5.2) and to suggest a parametric correction. However, this parametric correction relies on homoskedasticity of the residuals, an unappealing requirement in our setting (see also Card *et al.* (2013)).

[49]Recall that such assortative matching does not violate the no-sorting-on-match-effects assumption discussed above.

[50]Appendix Figure E.1 shows how these counterfactuals affect the distributions of effectiveness. Online Appendix

5.4 Robustness

We interpret the results in the previous sub-section as capturing the total, causal contribution of bureaucrats and organizations to the Russian state's effectiveness in procuring off-the-shelf goods. But are we adequately controlling for the precise item being purchased? And while prices paid are the primary metric of procurement effectiveness (see Sub-section 3.2), they are not the only one—what about spending quality?

Like-for-like comparisons If our goods classification based on contract texts is inaccurate, our estimates will conflate the true effects on prices with differences across bureaucrats and organizations in products bought. To probe this concern, we first show that our findings are similar in a sub-sample of goods that is by nature homogeneous—pharmaceuticals (see also Syverson, 2004; Bronnenberg *et al.*, 2015). We create barcode-level bins for pharmaceuticals as described in Sub-section 3.1 and make the same connectivity restrictions as in the full sample to create an analysis sample. Columns (4) and (5) of Table 1 summarize the sample. Table 3 presents the results of re-estimating (3) on the pharmaceuticals sample. Naturally, since the sample is more homogeneous and our barcode product categories are very precise, the share of the variation in prices explained by the good fixed effects is larger than in the broader sample. However, of the remaining variation in policy performance, all but the covariance shrinkage method attribute 30–40 percent to the combination of bureaucrats and organizations.[51] This is strikingly similar to the 40 percent found in the broader analysis sample. This is also what our theoretical framework suggests we should see, since we model the fulfillment costs imposed by bureaucrats and organizations on suppliers as proportional costs.

Second, our results are robust to focusing on more homogeneous subsets of goods in our full sample. In Figure 2 we split the sample into quintiles of good homogeneity as defined by the commonly-used measure of scope for quality differentiation developed by Sutton (1998).[52] We then reestimate (3) on successive subsamples. As we move from right to left, we restrict the sample to more and more homogeneous goods. As expected, the overall variance of average prices paid, shown by the grey shaded areas, decreases with good homogeneity. However, as shown by the blue line, the estimated share of the variance explained by bureaucrats and organizations remains very similar across the columns. In Appendix Figure E.2 we repeat this exercise using an alternative measure of scope for quality differentiation developed by Khandelwal (2010) and find the same result.[53]

Third, the right-most bar in Figure 2 shows that the results from our variance decomposition exercise are also essentially unaffected if we restrict the sample to items the text-based classification

E.3 compares these magnitudes to other studies of individuals' and organizations' effects on output in other settings.

[51] The covariance shrinkage method is less reliable in this sample since we have an order of magnitude fewer observations per connected set (an average of 1,411 vs 18,407) in this sample than in the sample used in Table 2. Despite this, the covariance shrinkage method attributes 20 percent of the variation to bureaucrats and organizations.

[52] We are able to match 70 percent of the items assigned an 10-digit HS code in Step 2 of the text analysis method with the Sutton (1998) measure.

[53] Another possibility is that organizations endogenously respond to the effectiveness of bureaucrats by purchasing more/fewer, or different types of, goods. This would lead us to underestimate the true variance in procurer effectiveness.

method is confidently able to assign a 10-digit Harmonized-System product code to.[54]

These results reassure us both that our text analysis procedure accurately classifies purchases into homogenous categories and that our broad sample of products is appropriate.

Spending Quality As discussed in Sub-section 3.2, procurers' primary goal is to achieve low prices without sacrificing on item quality. However, prices are not the only procurement outcome that matters. We study a form of procurement where non-price goals are a priori less important than they are in services or public works contracts. Nevertheless, bureaucrats who procure off-the-shelf, manufactured goods may also face multitasking problems in balancing price against other objectives. If buyers who achieve low prices do poorly on other measures of performance, we may erroneously conclude that they are effective when a more comprehensive evaluation would not.

To investigate, we first repeat our analysis using our spending quality index as the outcome instead of prices. Table 4 shows the results. Two key findings emerge. First, time and product effects explain a far smaller share of the variation in spending quality than in prices. This is unsurprising insofar as production costs vary significantly across products. More interestingly, it does not appear that the contracting problems and delays captured by our spending quality index are concentrated among a subset of products, perhaps because we restrict our sample to a broad range of similarly homogeneous manufactured goods. Second, the four estimation methods from Table 2 reveal that a significant share of the variation in spending quality is driven by the procurers. The most conservative, covariance-shrinkage method, attributes 24 percent of the variation to bureaucrats and organizations. This is expected since the components of the spending quality index—particularly contract renegotiations and cost overruns—are outcomes buyers have scope to influence. However, bureaucrats and organizations explain a smaller share of the variation in spending quality than in prices.

To study the multitasking issue, Figure 3 shows the correlation between the bureaucrats' (Panel A) and organizations' (Panel B) covariance-shrunk price and spending quality effects. The panels show binned scatterplots together with a regression line fitted on the underlying data, and the correlation between the two effects is shown in the upper left corner. The figure reveals a strong, positive relationship between procurers' impact on the two outcomes and a fairly linear relationship between the two (correlations of 0.42 for bureaucrats and 0.48 for organizations): bureaucrats who achieve low prices also perform well on spending quality, and similarly for organizations. Additionally, in Appendix Table E.1 we re-estimate the variance decomposition including the spending quality outcomes as controls (despite them more properly being considered endogenous to the bureaucrat and organization making the purchase), and show that the results are essentially unchanged from our baseline specification in Table 2.

Overall, the results suggest that while bureaucrats and organizations clearly influence spending

[54] The algorithm developed in Step 2 of the procedure outlined in Sub-section 3.1 and Online Appendix A assigns a 10-digit code to 37 percent of the items in our analysis sample with high confidence. The remaining items in the Analysis Sample are also clustered into homogeneous bins, but we cannot confidently assign a pre-existing 10-digit code to these items.

quality, the multitasking issue is not severe. In our subsequent analysis we thus use bureaucrats' and organizations' estimated effect on prices paid—their primary, legislated target (Yakovlev *et al.*, 2011)—as our preferred measure of their performance.

5.5 What do effective bureaucracies do differently?

We now analyze what it is that distinguishes effective bureaucracies from their ineffective peers. Our data contain detailed information on the evolution of each of the 6.5 million procurement processes in the sample.[55] We construct 85 potential explanatory variables for bureaucrats, and 114 for organizations, which we summarize in Table F.1. There are seven categories of predictor variables: spending quality measures (6 variables for bureaucrats, 5 for organizations), and features of the purchase request (12); the bureaucrat and organization (12 and 42); the auction (19); participating suppliers (35); and the region (5). We investigate which of these co-vary with the estimated price- and spending quality-effectiveness of the implementing bureaucrat and organization.

To avoid overfitting and for the sake of parsimony, we use a LASSO procedure to first select 30 predictor variables.[56] We then regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3) (and vice-versa for the organization effects, the results for which are shown in appendix Figures F.1 – F.6).[57] The left panels of Figures 4 and 5 show coefficients from a series of bivariate regressions of the bureaucrat price effect (in Figure 4) and spending quality effect (in Figure 5) on each of the selected observables. The right panels show the LASSO coefficients (as crosses) and those from the multivariate regression of the procurer effect on all of the selected variables (as circles). To facilitate comparison, all variables are standardized to have unit standard deviation. The coefficients can thus be interpreted as the association between a one-standard deviation change in the predictor and the impact of the procurer on prices.

Several interesting patterns emerge. First, effective bureaucrats do not rush at the end of the fiscal year, a practice associated with wasteful spending (Liebman & Mahoney, 2017). They also set lower reservation prices, perhaps by soliciting accurate commercial information from trusted suppliers and established market players.[58] Ultimately, effective bureaucrats attract and admit a larger and more diverse pool of bidders, as emphasized by the theoretical framework in Section 4. For each auction, we calculate the fraction of the pool of potential bidders who participate, and for each bureaucrat we calculate the Herfindahl-Hirschmann Index (HHI) of the suppliers they work

[55] In addition to the process, contract, firm, and spending quality data described in Section 2, we here also use data on corruption and other measures of institutions across regions from Schulze *et al.* (2016) and the ICSID Russian Regions data.

[56] The procedure selects the smallest model with *at least* 30 predictors so the actual number varies slightly from figure to figure. Table F.1 shows pairwise coefficients from regressing price-effectiveness on each of the 411 potential explanatory variables we start out with. Figures F.3 and F.4 instead show results from using the LASSO procedure to select 60 instead of 30 predictors. The patterns in the findings are very similar to those described here.

[57] To account for small firms not being covered by the *Ruslana* data and the strong correlation between some of our variables, we also use an elastic net regularizer (a weighted average of LASSO and Ridge regression). Figures F.7 and F.8 show that the results are not sensitive to placing more weight on the Ridge regression.

[58] A common procedure is to apply officially standardized algorithms to market research on the average price paid for a given good. Research has shown the use of flawed market information to be one of the main ways ineffective bureaucrats drive up prices in Russia (see Sapozhkov, Oleg. "Krivye Putyi Goszakazchikov." *Kommersant*, April 12, 2019).

with, and we find that both measures are strongly correlated with bureaucrat performance.[59] An example comes from a purchase of winter boots for a Saratov orphanage. The bureaucrat overseeing the request disqualified a firm from participating in the subsequent auction on the grounds that its application did not contain information on the height of the firm's boots' sole and heel. Only two bids were ultimately submitted in the auction, and the orphanage ended up paying a price per boot less than 10 percent below the maximum price.

Second, effectiveness appears to be very embodied in the individual procurers doing the work. Of the four categories of predictors we consider—features of the auction request; the bureaucrat; the auction itself; and the participating suppliers—characteristics of the bureaucrat have most predictive power. More experienced bureaucrats—for example those who run more auctions—are more effective, consistent with them having a larger network of contacts with suppliers to draw on. They also have fewer procurement processes fail due to no suppliers applying to participate.[60] Finally, bureaucrats who specialize more in particular products (as measured by the HHI of the products they buy)—another measure of bureaucrat experience—are also more effective.

Third, effective bureaucrats also end up purchasing from particular types of suppliers. They buy from suppliers that specialize in the products requested, and in selling to government (as measured by contracts won from state-owned enterprises). Effective bureaucrats also avoid middlemen: they are less likely to buy from wholesalers and exporters, but more likely to buy from firms that import the product they are purchasing. Finally, their suppliers are less likely to have the same postal code, or even to come from the same region.

Fourth, the correlates of bureaucratic effectiveness are strikingly similar when we look at prices paid and spending quality. Twenty of the 32 strongest predictors of price effectiveness included in Figure 4 are also among the 32 strongest predictors of spending quality effectiveness in Figure 5.[61] This is particularly true for features of the bureaucrats themselves, where all the features selected for the price outcome also appear in the quality figure. This is not surprising since we saw in Subsection 5.4 that bureaucrats' price effectiveness is highly positively correlated with their spending quality effectiveness. A final observation worth making is that there are some notable variables among those that are *not* selected by the LASSO. In particular, the wide range of regional measures of corruption have very weak predictive power. It thus appears that variation in bureaucratic procurement effectiveness in Russia is not primarily due to variation in corruption.

We conclude from these findings that a key part of what makes procurers effective is their ability to reduce entry barriers to participation in procurement auctions, and to attract firms to their auctions. The findings are very similar when we look at the determinants of organization effectiveness

[59] We treat all winners of contracts for the same 2-digit product in the previous semester as the pool of potential bidders.

[60] We label the fraction of the bureaucrat's purchases where this does not occur their "success rate". Purchase failure is an uncommon outcome, but effectiveness and success being positively correlated also assuages a potential selection concern about only observing successful purchases in our main dataset.

[61] Eighteen of these have the same sign in both cases. The two features for which this is not the case, the total value of auctions run by the bureaucrat in the entire sample and the total value of auctions won by the supplier, have very small coefficients, particularly for quality.

and so for conciseness we relegate those figures to Appendix F. One interesting additional finding that does emerge is that the single strongest predictor of an organization's effectiveness is its overall performance score in independent surveys and evaluations conducted by the Federal Treasury. This suggests both that our measure of effectiveness is correlated with what the federal government considers to be important, and that this effectiveness could be measured independently by central governments and then used to set procurement policy, the subject we turn to next.

6 Policy Design with a Heterogeneous Bureaucracy

We saw in Section 5 that a large share of the overall variation in performance under a constant policy regime is attributable to bureaucratic agents' effectiveness. But in many organizations—especially in the public sector—increasing productivity directly, through human resource practices, can be infeasible or costly. Such enterprises can instead change their task assignment, better tailoring work protocols to their workforce. In this section we study the introduction of a different policy regime in Russian procurement—a change in the bureaucracy's tasks. We show that the introduction of bid preferences favoring local manufacturers successfully shifted contracts to domestic producers, without significant impacts on prices or spending quality overall. However, these average treatment effects mask dramatic heterogeneity across "good" versus "bad" procurers, suggesting that there is significant scope for tailoring policy design to the effectiveness of the implementing bureaucracy.

6.1 Overall impact of bid preferences for locally manufactured goods

Many governments use bid preferences to attempt to steer demand towards favored firms. The impact of such policies is theoretically ambiguous (see e.g. McAfee & McMillan, 1989), though empirical studies in contexts with high state capacity tend to find price increases and participation decreases (Marion, 2007; Krasnokutskaya & Seim, 2011; Athey *et al.*, 2013). In Russia's case, as in many others, bid preferences favor local manufacturers. Its preferences policy imposed a bid penalty of 15 percent on foreign-manufactured goods (see Sub-section 2.3). In 2011–2014, the preferences only came into effect in May or June each year. Moreover, the policy applied only to a subset of goods—a subset that varied from year to year.[62] We exploit this variation in a generalized difference-in-differences design, estimating

$$y_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \delta \text{Preferenced}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt} \tag{5}$$

where y_{igt} is the outcome in purchase i of good g in month t, Preferenced_{gt} is a dummy indicating that g is a treated good in the year month t falls within, and PolicyActive_t is a dummy indicating that the year's list of preferenced goods has been published. $\mathbf{X}_{igt}$ are the same controls we use in Section 5, but for clarity we separate out the good and month fixed effects, μ_g and λ_t. ε_{igt} is an error term we allow to be clustered by month and good. Because there must be a minimum of one bidder

[62] Preferenced goods spanned many categories, including automobiles, clocks, various food products, medical equipment, pharmaceuticals, and textile and furs (see Table H.1 for the full list).

in the auction offering a Russian-made good and a minimum of one bidder offering a foreign-made good for preferences to apply, our estimates should be interpreted as Intent-to-Treat (ITT) effects.

Following Cengiz *et al.* (2019), we also stack all the events (the preference list being published) to estimate an event study analog of equation (5) in a window starting three months before and ending four months after each year's preference list is published (ListMonth$_s$):

$$p_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \sum_{s=-3}^{4} \delta_s \text{Preferenced}_{gt} \times \mathbf{1}\{t - \text{ListMonth}_t = s\} + \varepsilon_{igt} \tag{6}$$

To estimate (5) and (6), we expand the Analysis Sample and Pharmaceuticals Sub-sample to also include purchases where bid preferences apply, and which were managed by bureaucrats and organizations in these samples. The samples are summarized in columns (3) and (6) of Table 1. In the Analysis Sample, we define Preferenced$_{gt}$ as a dummy equal to one if good g is on that year's list. Since pharmaceuticals are always on the list, for pharmaceuticals we instead define Preferenced$_{gt}$ as equal to one if the drug is manufactured both in Russia and abroad.[63]

The estimated event study coefficients δ_s are all close to zero and statistically indistinguishable from zero in the months leading up to the publication of the preference list. Figure 6 shows this for prices in the Analysis Sample. This finding lends credibility to our difference-in-differences design's identifying assumption of parallel trends. The figure also shows no evidence of anticipation of the publication of the preference list. Figure G.1 shows the evolution of the share of purchases for preferenced items around the date of the publication of the list and also shows no evidence that buyers are able to manipulate the timing of their purchases to avoid or take advantage of preferences.

The preferences policy achieves its primary goal: the good purchased is 14 percent more likely to be domestically manufactured when bid preferences are in effect.[64] We show this first result from estimating (5) in Column (7) of Table 5. Columns (1) to (6) establish that it does so at little to no cost. Participation declines somewhat in both the full sample and the pharmaceuticals sample. However, prices are unaffected on average and spending quality increases slightly in the full sample; in the pharmaceuticals sample prices decrease and spending quality decreases somewhat on average.[65] The limited or even beneficial overall impact on prices suggests that the policy's discouragement of foreign manufacturers is offset by a combination of encouragement of local manufacturers and the mechanical decrease that applies when the winning bidder supplies foreign manufactured goods.

These findings contrast with studies of similar preference policies in the U.S. (see e.g. Marion, 2007; Krasnokutskaya & Seim, 2011; Athey *et al.*, 2013). Our analysis in the next sub-section points towards a possible explanation: U.S. procurers are probably more effective on average than Russian procurers. We estimate impacts similar to those found in the U.S.—increased prices—when

[63] Several drugs in use in Russia are manufactured either only abroad or only domestically.

[64] In Column (7) of Table 5 we restrict the sample to purchases in which an auction takes place in order to be consistent with Column (7) of Table 6. We find an increase in the probability of a domestic producer winning the auction of similar magnitude in the full pharmaceuticals sample (results available from the authors upon request).

[65] Recall that a higher number implies worse spending quality.

preferences are implemented by Russian procurers of high effectiveness, but when procurers are ineffective, we find the opposite impact.

6.2 Bureaucratic performance heterogeneity under different policy regimes

The model in Section 4 implies that bid preferences will compress the procurement performance of the bureaucratic apparatus. Proposition 2 describes how the variation in the entry costs buyers impose on suppliers that drives bureaucracies' effectiveness can also lead to patterns of heterogeneity in the treatment effect of introducing bid preferences. Such a finding would have striking implications for procurement policy design across contexts. To test this proposition in our data, we now compare treatment effects among effective and ineffective buyers. Estimates of effectiveness (in the absence of bid preferences) come from our analysis in Section 5.

We extend (5) to estimate heterogeneous treatment effects as follows:

$$
\begin{aligned}
y_{igt} = & \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \theta_b \hat{\alpha}_b + \theta_j \hat{\psi}_j + \delta \text{Preferenced}_{gt} \times \text{PolicyActive}_t + \rho_b \text{Preferenced}_{gt} \hat{\alpha}_b \\
& + \rho_j \text{Preferenced}_{gt} \hat{\psi}_j + \eta_b \text{PolicyActive}_t \hat{\alpha}_b + \eta_j \text{PolicyActive}_t \hat{\psi}_j \\
& + \pi_b \text{Preferenced}_{gt} \times \text{PolicyActive}_t \hat{\alpha}_b + \pi_j \text{Preferenced}_{gt} \times \text{PolicyActive}_t \hat{\psi}_j + \varepsilon_{igt}
\end{aligned} \tag{7}
$$

Table 6 shows the results. The small negative average price effect from Sub-section 6.1 masks substantial heterogeneity in the impact of bid preferences across bureaucracies. Consistent with Proposition 2, we find that prices drop significantly more for bureaucrats who pay higher prices when there are no bid preferences (i.e., who have a higher $\hat{\alpha}_b$). Columns (1) and (4) of Table 6 show this stark pattern both in the full sample and the pharmaceuticals sample. As we return to in Sub-section 6.4, the estimated coefficient on Bureaucrat FE $\times$ Preferenced $\times$ Policy Active is large (in absolute value) in both samples, and especially so in the pharmaceuticals sample.

Consistent with the model, these price improvements are accompanied by increases in participation (columns (2) and (5) of Table 6) and do not come at the expense of spending quality (columns (3) and (6)). Column (7) shows that the increases in the probability of a domestic winner are also concentrated among the least effective bureaucrats.[66]

While we find support for all the model's predictions on heterogeneous effects of bid preferences by bureaucrat effectiveness, we do not see this for organization effectiveness $\hat{\psi}_j$.[67] The model offers

[66] In fact, Proposition 2 predicts a U-shaped relationship between the probability a domestic good is supplier and bureaucrat type. Panel B of Figure G.4 shows that this is indeed what we see. We do not see a similar pattern for organizations and in fact the negative coefficient in Column (7) is not picking up a strong pattern of smaller effects for less effective organizations (results available upon request).

[67] That is, we see price decreases that are largest for the least effective bureaucrats; changes in participation that are larger for the least effective bureaucrats; and a U-shaped relationship between the probability a domestic good is supplied and bureaucrat type (results available upon request). When we look at heterogeneity by organization effectiveness, we do not see evidence consistent with any of these predictions. In the pharmaceuticals sample the coefficient on Organization FE $\times$ Preferenced $\times$ Policy Active is in fact positive, but very imprecisely estimated. In the full sample the differential effect for effective organizations is positive and marginally significant, but small in magnitude and in particular much smaller than the opposite-signed effect for effective bureaucrats.

a potential explanation: heterogeneity of the effect of bid preferences is driven by participation costs, but differences in organization effectiveness may to a greater extent be due to contract fulfillment costs than participation costs.[68]

We next estimate a less parametric version of (7) by including separate triple-interaction terms for each decile of bureaucrat effectiveness $\hat{\alpha}_b$ and organization effectiveness $\hat{\psi}_j$:

$$y_{igt} = \sum_{k=1}^{10} \left\{ D_{kj} + D_{kb} \times \left(\rho_k \text{Preferenced}_{gt} + \eta_k \text{PolicyActive}_t + \pi_k \text{Preferenced}_{gt} \times \text{PolicyActive}_t \right) \right\} + \mathbf{X}_{igt}\beta + \mu_g + \mu_t + \varepsilon_{igt} \tag{8}$$

where D_{kj} and D_{kb} are indicators for organization j and bureaucrat b belonging to decile k of their respective distributions of effectiveness. We also extend the event study (6) to estimate effects separately by quartile of bureaucrat- and organization-effectiveness. In these, rather than normalizing the reference month (the month before the preference list is published) to zero, we normalize it to the baseline performance in each group to better highlight how different their performance was before the preferences are introduced, and how their performance converges as a result of the preferences.

The price decreases in Table 6 are concentrated among the least effective 60 percent of bureaucrats. Figure 7 shows this graphically. We see a clear pattern of larger price drops for ineffective bureaucrats in Panel A, with the estimated price effect of the policy decreasing throughout the six lowest deciles of bureaucrat effectiveness. The figure also shows more suggestive evidence of price *increases* when the policy is administered by effective bureaucrats.[69] The event studies in Panel B of Figure 7 help rule out potential confounds like mean reversion or differences in seasonality across different types of bureaucrats. The graph shows no discernible trends in prices before the introduction of bid preferences and then a marked divergence of prices paid by the two groups—high versus low effectiveness bureaucrats—after the introduction of preferences. These patterns provide compelling evidence that the estimates in Table 6 capture the causal differential of interest.

[68] In the model, buyers impose two types of costs on potential suppliers: fulfillment costs (α_θ for bureaucrats and ψ_θ for organizations) and participation costs (α_c and ψ_c). As Proposition 1 states, both costs affect prices at baseline (without bid preferences) in the same way (though with different coefficients on fulfillment and participation costs) and so we subsumed them into the composite terms $\tilde{\alpha}$ and $\tilde{\psi}$ that enter equation (1). By contrast, as Proposition 2 states, the heterogeneity of the effects of bid preferences is governed by the participation costs and not the fulfillment costs. Hence, if most of the variation in baseline performance of organizations is driven by fulfillment costs, while most of the variation in baseline performance of bureaucrats is driven by participation costs, then we would expect Proposition 2 to be consistent with the heterogeneity of the estimated treatment effects by bureaucrat effectiveness but not by organization effectiveness, which is what we see. Consistent with this, when we compare the features that predict baseline effectiveness for bureaucrats vs organizations, we see differences. The predictors of organization effectiveness are less to do with participation (the types of costs we think may be incorporated in α_c and ψ_c) than those for bureaucrat effectiveness. Those that *do* predict organization effectiveness are more to do with the end user organization itself and, potentially, their idiosyncratic fulfillment costs.

[69] Appendix Figure G.2 shows the analogous results for organizations, confirming the findings in table 6 that there is limited heterogeneity. Appendix Figure G.3 shows that consistent with the findings for prices, we see strong heterogeneity in the impact of the policy on participation by bureaucrat effectiveness. More effective bureaucrats experience large drops in participation, while less effective bureaucrats do not experience these participation drops, and may even see participation increases. Similarly to the effects on prices, there is little evidence of heterogeneity by organization effectiveness.

Overall, these results suggest that, from the perspective of a government trying to minimize the prices it pays for its goods while simultaneously steering government demand towards domestic manufacturers, a "buy local" procurement policy of the form used in Russia is a more effective policy tool when the bureaucrats administering the policy are *less* effective at their job, consistent with the logic of our model in Section 4. We trace out the policy design implications in Sub-section 6.4, after examining what explains this heterogeneity in policy impact in Sub-section 6.3.

6.3 Drivers of performance heterogeneity under different policy regimes

To unpack the relationship between bureaucratic heterogeneity and performance under different policy regimes, we turn again to our data on procurement processes and take an approach similar to the one we used in Sub-section 5.5 to study the drivers of performance in the baseline policy regime. Bureaucratic effectiveness can affect policy performance under different policy regimes in two ways. First, the attributes that are associated with bureaucratic effectiveness in the baseline policy regime may assume a different significance under the preference policy regime even without the bureaucrats or organizations changing the way they carry out their work. Second, new attributes may become important under the preference regime, and so bureaucrats and organizations that are able to change these attributes may benefit the most from the policy change.

We estimate a triple difference regression akin to (7):

$$\begin{aligned} y_{igt} = {} & \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \mathbf{Z}_{igt}\boldsymbol{\theta} + \text{Preferenced}_{gt} \times \mathbf{Z}_{igt}\boldsymbol{\gamma} + \text{PolicyActive}_t \times \mathbf{Z}_{igt}\boldsymbol{\eta} \\ & + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \mathbf{Z}_{igt}\boldsymbol{\pi} + \varepsilon_{igt} \end{aligned} \tag{9}$$

where all terms are as defined in equation (7) but we replace the interactions with bureaucrat ($\hat{\alpha}_b$) and organization ($\hat{\psi}_j$) effectiveness with a vector of observables $\mathbf{Z}_{igt}$. Since our data contain a large number of these (see Table F.1), the vector $\mathbf{Z}_{igt}$ is chosen by the same regularization procedure used in Sub-section 5.5.[70] Comparing the variables in the vector $\mathbf{Z}_{igt}$ selected here to those selected when studying the correlates of baseline performance in Sub-section 5.5 allows us to tell apart the two channels discussed above.

Figure 8 shows the results for prices paid. The variables that affect the policy's impact without changes in bureaucratic behavior are those relating to the requests (5 of 6 variables in $\mathbf{Z}_{igt}$ are also in Figure 4), bureaucrats (3/4 variables), and organizations (4/5 variables). These variables describe the buyers and the way they draft the request documents rather than who participates and how the auction plays out. Without changes in behavior, these upstream factors influence how the policy change affects procurement performance more than downstream factors like the types of participants and how auctions play out. Conversely, the variables that become important under preferences are those relating to the auctions (only 3 of 8 variables in $\mathbf{Z}_{igt}$ are also in Figure 4) and the suppliers (3/8 variables). Particularly noteworthy is that the share of bidders in the auctions

[70] We first run a LASSO procedure with the full set of observables in our data to select the elements of $\mathbf{Z}_{igt}$. For the selected variables, we run regression (9). As in Sub-section 5.5, we also use an elastic net procedure so that the regularization takes greater account of the correlation between the observables. Figures G.6 (for prices) and G.7 (for quality) show that the results are very robust to how much weight we place on the ridge criterion in the elastic net.

who have experience importing or exporting become relevant, presumably since the preference policy drives a wedge between foreign and domestic products.[71]

Summarizing, the same upstream characteristics of the buyers and the way they write requests drive baseline performance and the impacts of policy change. But under the new policy, different characteristics of the auction and the supplier matter. Ultimately, this suggests that there is significant scope for tailoring policy design to the capacity of the implementing bureaucracy since it is these deeper characteristics embodied in the buyers that appear to matter under both policy regimes. These results illuminate *why* the potential scope for and benefits of tailoring policy design to the capacity of implementing bureaucrats are as large as the results in Sub-section 6.2 suggest.

6.4 Implications for policy design

We have seen that deviations from mechanistic, uniform performance—Weber's ideal—depend not just on a bureaucracy's workforce, but also on the policies that these individuals and organizations are asked to carry out. The model in Section 4 illustrates why commonly observed preferences for domestic producers may plausibly achieve public procurement goals in polities with ineffective bureaucracies, but not in polities with effective bureaucracies. We found evidence that this is in fact the case for a 15 percent preference rate in Figure 7A: the policy decreased prices by up to 14 percent when implemented by the least effective Russian bureaucrats, but for more effective bureaucrats, prices *increased*. The adverse impact when the policy is administered by effective bureaucrats in Russia is comparable to that for similar preference policies in the U.S. (Marion, 2007; Krasnokutskaya & Seim, 2011; Athey *et al.*, 2013). This raises the question of whether policy makers may want to pick different bid preference policies depending on the effectiveness of their implementing bureaucracy.

To shed light on this question, we again need both the bureaucrat effectiveness estimates from Section 5 and the heterogeneity-in-impact estimates from Sub-section 6.2. We combine our estimates from Figure 7 of the effect of the 15 percent preference in each decile of the overall effectiveness distribution with the distributions of effectiveness in a range of subgroups of bureaucrats. Assuming that the semi-elasticity of prices with respect to the preference rate is locally constant, we can then estimate the level of the preference rate that would achieve the same effect in each subgroup as we observe on average across the full sample.

Specifically, we assume that for each decile k of effectiveness, the semi-elasticity of prices with respect to the preference rate $1-\gamma$ is equal to the average treatment effect of the 15 percent preference rate we estimate for that decile so that log prices are locally linear in the preference rate, with slope $TE_k/0.15$, where TE_k is the treatment effect for decile k estimated using equation (8) shown in Figure 7A. This is a strong assumption, and the model in Section 4 does not imply this constant elasticity, but we show in Appendix Figure G.8 that such a simplification is nevertheless reasonable locally.

We can then ask, for any subgroup g with distribution of bureaucrats $w_{kg}, k=1,...,10$ across the deciles of effectiveness, what preference rate $1-\gamma_g^*$ would achieve the same impact in that subgroup

[71] Figure G.5 shows that the results are very similar when we study spending quality as the outcome instead of prices.

as the 15 percent rate achieves in the overall sample. For each subgroup, our estimates in Figure 7A imply a treatment effect of a 15 percent bid penalty of $TE_g = \sum_{k=1}^{10} w_{kg} TE_k$, and our constant elasticity assumption implies that we can find the equivalent policy by solving

$$d\log(p_g) = \overline{TE} - TE_g = \left(1 - \gamma_g^* - 0.15\right) \frac{TE_g}{0.15} \Leftrightarrow 1 - \gamma_g^* = \frac{0.15\overline{TE}}{TE_g} \tag{10}$$

where $\overline{TE} = \sum_{k=1}^{10} TE_k$ is the treatment effect in the overall sample. Applying equation (10) in different subgroups allows us to provide a back of the envelope estimate of how policy-makers overseeing different bureaucracies can achieve a given policy goal, in this case a particular overall effect on prices, by tailoring the preference policy to the effectiveness of the implementing bureaucracy.

We consider subgroups of bureaucrats distinguished by the government department they are working with, the level of government they work with, their experience (the volume of transactions they undertake), and whether they work in-house or externally. These are observable markers that we consider in our analysis of the drivers of bureaucratic performance, and which policy-makers might plausibly consider when designing policy. Figure 9 plots these subgroups' equivalent bid penalty $1 - \gamma_g^*$ from equation (10) against the group's average baseline performance, excluding groups for which the 95% confidence interval on $1 - \gamma_g^*$ is wider than 0.3.[72]

Figure 9 shows a wide range of equivalent policies, ranging from 23 percent for the most effective subgroup, to 10 percent for the least effective subgroup. These numbers are, of course, obtained under highly restrictive assumptions, but they nevertheless serve to illustrate the usefulness of considering individual policy-implementers' effectiveness in policy design.

7 Conclusion

In this paper we have presented evidence that, contrary to the mechanistic view of the bureaucracy in much of the existing literature, the individuals and organizations tasked with implementing policy are important sources of variation in states' productivity. Bureaucrats and public sector organizations together account for a full 39 percent of the variation in quality-adjusted prices paid by the Russian government for its inputs. Consistent with a simple endogenous entry model of procurement, effective public procurers engage in practices that lower entry costs for potential suppliers and attract a larger and more diverse pool of participants, allowing them to achieve lower prices. However, in many contexts, the performance of individuals and organizations cannot be directly improved, but the tasks bureaucrats are directed to carry out can. Studying the impact of a "buy local" policy that provides bid preferences for locally manufactured goods, we show that participation increases and prices decrease when the policy is implemented by less effective bureaucrats, while performance is essentially unaffected when the policy is implemented by more effective bureaucrats, consistent with our model.

[72] These tend to be groups with very good baseline performance with many bureaucrats in deciles with estimated treatment effects vey close to zero, leading to noisy estimates when we divide through by them.

These findings have important implications. First, they suggest that there are huge returns to the state from employing more bureaucrats at the high end of the observed performance range, training bureaucrats better, or improving organization-wide characteristics such as management quality—if such changes are possible. Second, our findings imply that the nature of the policy regime in place determines the extent to which differences in bureaucratic effectiveness manifest themselves in differences in public sector output. In turn, this suggests that policies that are suboptimal when state effectiveness is high may become second-best optimal when state effectiveness is low. Achieving the *best* policy outcomes likely requires both improving the effectiveness of the bureaucratic apparatus and choosing policies that are tailored to the effectiveness of their implementers.

References

ABOWD, JOHN M., KRAMARZ, FRANCIS, & MARGOLIS, DAVID. 1999. High Wage Workers and High Wage Firms. *Econometrica*, **67**, 251–333.

ABOWD, JOHN M., CREECY, ROBERT H., & KRAMARZ, FRANCIS. 2002. *Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data*. Census Bureau Technical Paper TP-2002-06.

ALLCOTT, HUNT. 2015. Site Selection Bias in Program Evaluation. *Quarterly Journal of Economics*, **130**, 1117–1165.

ANDREWS, MARTYN J., GILL, LEONARD, SCHANK, THORSTEN, & UPWARD, RICHARD. 2008. High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias. *Journal of the Royal Statistical Society*, **171**(3), 673–697.

ANGRIST, JOSHUA D. 2004. Treatment effect heterogeneity in theory and practice. *The Economic Journal*, C52–C83.

ATHEY, SUSAN, COEY, DOMINIC, & LEVIN, JONATHAN. 2013. Set-Asides and Subsidies in Auctions. *American Economic Journal: Microeconomics*.

BANDIERA, ORIANA, PRAT, ANDREA, & VALLETTI, TOMMASO. 2009. Active and Passive Waste in Government Spending: Evidence from a Policy Experiment. *American Economic Review*, **99**(4), 1278–1308.

BARABASHEV, ALEXEI, & STRAUSSMAN, JEFFREY D. 2007. Public service reform in Russia, 1991–2006. *Public Administration Review*, **67**(3), 373–382.

BERNARD, ANDREW B, JENSEN, J BRADFORD, REDDING, STEPHEN J, & SCHOTT, PETER K. 2007. Firms in international trade. *Journal of Economic perspectives*, **21**(3), 105–130.

BERTRAND, MARIANNE, & SCHOAR, ANTOINETTE. 2003. Managing with style: The effect of managers on firm policies. *Quarterly Journal of Economics*.

BERTRAND, MARIANNE, BURGESS, ROBIN, CHAWLA, ARUNISH, & XU, GUO. 2020. The Glittering Prizes: Career Incentives and Bureaucrat Performance. *Review of Economic Studies*, **87**, 626–655.

BESLEY, TIMOTHY, & PERSSON, TORSTEN. 2009. The origins of state capacity: Property rights, taxation, and politics. *American Economic Review*, **99**, 1218–44.

BEST, MICHAEL CARLOS, BROCKMEYER, ANNE, KLEVEN, HENRIK, SPINNEWIJN, JOHANNES, & WASEEM, MAZHAR. 2015. Production vs Revenue Efficiency With Limited Tax Capacity: Theory and Evidence From Pakistan. *Journal of Political Economy*, **123**, 1311–1355.

BLOOM, NICHOLAS, SONG, JAE, PRICE, DAVID, GUVENEN, FATIH, & VON WACHTER, TILL. 2019. Firming up inequality. *Quarterly Journal of Economics*, **134**(1), 1–50.

BOLD, TESSA, KIMENYI, MWANGI, MWABU, GERMANO, NG'ANG'A, ALICE, & SANDEFUR, JUSTIN. 2018. Experimental Evidence on Scaling Up Education Reforms in Kenya. *Journal of Public Economics*, **168**, 1–20.

BRONNENBERG, BART, DUBÉ, JEAN-PIERRE, GENTZKOW, MATTHEW, & SHAPIRO, JESSE M. 2015. Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium. *Quarterly Journal of Economics*.

CARD, DAVID, HEINING, JÖRG, & KLINE, PATRICK. 2013. Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics*, **128**, 967–1015.

CARD, DAVID, CARDOSO, ANA RUTE, & KLINE, PATRICK. 2016. Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *Quarterly Journal of Economics*.

CARD, DAVID, CARDOSO, ANA RUTE, HEINING, JÖRG, & KLINE, PATRICK. 2018. Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics*, **36**, s13–s70.

CARNEIRO, ANABELA, GUIMARÃES, PAULO, & PORTUGAL, PEDRO. 2012. Real Wages and the Business Cycle: Accounting for Worker, Firm and Job Title Heterogeneity. *American Economic Journal: Macroeconomics*, **4**, 133–152.

CENGIZ, DORUK, DUBE, ARINDRAJIT, LINDNER, ATTILA, & ZIPPERER, BEN. 2019. The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics*, **134**, 1405–1454.

CHETTY, RAJ, FRIEDMAN, JOHN N., & ROCKOFF, JONAH E. 2014. Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates. *American Economic Review*, **104**, 2593–2632.

COVIELLO, DECIO, MORETTI, LUIGI, SPAGNOLO, GIANCARLO, & VALBONESI, PAOLA. 2017. Court Efficiency and Procurement Performance. *Scandinavian Journal of Economics*.

COVIELLO, DECIO, GUGLIELMO, ANDREA, & SPAGNOLO, GIANCARLO. 2018. The Effect of Discretion on Procurement Performance. *Management Science*.

DAL BO, ERNESTO, FINAN, FRED, & ROSSI, MARTIN. 2013. Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *Quarterly Journal of Economics*.

DECAROLIS, FRANCESCO, GIUFFRIDA, LEONARDO M., IOSSA, ELISABETTA, MOLLISI, VINCENZO, & SPAGNOLO, GIANCARLO. 2018. *Bureaucratic Competence and Procurement Outcomes*. Mimeo Universita Bocconi.

DEHEJIA, RAJEEV, POP-ELECHES, CRISTIAN, & SAMII, CYRUS. forthcoming. From Local to Global: External Validity in a Fertility Natural Experiment. *Journal of Business and Economic Statistics*.

DUFLO, ESTHER, GREENSTONE, MICHAEL, PANDE, ROHINI, & RYAN, NICHOLAS. 2013. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *Quarterly Journal of Economics*, **128**, 1499–1545.

DUFLO, ESTHER, GREENSTONE, MICHAEL, PANDE, ROHINI, & RYAN, NICHOLAS. 2018. The Value of Regulatory Discretion: Estimates from Environmental Inspections in India. *Econometrica*, **86**(6), 2123–2160.

ENIKOLOPOV, RUBEN, & ZHURAVSKAYA, EKATERINA. 2007. Decentralization and political institutions. *Journal of Public Economics*.

FERRAZ, CLAUDIO, FINAN, FREDERICO, & SZERMAN, DIMITRI. 2015. *Procuring Firm Growth: The Effects of Government Purchases on Firm Dynamics*. mimeo, UC Berkeley.

FINAN, FREDERICO, OLKEN, BENJAMIN A., & PANDE, ROHINI. 2017. The Personnel Economics of the Developing State. *In:* BANERJEE, ABHIJIT, & DUFLO, ESTHER (eds), *Handbook of Field Experiments, Volume II*. North Holland.

FINKELSTEIN, AMY, GENTZKOW, MATTHEW, & WILLIAMS, HEIDI. 2016. Sources of Geographic Variation in Health Care: Evidence from Patient Migration. *Quarterly Journal of Economics*.

GAURE, SIMEN. 2013. OLS with multiple high dimensional category variables. *Computational Statistics and Data Analysis*, **66**, 8–18.

GENTRY, MATTHEW, & LI, TONG. 2014. Identification in auctions with selective entry. *Econometrica*, **82**(1), 315–344.

GOLDBERG, PINELOPI KOUJIANOU, & DE LOECKER, JAN. 2014. Firm Performance in a Global Market. *Annual Review of Economics*.

HANSMAN, CHRISTOPHER, HJORT, JONAS, & LEÓN, GIANMARCO. 2019. Interlinked Firms and the Consequences of Piecemeal Regulation. *Journal of the European Economic Association*.

HECKMAN, JAMES, & SMITH, JEFFREY. 1997. Making the Most out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts. *Review of Economic Studies*, 487–535.

HOBERG, GERARD, & PHILLIPS, GORDON. 2016. Text-Based Network Industries and Endogenous Product Differentiation. *Journal of Political Economy*, **124**(5), 1423–1465.

JONES, BENJAMIN F., & OLKEN, BENJAMIN A. 2005. Do leaders matter? National leadership and growth since World War II. *Quarterly Journal of Economics*.

KANE, THOMAS J., & STAIGER, DOUGLAS O. 2008. *Estimating teacher impacts on student achievement: An experimental evaluation*. NBER working Paper.

KHAN, ADNAN, KHWAJA, ASIM, & OLKEN, BEN. 2016. Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors. *Quarterly Journal of Economics*.

KHAN, ADNAN, KHWAJA, ASIM, & OLKEN, BEN. 2018. Making Moves Matter: Experimental Evidence on Incentivizing Bureaucrats through Performance-Based Postings. *American Economic Review*.

KHANDELWAL, AMIT. 2010. The Long and Short (of) Quality Ladders. *The Review of Economic Studies*, **77**(4), 1450–1476.

KLING, JEFFREY R, LIEBMAN, JEFFREY B, & KATZ, LAWRENCE F. 2007. Experimental Analysis of Neighborhood Effects. *Econometrica*, **75**, 83–119.

KRASNOKUTSKAYA, ELENA, & SEIM, KATJA. 2011. Bid Preference Programs and Participation in Highway Procurement Auctions. *American Economic Review*, 2653–2686.

LACETERA, NICOLA, LARSEN, BRADLEY, POPE, DEVING G., & SYDNOR, JUSTIN. 2016. Bid Takers or Market Makers? The Effect of Auctioneers on Auction Outcomes. *American Economic Journal: Microeconomics*, **8**, 195–229.

LAFFONT, JEAN-JAQUES. 2005. *Regulation and Development*. Cambridge University Press.

LANCASTER, TONY. 2000. The Incidental Parameter Problem since 1948. *Journal of Econometrics*, **95**, 391–413.

LEWIS-FAUPEL, SEAN, NEGGERS, YUSUF, OLKEN, BENJAMIN, & PANDE, ROHINI. 2016. *Can Electronic Procurement Improve Infrastructure Provision? Evidence from Public Works in India and Indonesia*. American Economic Journal: Economic Policy.

LIEBMAN, JEFFREY B., & MAHONEY, NEALE. 2017. Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement. *American Economic Review*, **107**, 3510–3549.

MARION, JUSTIN. 2007. Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions. *Journal of Public Economics*, 1591–1624.

MCAFEE, PRESTON R., & MCMILLAN, JOHN. 1989. Government procurement and international trade. *Journal of International Economics*, **26**(3), 291–308.

MILGROM, PAUL. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.

PAGE, LUCY, & PANDE, ROHINI. 2018. Ending Global Poverty: Why Money Isn't Enough. *Journal of Economic Perspectives*, **32**, 173–200.

RASUL, IMRAN, & ROGGER, DANIEL. 2018. Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service. *Economic Journal*.

RAUCH, JAMES E. 1999. Networks Versus Markets in International Trade. *Journal of International Economics*, 7–35.

SAMUELSON, WILLIAM F. 1985. Competitive bidding with entry costs. *Economics Letters*, **17**(1), 53–57.

SCHAPPER, PAUL R., MALTA, JOAO NUNO VEIGA, & GILBERT, DIANE L. 2009. Analytical Framework for the Management and Reform of Public Procurement. *In:* THAI, KHI V. (ed), *International Handbook of Public Procurement*. CRC Press.

SCHULZE, GÜNTHER G, SJAHRIR, BAMBANG SUHARNOKO, & ZAKHAROV, NIKITA. 2016. Corruption in Russia. *The Journal of Law and Economics*, **59**(1), 135–171.

SILVER, DAVID. 2016. *Haste or Waste? Peer Pressure and the Distribution of Marginal Returns to Health Care*. Mimeo: UC Berkeley.

SUTTON, JOHN. 1998. *Technology and Market Structure: Theory and History*. Cambridge: MIT Press.

SYVERSON, CHAD. 2004. Market Structure and Productivity: A Concrete Example. *Journal of Political Economy*.

SZAKONYI, DAVID. 2018. Businesspeople in Elected Office: Identifying Private Benefits from Firm-Level Returns. *American Political Science Review*, **112**(2), 322–338.

TRANSPARENCY INTERNATIONAL. 2016. *Corruption Perceptions Index*. 2016 version.

WEBER, MAX. 1921. *Economy and Society*.

XU, GUO. 2018. The Costs of Patronage: Evidence from the British Empire. *American Economic Review*, **108**(11), 3170–3198.

YAKOVLEV, ANDREI, DEMIDOVA, OLGA, & BALAEVA, OLGA. 2010. The System of Public Procurements in Russia: On the Road of Reform. *HSE Policy Paper*, 1–23.

YAKOVLEV, ANDREI, YAKOBSON, LEV, & YUDKEVICH, MARIA. 2011. *The Public Procurement System in Russia: Road Toward A New Quality*. Unpublished Working Paper.

YAKOVLEV, EVGENY, & ZHURAVSKAYA, EKATERINA. 2014. The Unequal Enforcement of Liberalization: Evidence from Russia's Reform of Business Regulation. *Journal of the European Economic Association*.

FIGURE 1: EVENT STUDY OF PROCUREMENT PRICES AROUND TIMES ORGANIZATIONS SWITCH BUREAUCRATS

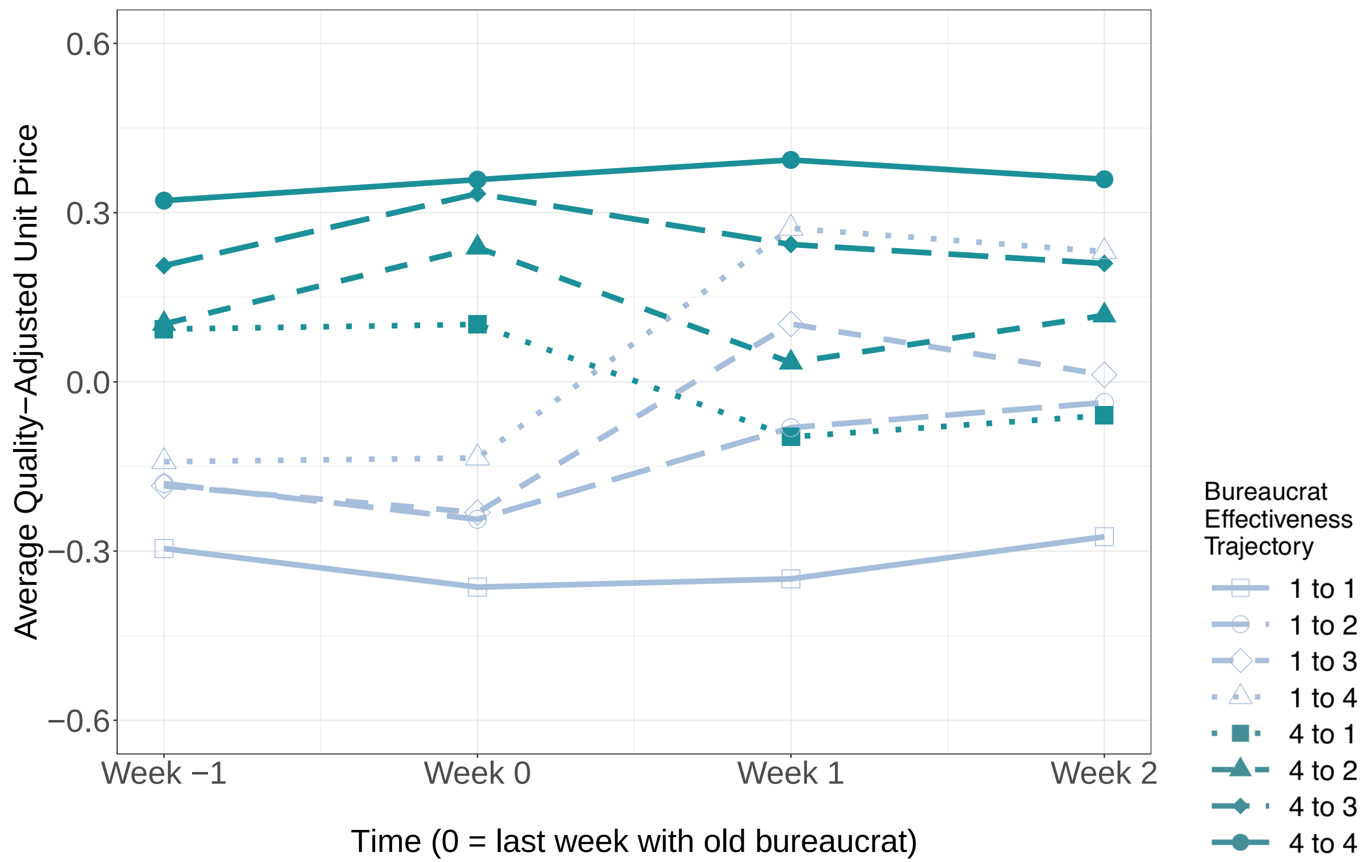

Notes: The figure shows time trends in prices around the time that organizations switch which bureaucrat makes purchases on their behalf. The horizontal axis indexes weeks in which bureaucrat-organization pairs work together, with time 0 being the last week in which the organization works with the old bureaucrat just before switch, and time 1 being the first week the organization works with the new bureaucrat after the switch. The y axis measures average residualized prices paid by the bureaucrat-organization pair where prices are residualized by regressing log unit prices on good and month fixed effects. We create a balanced panel in which we require each bureaucrat-organization pair to work together in at lesat two separate weeks and each bureaucrat to work with at least one other organization in the quarter containing time 0 (for the "old" bureaucrat the organization works with before the switch) or time 1 (for the "new" bureaucrat the organization works with after the switch). Bureaucrats are classified into quartiles according to the average (residualized) prices they achieve with the *other* organizations they work with in the quarter containing time 0 (for the old bureaucrat) or the quarter containing time 1 (for the new bureaucrat).

FIGURE 2: ROBUSTNESS TO USING SUBSAMPLES OF INCREASINGLY HETEROGENEOUS GOODS

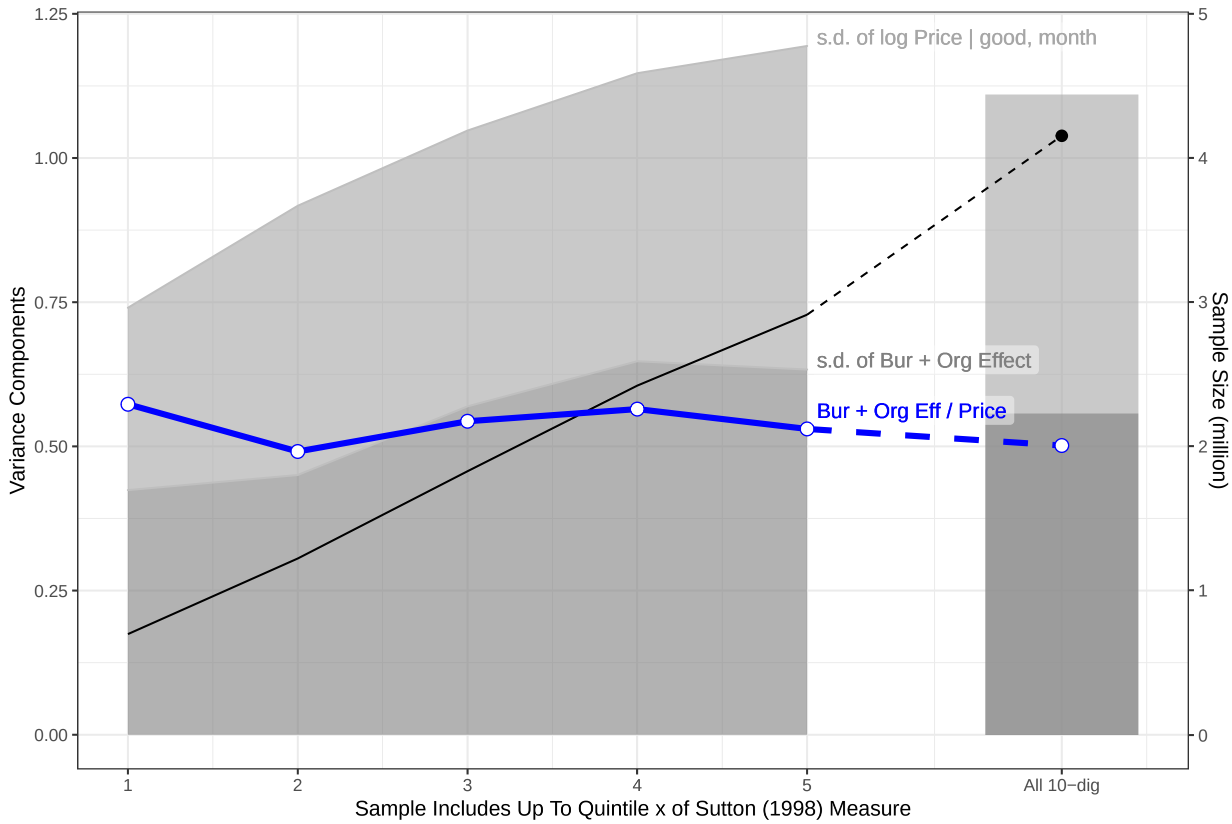

Notes: The figure shows the components of the variance of prices due to bureaucrats and organizations estimated by implementing the variance decomposition in equation (4) (see notes to Table 2 for details). The right-most bar uses the sub-sample consisting of all auctions for goods that our text analysis classification method is able to assign a 10-digit product code to. The left portion of the figure uses the sub-set of the sample that we can match to the scope-for-quality-differentiation ladder developed by Sutton (1998). Moving from right to left we remove quintiles of the data with the highest scope for quality differentiation, as shown by the black line, which indicates the sample size used. The dark shaded region is the variance of prices attributable to the bureaucrats and organizations. The dark and light shaded regions show the total variance of prices. The blue line shows the fraction of the overall variance attributable to bureaucrats and organization, highlighting that it remains roughly constant as we add more heterogeneous goods to the sample.

Figure 3: Buyers Who Achieve Low Prices Also Achieve Better Spending Quality

Panel A: Bureaucrats

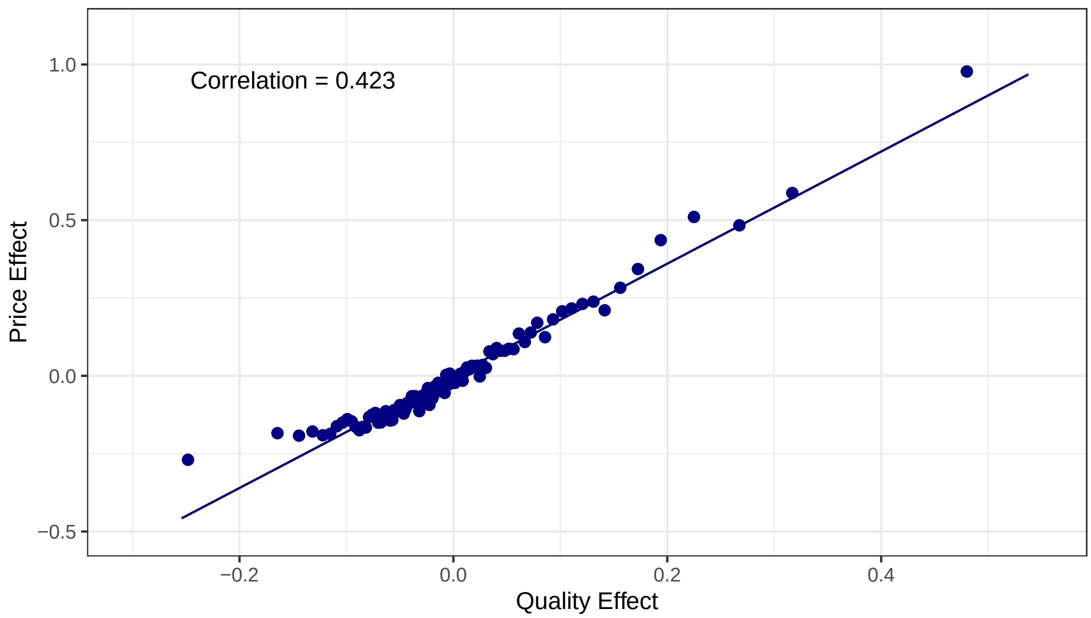

Panel B: Organizations

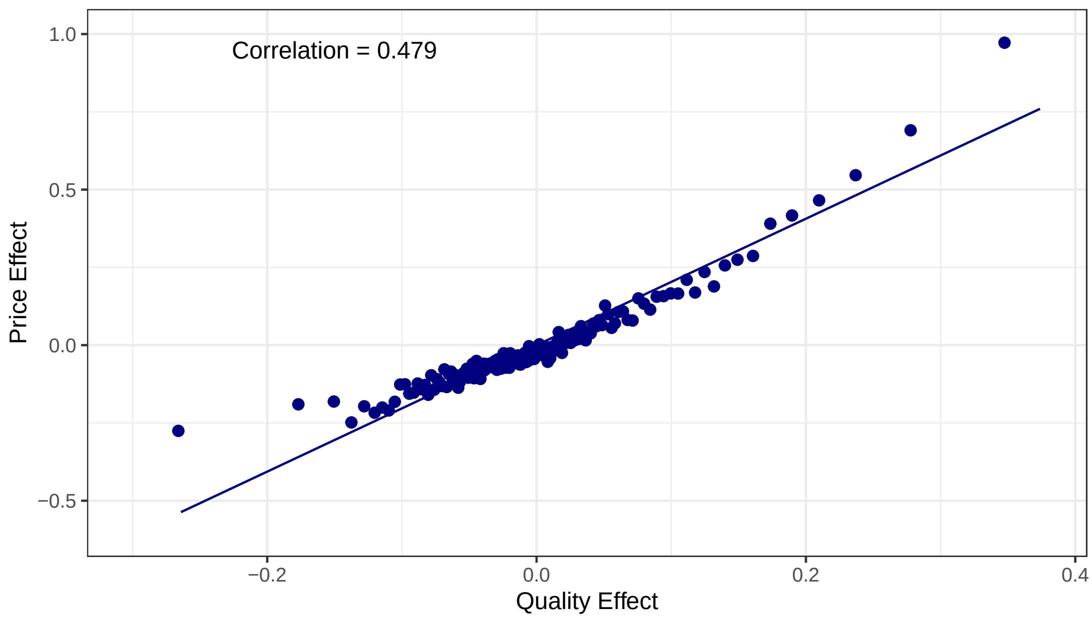

Notes: The figure shows the correlation between bureaucrats' (panel A) and organizations' (panel B) covariance-shrunk price effects and their covariance-shrunk spending quality effects. They are estimated by implementing the variance decomposition in equation (4) and then implementing our covariance-shrinkage method (see notes to Table 2 for details). The panels show binned scatterplots together with a regression line fitted on the underlying data, and the correlation between the two effects shown in the upper left corner.

FIGURE 4: CORRELATES OF BUREAUCRAT EFFECTIVENESS (PRICE)

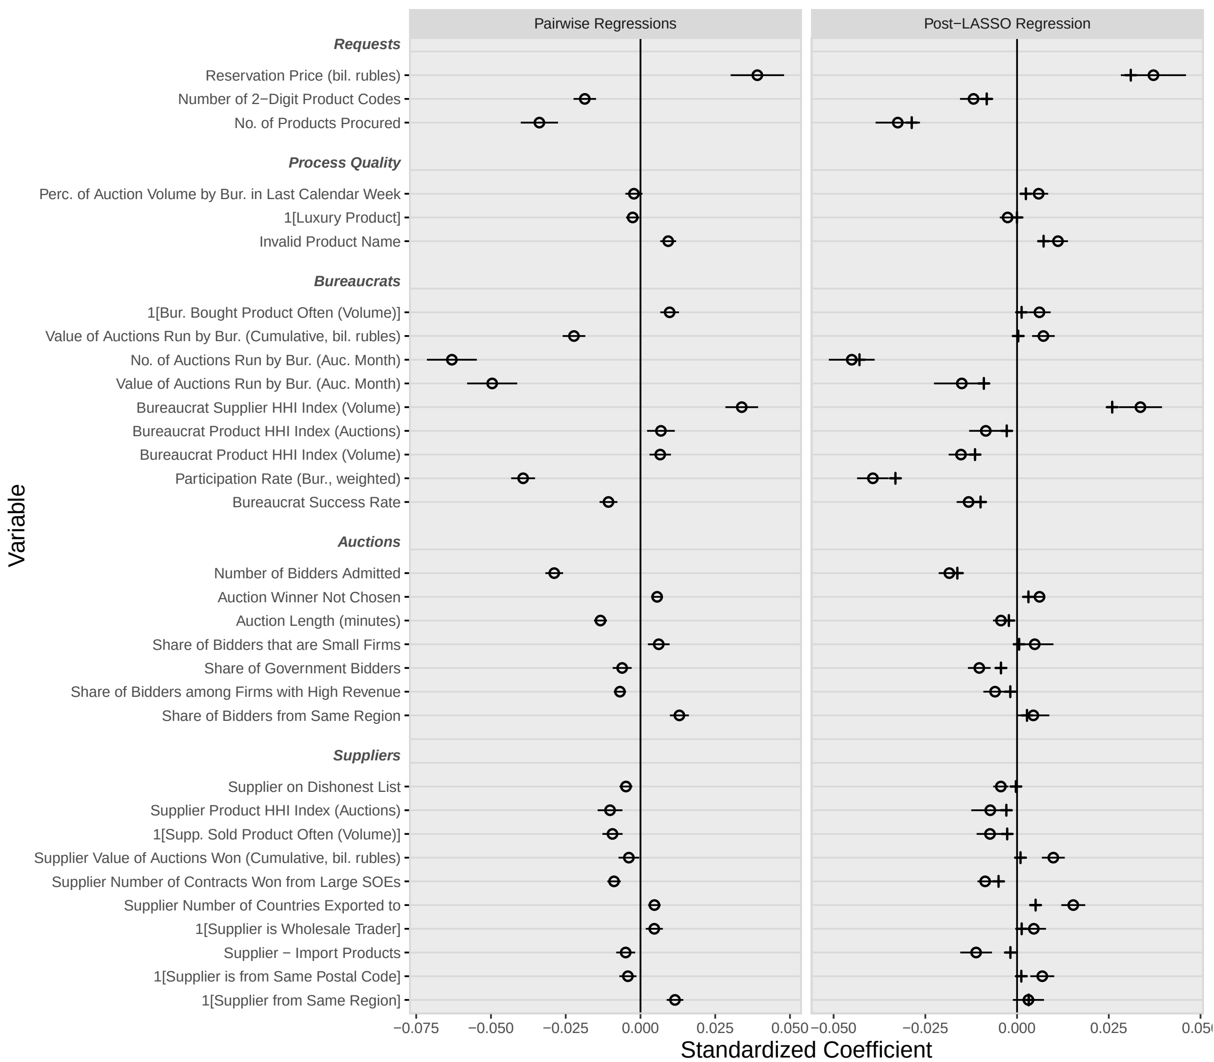

Notes: The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (3) for prices: $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. As described in section 5.5, since our data contain a large number of observables (see Table F.1), we use a LASSO procedure to select 30 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients (in circles) and confidence intervals from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients (in circles) and confidence from the multivariate regression of the effects on all of the selected variables as well as the LASSO coefficients (as crosses). To facilitate comparison of effect sizes across variables, all variables are standardized to have unit standard deviation.

FIGURE 5: CORRELATES OF BUREAUCRAT EFFECTIVENESS (QUALITY)

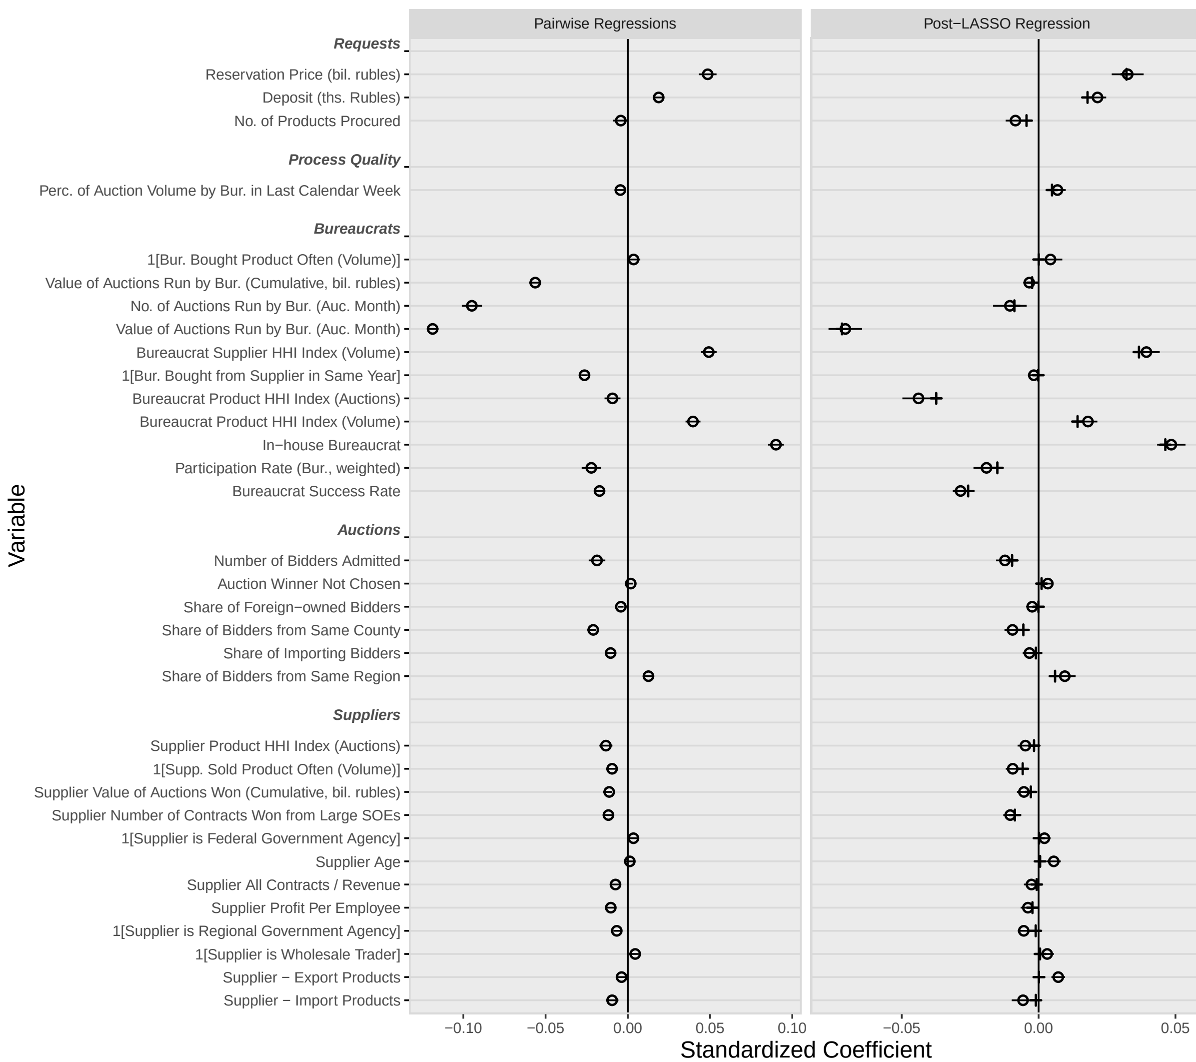

Notes: The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (3) for spending quality as discussed in Section 5.4: $q_i = \mathbf{X}_i \boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. As described in section 5.5, we use a LASSO procedure to select 30 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients (in circles) and confidence intervals from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients (in circles) and confidence from the multivariate regression of the effects on all of the selected variables as well as the LASSO coefficients (as crosses). To facilitate comparison of effect sizes across variables, all variables are standardized to have unit standard deviation.

FIGURE 6: EVENT STUDY OF EFFECT OF BID PREFERENCES ON AVERAGE PRICES

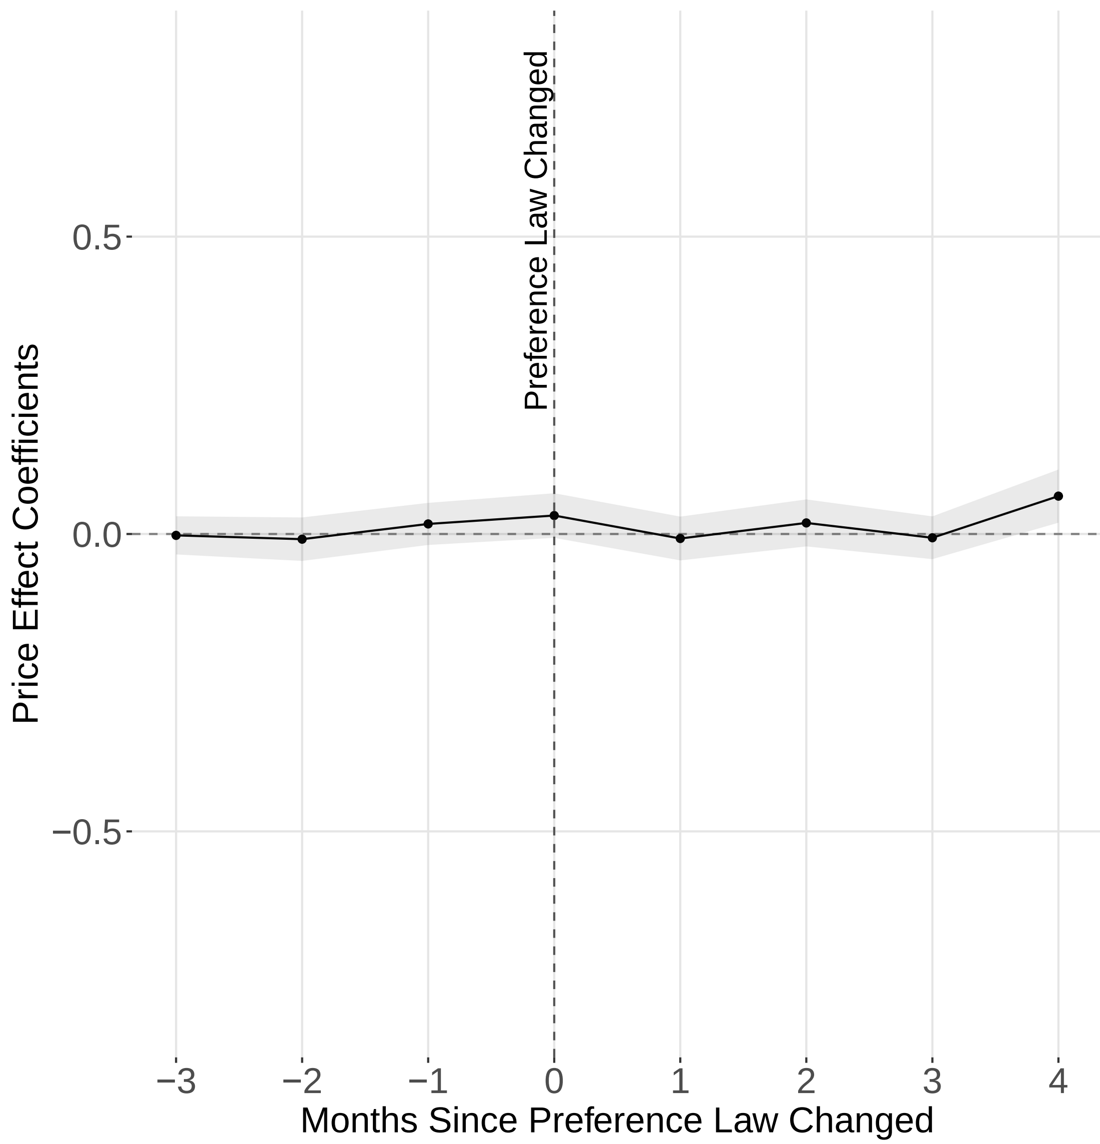

Notes: The figure shows the results of an event study analysis of the impact of the preferences policy on prices. Following Cengiz *et al.* (2019), we stack all the events (the preference list being published) and focus on a window starting three months before and ending four months after each year's preference list is published. We estimate equation (6): $p_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \sum_{s=-3}^{4} \delta_s \text{Preferenced}_{gt} \times \mathbf{1}\{t - \text{ListMonth}_t = s\} + \varepsilon_{igt}$ where p_{igt} is the log price paid in transaction i for good g in month t; $\mathbf{X}_{igt}$ are the same controls we use in Section 5, but for clarity we separate out the good and month fixed effects, μ_g and λ_t; Preferenced_{gt} is a dummy indicating that g is on the preferences list in the year month t falls within, ListMonth_t is the month closest to month t in which a preference list is published; and ε_{igt} is an error term we allow to be clustered by month and good. The figure shows the estimated δ_s coefficients and their 95% confidence intervals.

Figure 7: Heterogeneity of Bid Preferences' Effect by Bureaucrat Effectiveness

Panel A: Difference in Differences by Bureaucrat Effectiveness Decile

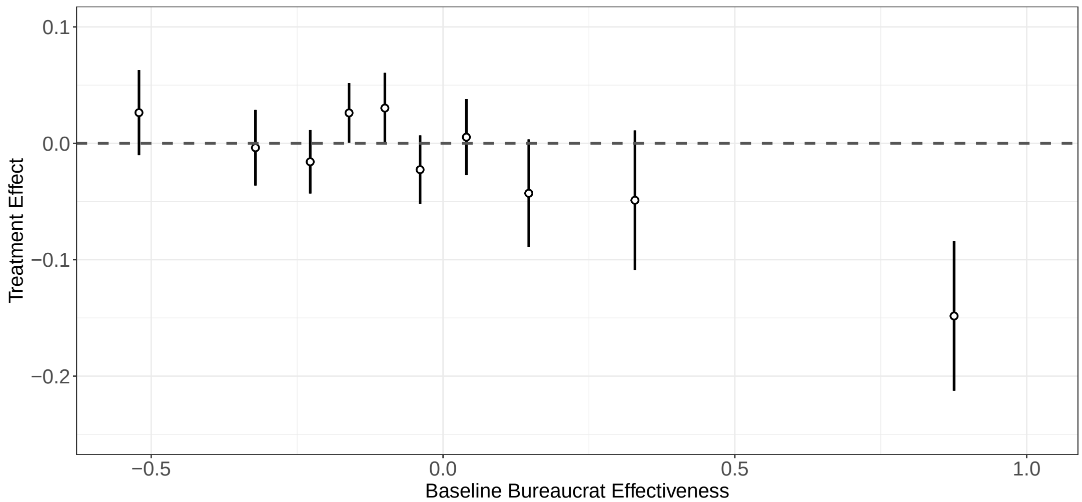

Panel B: Event Study by Bureaucrat Effectiveness

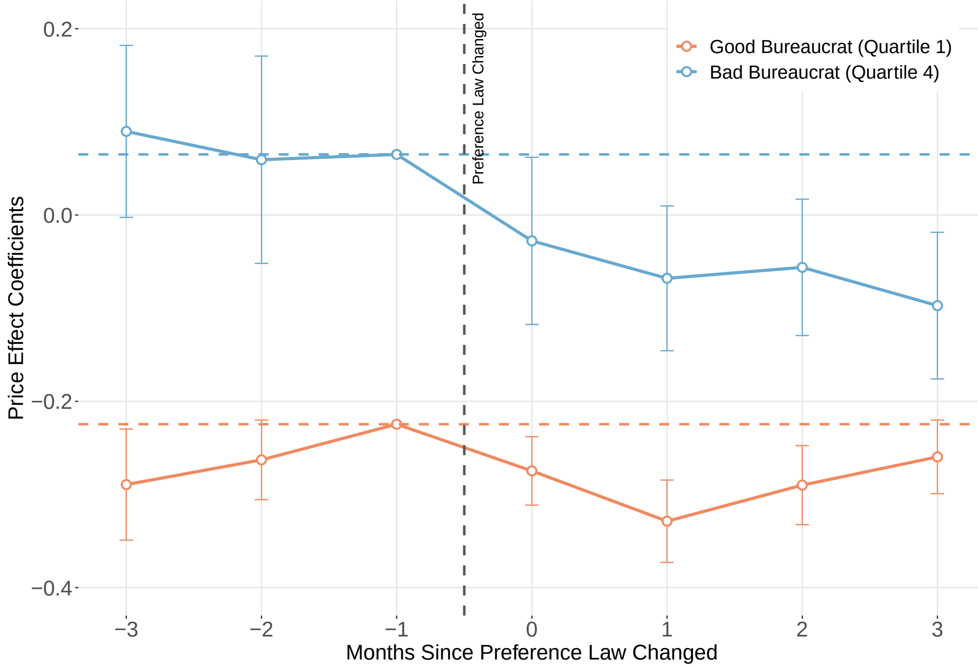

Notes: The figure shows how the impacts of the introduction of bid preferences varies by the effectiveness of the implementing bureaucrat. Panel A shows estimates from implementing the triple difference model (8) to estimate separate effects for each decile of bureaucrat effectiveness: $y_{igt} = \sum_{k=1}^{10} \{D_{kj} + D_{kb} \times (\rho_k \text{Preferenced}_{gt} + \eta_k \text{PolicyActive}_t + \pi_k \text{Preferenced}_{gt} \times \text{PolicyActive}_t)\} + \mathbf{X}_{igt}\beta + \mu_g + \mu_t + \varepsilon_{igt}$ where D_{kj} and D_{kb} are indicators for organization j and bureaucrat b belonging to decile k of their respective distributions of effectiveness. The horizontal axis plots the average effectiveness within the relevant decile, while the vertical axis plots the estimated treatment effects π_k with their 95% confidence intervals. Panel B extends the event study (6) shown in figure 6 (see notes to figure 6 for details) to estimate separate effects for the top and bottom quartile of bureaucrats. Rather than normalizing the reference month (the month before the preference list is published) to zero, we normalized it to the baseline performance in each group to better highlight how different their performance was before the preferences were introduced, and how their performance converges as a result of the preferences.

Figure 8: Predictors of Heterogeneity of Effect of Bid Preferences for Domestic Producers

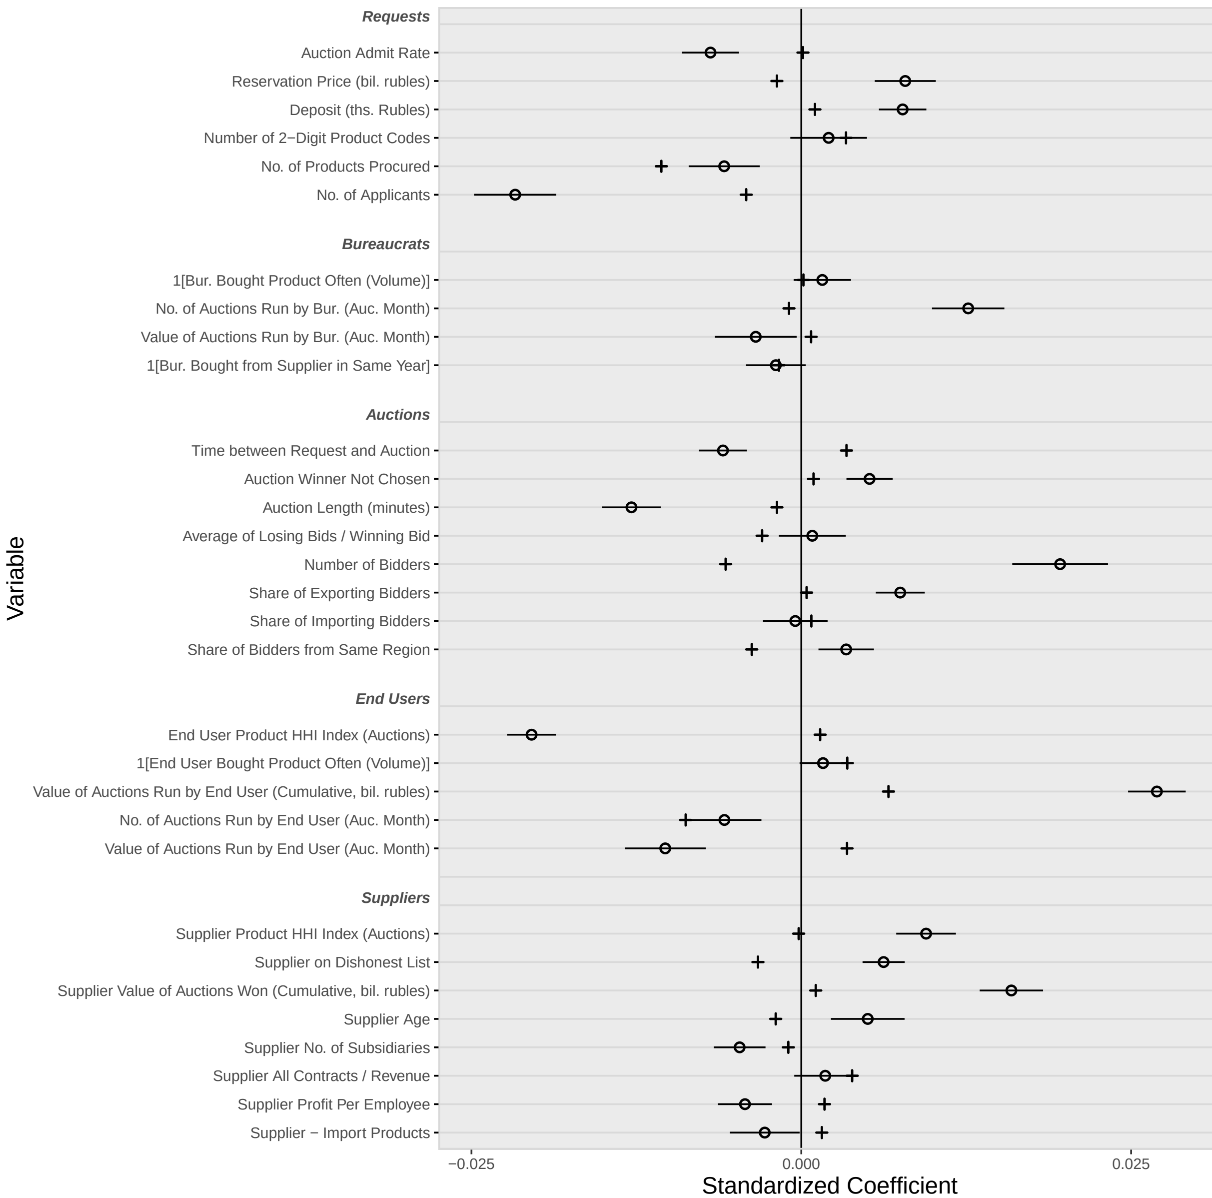

Notes: The figure shows the results of estimating our triple-differences specification for heterogeneity of the effect of bid preferences (9): $y_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \mathbf{Z}_{igt}\boldsymbol{\theta} + \text{Preferenced}_{gt} \times \mathbf{Z}_{igt}\boldsymbol{\gamma} + \text{PolicyActive}_t \times \mathbf{Z}_{igt}\boldsymbol{\eta} + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \mathbf{Z}_{igt}\boldsymbol{\pi} + \varepsilon_{igt}$. Since our data contain a large number of these (see Table F.1), the vector $\mathbf{Z}_{igt}$ is chosen by the same regularization procedure used in figure 4 and described in Sub-section 5.5 to return 30 non-zero coefficients. The coefficients from the LASSO are shown as crosses, while the circles show the coefficients and 95% confidence intervals of a multivariate regression including the 30 observables.

FIGURE 9: TAILORING BID PREFERENCES TO BUREAUCRATIC CAPACITY

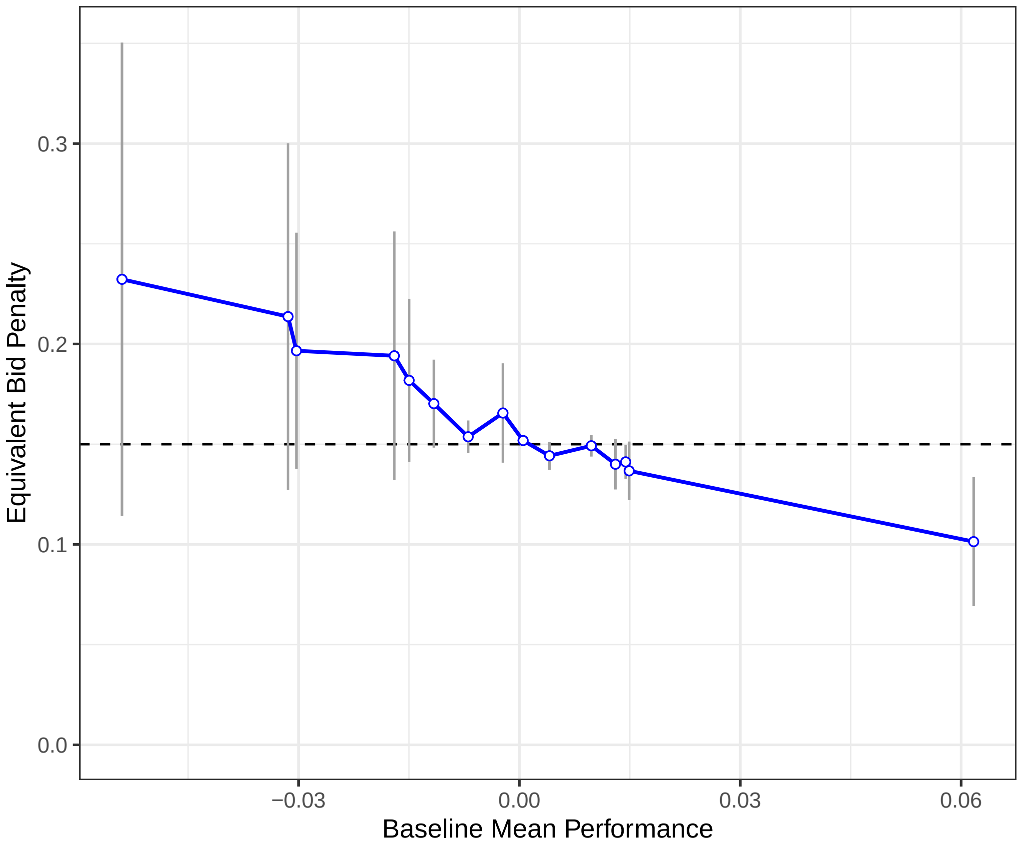

Notes: The figure shows estimates of the preference policy that attains the same impact in subsamples of bureaucrats as the 15 percent policy achieves in the overall sample. We combine our estimates in figure 7 of the effect of the 15 percent preference in each decile of the overall effectiveness distribution with the distributions of bureaucratic effectiveness in a range of subgroups of our data. The subgroups considered are, from left to right: sport department; culture department; regional government; housing department; non in-house bureaucrats; bureaucrats with high auction volume; buyers far from their regional capital; buyers near their regional capital; bureaucrats with low auction volume; in-house bureaucrats; municipal government; education department; internal affairs department; other departments; and federal government. We assume that for each decile k of effectiveness, log prices are locally linear in the preference rate, with slope $TE_k/0.15$, where TE_k is the treatment effect for decile k estimated using equation (8) shown in Figure 7A. For any subgroup g with a distribution of bureaucrats $w_{kg}, j=1,...,10$ across the deciles of effectiveness, we can find the preference rate $1-\gamma_g^*$ that would achieve the same impact in that subgroup as the 15 percent rate achieves in the overall sample as follows. For each subgroup, our estimates in figure 7A imply a treatment effect of a 15 percent bid penalty of $TE_g = \sum_{k=1}^{10} w_{jg} TE_k$, and our constant elasticity assumption implies that the equivalent policy solves (10) yielding $1-\gamma_g^* = 0.15\overline{TE}/TE_g$ where $\overline{TE} = \sum_{k=1}^{10} TE_k$ is the treatment effect in the overall sample. These are shown on the vertical axis of the figure along with their 95% confidence intervals

TABLE 1: SUMMARY STATISTICS

	All Products			Pharmaceuticals Subsample		
	(1) No Preferences Full Sample	(2) No Preferences Analysis Sample	(3) Analysis Sample With Preferences	(4) No Preferences Full Sample	(5) No Preferences Analysis Sample	(6) Analysis Sample With Preferences
(1) # of Bureaucrats	115,854	37,722	37,722	5,561	2,473	2,473
(2) # of Organizations	88,306	44,560	44,560	3,662	1,866	1,866
(3) # of Connected Sets	26,234	616	616	0	129	129
(4) # of Bureaucrats with >1 Org.	14,090	11,063	11,320	965	926	1,095
(5) # of Organizations with >1 Bur.	54,575	37,306	37,536	2,076	1,449	1,596
(6) Mean # of Bureaucrats per Org.	3.96	5.59	6.02	3.1	4.32	6.3
(7) Mean # of Organizations per Bur.	3.02	6.6	7.12	2.04	3.26	4.75
(8) # of Federal Organizations	12,889	1,583	1,583	496	26	26
(9) # of Regional Organizations	25,162	15,530	15,530	2,786	1,599	1,599
(10) # of Municipal Organizations	50,255	27,447	27,447	380	241	241
(11) # of Health Organizations	10,167	7,231	7,231	3,172	1,705	1,705
(12) # of Education Organizations	42,045	25,271	25,271	109	61	61
(13) # of Internal Affairs Organizations	3,126	668	668	105	3	3
(14) # of Agr/Environ Organizations	1,032	255	255	26	1	1
(15) # of Other Organizations	31,936	11,135	11,135	250	96	96
(16) # of Goods	16,373	14,875	15,649	4,220	3,861	4,351
(17) Mean # of Goods Per Bur.	35	72.5	93.2	31.6	42.5	82.3
(18) # of Regions	86	86	86	85	79	79
(19) Mean # of Regions per Bur.	1	1	1	1	1	1
(20) # of Auction Requests	1,733,433	1,199,363	1,871,717	62,755	42,874	114,807
(21) Mean # of Requests per Bur.	15	31.8	49.6	11.3	17.3	46.4
(22) Mean # of Applicants	3.6	3.6	3.46	2.98	3.03	3.02
(23) Mean # of Bidders	2.06	2.07	2.07	1.94	1.98	2
(24) Mean Reservation Price	0.149	25,140	0.134	0.096	0.062	0.055
(25) Quantity Mean	1,131	1,053	1,124	1,201	1,719	975
Median	20	25	27	40	45	50
SD	80,563	90,917	174,951	136,260	172,145	108,598
(26) Total Price Mean (bil. USD)	93.3	80.1	81.2	128	91.1	101
Median	4.67	4.32	4.74	6.23	6.7	7.06
SD	578	493	482	5,745	493	525
(27) Unit Price Mean (bil. USD)	72.1	61.3	55.6	20.2	25.4	28.8
Median	0.21	0.167	0.18	0.175	0.18	0.18
SD	21,248	23,015	19,168	226	265	281
(28) Mean # of Contract Renegotiations (log)	0.126	0.121	0.133	0.15	0.141	0.178
(29) Mean Size of Cost Over-run	-0.002	-0.002	-0.002	-0.003	-0.003	-0.003
(30) Mean Length of Delay in Days (log)	0.061	0.064	0.057	0.077	0.076	0.069
(31) Mean 1[End User Complained about Contract]	0.001	0.001	0.001	0.001	0	0.001
(32) Mean 1[Contract Cancelled]	0.012	0.012	0.009	0.016	0.016	0.01
(33) Mean 1[Product is of Substandard Quality]	0.005	0.005	0.009	0.075	0.058	0.041
(34) # of Observations	15,096,254	11,339,188	16,348,332	290,483	181,961	460,531
(35) Total Procurement Volume (bil. USD)	516	395	629	14.5	9.38	19.9

Notes: The table reports summary statistics for six samples. The All Products columns show statistics for purchases of all off-the-shelf goods, while the Pharmaceuticals Subsample columns restrict attention to purchases of medicines. Full Sample denotes all unpreferenced auctions. Analysis Sample denotes all unpreferenced auctions in connected sets that fulfill the restrictions discussed in section 5.2: singleton bureaucrat-organization, bureaucrat-good, organization-good pairs, and levels of our control fixed effects are removed; each procurer (bureaucrats and organizations) implements a minimum of five purchases; and connected sets have at least three bureaucrats and organizations. With Bid Preferences denotes all preferenced auctions that fulfill the same three restrictions. Organizations working in Education include schools, universities, pre-schools, and youth organizations. Organizations working in Internal Affairs include police, emergency services, local administration, taxes, and transportation. Organizations working in Agriculture or the Environment include environmental protection funds, agricultural departments and nature promotion agencies. The Other category includes funds, monitoring agencies, and land cadasters, among many others. All sums are measured in billions of US dollars at an exchange rate of 43 rubles to 1 US dollar.

TABLE 2: SHARE OF VARIATION IN POLICY PERFORMANCE EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS

	Fixed Effects (1)	(s.e.) (2)	Split Sample (3)	(s.e.) (4)	Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	1.199	(0.0343)	1.253	(0.0322)	0.824	0.425
(2) s.d. of Organization Effects (across orgs)	1.133	(0.0416)	1.179	(0.0476)	0.786	0.377
(3) s.d. of Bureaucrat Effects (across items)	0.795	(0.0321)	0.830	(0.0414)	0.601	0.261
(4) s.d. of Organization Effects (across items)	0.931	(0.0469)	0.970	(0.0576)	0.709	0.355
(5) Bur-Org Effect Correlation (across items)	-0.726	(0.0156)	-0.557	(0.0395)	-0.669	0.297
(6) s.d. of Bur + Org Effects Within CS (across items)	0.651	(0.0209)	0.655	(0.0221)	0.542	0.499
(7) s.d. of log unit price	2.188		2.188		2.188	2.188
(8) s.d. of log unit price | good, month	1.280		1.280		1.280	1.280
(9) Adjusted R-squared	0.963		0.963		0.963	0.963
(10) Number of Bureaucrats	37,722		37,722		37,722	37,722
(11) Number of Organizations	44,560		44,560		44,560	44,560
(12) Number of Bureaucrat-Organization Pairs	248,898		248,898		248,898	248,898
(13) Number of Connected Sets	616		616		616	616
(14) Number of Observations	11,339,188		11,339,188		11,339,188	11,339,188

Notes: The table shows the components of the variance due to bureaucrats, organizations, and connected sets estimated by implementing the variance decomposition in equation (4). The sample used is the All Products-Analysis Sample summarized in Table 1. Rows 1 & 2 show the s.d. of the bureaucrat, organization and connected set effects. Rows 3–6 show the components of the variance of prices across purchases, effectively weighting the estimates in rows 1–3 by the number of purchases they conduct. Column 1 uses the fixed effect estimates from equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates. Columns 2 and 4 show standard errors of the estimates in columns 1 and 3, respectively, estimated by bootstrapping 100 times. Column 5 uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2, and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2/(\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column 1, and analogously for the organization effects. Column 6 shows our preferred estimates, which form predictions of the bureaucrat and organization effects that minimize the expected sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates. Formally, the covariance shrinkage predictors solve $\min_{\boldsymbol{\Lambda}} \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}})'(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}})\right]$ where $\hat{\boldsymbol{\theta}}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 5.2.

Table 3: Robustness to Restricting to Pharmaceuticals Subsample with Barcode Information

	Fixed Effects (1)	(s.e.) (2)	Split Sample (3)	(s.e.) (4)	Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	0.266	(0.0116)	0.244	(0.0152)	0.124	0.0803
(2) s.d. of Organization Effects (across orgs)	0.207	(0.00752)	0.210	(0.00792)	0.0883	0.0575
(3) s.d. of Bureaucrat Effects (across items)	0.184	(0.0293)	0.191	(0.0317)	0.111	0.0665
(4) s.d. of Organization Effects (across items)	0.197	(0.0359)	0.206	(0.0376)	0.106	0.0544
(5) Bur-Org Effect Correlation (across items)	-0.544	(0.0854)	-0.192	(0.0365)	-0.276	-0.0304
(6) s.d. of Bur + Org Effects Within CS (across items)	0.183	(0.00625)	0.183	(0.00688)	0.130	0.0846
(7) s.d. of log unit price	1.914		1.914		1.914	1.914
(8) s.d. of log unit price | good, month	0.430		0.430		0.430	0.430
(9) Adjusted R-squared	0.996		0.996		0.996	0.996
(10) Number of Bureaucrats	2,473		2,473		2,473	2,473
(11) Number of Organizations	1,866		1,866		1,866	1,866
(12) Number of Bureaucrat-Organization Pairs	8,067		8,067		8,067	8,067
(13) Number of Connected Sets	129		129		129	129
(14) Number of Observations	181,961		181,961		181,961	181,961

Notes: The table shows the components of the variance due to bureaucrats, organizations, and connected sets estimated by implementing the variance decomposition in equation (4). The sample used is the Pharmaceuticals-Analysis Sample summarized in Table 1. The table is constructed analogously to table 2 (whose notes contain further details). All methods are described fully in Section 5.2.

TABLE 4: SPENDING QUALITY VARIANCE DECOMPOSITION

	Fixed Effects (1)	(s.e.) (2)	Split Sample (3)	(s.e.) (4)	Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	0.378	(0.0239)	0.421	(0.0265)	0.187	0.0999
(2) s.d. of Organization Effects (across orgs)	0.379	(0.0407)	0.418	(0.0458)	0.194	0.0888
(3) s.d. of Bureaucrat Effects (across items)	0.338	(0.0412)	0.362	(0.0441)	0.187	0.0806
(4) s.d. of Organization Effects (across items)	0.373	(0.0593)	0.396	(0.0638)	0.211	0.089
(5) Bur-Org Effect Correlation (across items)	-0.809	(0.0287)	-0.596	(0.0854)	-0.701	0.310
(6) s.d. of Bur + Org Effects Within CS (across items)	0.222	(0.0229)	0.227	(0.0229)	0.155	0.137
(7) s.d. of quality index	0.592		0.592		0.592	0.592
(8) s.d. of quality index | good, month	0.570		0.570		0.570	0.570
(9) Adjusted R-squared	0.946		0.946		0.946	0.946
(10) Number of Bureaucrats	37,722		37,722		37,722	37,722
(11) Number of Organizations	44,560		44,560		44,560	44,560
(12) Number of Bureaucrat-Organization Pairs	248,898		248,898		248,898	248,898
(13) Number of Connected Sets	616		616		616	616
(14) Number of Observations	11,339,188		11,339,188		11,339,188	11,339,188

Notes: The table shows the components of the variance due to bureaucrats, organizations, and connected sets estimated by implementing the variance decomposition in equation (4) but with spending quality as the outcome, as discussed in section 5.4. The sample used is the All Products-Analysis Sample summarized in Table 1. Rows 1 & 2 show the s.d. of the bureaucrat, organization and connected set effects. Rows 3–6 show the components of the variance of prices across purchases, effectively weighting the estimates in rows 1 & 2 by the number of purchases they conduct. Column 1 uses the fixed effect estimates from equation (3): $q_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates. Columns 2 and 4 show standard errors of the estimates in columns 1 and 3, respectively, estimated by bootstrapping 100 times. Column 5 uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2, and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2/(\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column 1, and analogously for the organization effects. Column 6 shows our preferred estimates, which form predictions of the bureaucrat and organization effects that minimize the expected sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates. Formally, the covariance shrinkage predictors solve $\min_{\mathbf{\Lambda}} \mathbb{E}\left[(\boldsymbol{\theta} - \mathbf{\Lambda}\hat{\boldsymbol{\theta}})'(\boldsymbol{\theta} - \mathbf{\Lambda}\hat{\boldsymbol{\theta}})\right]$ where $\hat{\boldsymbol{\theta}}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 5.2.

TABLE 5: BID PREFERENCES INCREASE DOMESTIC WINNERS WITH LIMITED IMPACT ON PRICES OR PARTICIPATION

	All Products			Pharmaceuticals			
	Log Price (1)	Num. Bidders (2)	Quality (3)	Log Price (4)	Num. Bidders (5)	Quality (6)	Domestic Winner (7)
log Standardized Quantity	-0.308***	0.043***	-0.002**	-0.030***	0.012***	-0.011***	0.004***
	(0.015)	(0.002)	(0.0006)	(0.002)	(0.002)	(0.0006)	(0.0008)
Preferenced * Policy Active	-0.004	-0.041***	-0.020***	-0.025**	-0.024	0.010	0.042***
	(0.010)	(0.011)	(0.005)	(0.011)	(0.020)	(0.007)	(0.005)
R^2	0.653	0.266	0.222	0.948	0.271	0.261	0.736
Observations	16,348,332	16,348,332	16,348,332	460,531	460,531	460,531	460,531
Outcome Mean	5.557	2.065	0.075	6.279	1.942	0.178	0.385
Constituent Terms	✓	✓	✓	✓	✓	✓	✓
Good fixed effects	✓	✓	✓	✓	✓	✓	✓
Month fixed effects	✓	✓	✓	✓	✓	✓	✓
*Year*Product*Size*Region* fixed effects	✓	✓	✓	✓	✓	✓	✓

This table estimates the Intent to Treat (ITT) of the bid preference policy from equation (5): $y_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt}$. The sample used is summarized in columns (3) and (6) of Table 1. In the All Products sample an item has Preferenced_{gt} = 1 if the type of good appears on the list of goods covesred by the preferences policy for that year. In the Pharmaceuticals sample, Preferenced_{gt} = 1 if the drug purchased is made both in Russia and abroad. PolicyActive_t = 1 during the part of the relevant year that the preferences policy was in effect. Standard errors are clustered by month and good.

TABLE 6: BID PREFERENCES ARE MORE EFFECTIVE WHEN IMPLEMENTED BY LESS EFFECTIVE BUREAUCRATS

	All Products			Pharmaceuticals			
	Log Price (1)	Num. Bidders (2)	Quality (3)	Log Price (4)	Num. Bidders (5)	Quality (6)	Domestic Winner (7)
log Standardized Quantity	-0.309***	0.042***	-0.0006	-0.027***	0.008***	-0.006***	0.003***
	(0.015)	(0.002)	(0.0006)	(0.002)	(0.002)	(0.0006)	(0.0007)
Bureaucrat FE * Preferenced * Policy Active	-0.082***	0.093***	-0.131***	-0.466***	1.12***	0.061	0.216***
	(0.019)	(0.020)	(0.044)	(0.090)	(0.201)	(0.066)	(0.048)
Organization FE * Preferenced * Policy Active	0.026*	-0.007	-0.006	0.084	0.451**	-0.006	-0.137***
	(0.015)	(0.010)	(0.036)	(0.103)	(0.193)	(0.099)	(0.044)
R^2	0.658	0.271	0.240	0.950	0.288	0.326	0.734
Observations	16,348,332	16,348,332	16,348,332	460,531	460,531	460,531	292,366
Outcome Mean	5.557	2.065	0.075	6.279	1.942	0.178	0.385
Constituent Terms	✓	✓	✓	✓	✓	✓	✓
Good fixed effects	✓	✓	✓	✓	✓	✓	✓
*Year*Product*Size*Region* fixed effects	✓	✓	✓	✓	✓	✓	✓
Month fixed effects	✓	✓	✓	✓	✓	✓	✓
ConnectedSet fixed effects	✓	✓	✓	✓	✓	✓	✓

This table estimates the triple-difference from equation (7): $y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \theta_b\hat{\alpha}_b + \theta_j\hat{\psi}_j + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \rho_b\text{Preferenced}_{gt}\hat{\alpha}_b + \rho_j\text{Preferenced}_{gt}\hat{\psi}_j + \eta_b\text{PolicyActive}_t\hat{\alpha}_b + \eta_j\text{PolicyActive}_t\hat{\psi}_j + \pi_b\text{Preferenced}_{gt} \times \text{PolicyActive}_t\hat{\alpha}_b + \pi_j\text{Preferenced}_{gt} \times \text{PolicyActive}_t\hat{\psi}_j + \varepsilon_{igt}$. The sample used is summarized in columns (3) and (6) of Table 1. In the All Products sample an item has Preferenced_{gt} = 1 if the type of good appears on the list of goods covered by preferences that year. In the Pharmaceuticals sample, Preferenced_{gt} = 1 if the drug purchased is made both in Russia and abroad. PolicyActive_t = 1 during the part of the year that the preferences policy was in effect. Bureaucrat and Organization FEs are the covariance-shrunk bureaucrat and organization effects estimated in section 5. Standard errors are clustered by month and good.

Online Appendix (For Web Publication Only)

Table of Contents

A Details on Text Analysis

This appendix provides some of the details of the procedure we use to categorize procurement purchases into groups of homogeneous products. We proceed in three steps. First, we transform the raw product descriptions in our data into vectors of word tokens to be used as input data in the subsequent steps. Second, we develop a transfer learning procedure to use product descriptions and their corresponding Harmonized System product codes in data on the universe of Russian imports and exports to train a classification algorithm to assign product codes to product descriptions. We then apply this algorithm to the product descriptions in our procurement data. Third, for product descriptions that are not successfully classified in the second step, either because the goods are non-traded, or because the product description is insufficiently specific, we develop a clustering algorithm to group product descriptions into clusters of similar descriptions.

Once our data is grouped into products, we create our main outcome of interest–unit prices—in three steps. First, we standardize all units to be in SI units (e.g. convert all lengths to meters). Second, for each good, we keep only the most frequent standardized units i.e. if a good is usually purchased by weight and sometimes by volume, we keep only purchases by weight. Third, we drop the top and bottom 5% of the unit prices for each good since in some cases the number of units purchased is off by an order of magnitude spuriously creating very large or very small unit prices due to measurement error in the quantity purchased.

A.1 Preparing Text Data

The first step of our procedure 'tokenizes' the sentences that we will use as inputs for the rest of the procedure. We use two datasets of product descriptions. First, we use the universe of customs declarations on imports and exports to & from Russia in 2011–2013. Second, we use the product descriptions in our procurement data described in Subsection 3.1. Each product description is parsed in the following way, using the Russian libraries for Python's Natural Language Toolkit[73]

1. Stop words are removed that are not core to the meaning of the sentence, such as "the", "and", and "a".
2. The remaining words are lemmatized, converting all cases of the same word into the same 'lemma' or stem. For example, 'potatoes' become 'potato'.
3. Lemmas two letters or shorter are removed.

We refer to the result as the *tokenized* sentence. For example the product description "NV-Print Cartridge for the Canon LBP 2010B Printer" would be broken into the following tokens: [cartridge, NV-Print, printer, Canon, LBP, 3010B]. [74] Similarly, the product description "sodium bicarbonate -

[73]Documentation on the Natural Language Toolkit (NLTK) can be found at http://www.nltk.org/

[74]The original Russian text reads as "картридж NV-Print для принтера Canon LBP 3010B" with the following set of Russian tokens: [картридж, NV-Print, принтер, Canon, LBP, 3010B].

solution for infusion 5%,200ml" would result in the following tokens: [sodium, bicarbonate, solution, infusion, 5%, 200ml].[75]

A.2 Classification

In the second step of our procedure we train a classification algorithm to label each of the sentences in the customs data with one of the H_C labels in the set of labels in the customs dataset, $\mathcal{H}_C$. To prepare our input data, each of the N_C tokenized sentences $\mathbf{t}_i$ in the customs dataset is transformed into a vector of token indicators and indicators for each possible bi-gram (word-pair), denoted by $\mathbf{x}_i \in \mathcal{X}_C$.[76] Each sentence also has a corresponding good classification $g_i \in \mathcal{G}_C$, so we can represent our customs data as the pair $\{\mathbf{X}_C, \mathbf{g}_C\}$ and we seek to find a classifier $\hat{g}_C(\mathbf{x}) : \mathcal{X}_C \to \mathcal{H}_C$ that assigns every text vector $\mathbf{x}$ to a product code.

As is common in the literature, rather than solving this multiclass classification problem in a single step, we pursue a "one-versus-all" approach and reduce the problem of choosing among G possible good classifications to G_C binary choices between a single good and all other goods, and then combine them (Rifkin & Klautau, 2004). We do this separately for each 2-digit product category. Each of the G_C binary classification algorithms generates a prediction $p_g(\mathbf{x}_i)$, for whether sentence i should be classified as good g. We then classify each sentence as the good with the highest predicted value:

$$\hat{g}_C(\mathbf{x}_i) = \underset{g \in \mathcal{G}_C}{\operatorname{argmax}} p_g(\mathbf{x}_i) \tag{A.1}$$

Each binary classifier is a logistic regression solving

$$\min_{\mathbf{w}_g, a_g} \frac{1}{N_C} \sum_{i=1}^{N_C} \frac{1}{\ln 2} \ln\left(1 + \mathrm{e}^{-y_{gi} \cdot (\mathbf{w}_g \cdot \mathbf{x}_i + a_g)}\right) \tag{A.2}$$

where

$$y_{gi} = \begin{cases} 1 & \text{if } g_i = g \\ -1 & \text{otherwise} \end{cases}$$

The minimands $\hat{\mathbf{w}}_g$ and $\hat{a}_g$ are then used to compute $p_g(\mathbf{x}_i) = \hat{\mathbf{w}}_g \cdot \mathbf{x}_i + \hat{a}_g$ with which the final classification is formed using equation (A.1). We implement this procedure using the Vowpal Wabbit library for Python.[77] This simple procedure is remarkably effective; when trained on a randomly selected half of the customs data and then implemented on the reamining data for validation, the classifications are correct 95% of the time. Given this high success rate without regularization,

[75] The original Russian text reads as "натрия гидрокарбонат - раствор для инфузий 5%,200мл" with the set of Russian tokens as: [натрия, гидрокарбонат, раствор, инфузия, 5%, 200мл].

[76] The customs entry "Electric Table Lamps Made of Glass" is transformed into the set of tokens: [electric, table, lamp, glass]. The original Russian reads as "лампы электрические настольные из стекла" and the tokens as: [электрический, настольный, ламп, стекло].

[77] See http://hunch.net/~vw/.

we decided not to try and impose a regularization penalty to improve out of sample fit. We also experimented with two additional types of classifiers. First, we trained a linear support vector machine with a hinge loss function.[78] That is, a classifier that solves

$$\min_{\mathbf{w}_g,a_g} \frac{1}{N_C}\sum_{i=1}^{N_C} \max\{0, 1 - y_{gi}\cdot(\mathbf{w}_g\cdot\mathbf{x}_i + a_g)\} \tag{A.3}$$

Second, we trained a set of hierarchical classifiers exploiting the hierarchical structure of the HS product classification. Each classifier is a sequence of sub-classifiers. The first sub-classifier predicts which 4-digit HS code corresponds to the text. Then, within each 4-digit code, the next classifier predicts the corresponding 6-digit code, etc, until the last classifier that predicts the full 10-digit code within each 8-digit category. Our main analysis of section 5.3 presented in figure 1 and table 2 is repeated using these alternative classifiers in figure D.1 panels C and D and in table E.4. As they show, the results are robust to these alternative classification methods.

Having trained the algorithm on the customs dataset, we now want to apply it to the procurement dataset wherever possible. This is known as transfer learning (see, for example Torrey & Shavlik (2009)). Following the terminology of Pang & Yang (2010), our algorithm $\hat{g}_C$ performs the task $\mathcal{T}_C = \{\mathcal{H}_C, g_C(\cdot)\}$ learning the function $g_C(\cdot)$ that maps from observed sentence data X to the set of possible customs labels $\mathcal{G}_C$. The algorithm was trained in the domain $\mathcal{D}_C = \{\mathcal{X}_C, F(X)\}$ where $F(\mathbf{X})$ is the probability distribution of $\mathbf{X}$. We now seek to transfer the algorithm to the domain of the procurement dataset, $\mathcal{D}_B = \{\mathcal{X}_B, F(\mathbf{X})\}$ so that it can perform the task $\mathcal{T}_B = \{\mathcal{H}_B, g_B(\cdot)\}$. Examples of the classification outcomes can be found in Tables A.1 (translated into English) and A.2 (in the original Russian). The three columns on the left present the tokens from the descriptions of goods in the procurement data, along with an identifying contract number and the federal law under which they were concluded. The columns on the right indicate the 10-digit HS code ('13926100000 - Office or school supplies made of plastics') that was assigned to all four of the goods using the machine learning algorithm. In addition, we present the tokenized customs entries that correspond to this 10 digit HS code.

The function to be learned and the set of possible words used are unlikely to differ between the two domains—A sentence that is used to describe a ball bearing in the customs data will also describe a ball bearing in the procurement data—so $\mathcal{X}_C = \mathcal{X}_B$, and $h_C(\cdot) = h_B(\cdot)$. The two key issues that we face are first, that the likelihoods that sentences are used are different in the two samples so that $F(\mathbf{X})_C \neq F(\mathbf{X})_B$. This could be because, for example, the ways that importers and exporters describe a given good differs from the way public procurement officials and their suppliers describe that same good. In particular, the procurement sentences are sometimes not as precise as those used in the trade data. The second issue is that the set of goods that appear in the customs data differs from the goods in the procurement data so that $\mathcal{H}_C \neq \mathcal{H}_B$. This comes about because non-traded

[78] A description of the support vector loss function (hinge loss), which estimates the mode of the posterior class probabilities, can be found in Friedman *et al.* (2013, 427)

Table A.1: Example Classification - English

Contract ID	Law	Product Description	HS10 Code	Example Import Entries
5070512	94FZ	folder, file, Erich, Krause, Standard, 3098, green	3926100000	product, office, made of, plastic
15548204	44FZ	cover, plastic, clear	3926100000	office, supply, made of, plastic, kids, school, age, quantity
16067065	44FZ	folder, plastic	3926100000	supply, office, cover, plastic, book
18267299	44FZ	folder, plastic, Brauberg	3926100000	collection, office, desk, individual, plastic, packaging, retail, sale

Table A.2: Example Classification - Russian

Contract ID	Law	Product Description	HS10 Code	Example Import Entries
5070512	94FZ	Папка, файл, Erich, Krause, Standard, 3098, зелёная	3926100000	изделие, канцелярский, изготовленный, пластик
15548204	44FZ	Обложка, пластиковый, прозрачный	3926100000	канцелярский, принадлежность, изготовленный, пластик, дети, школьный, возрасть, количество
16067065	44FZ	Скоросшиватель, пластиковый	3926100000	принадлежность, канцелярский, закладка, пластиковый, книга
18267299	44FZ	Скоросшиватель, пластиковый, Brauberg	3926100000	набор, канцелярский, настольный, индивидуальный, пластмассовый, упаковка, розничный, продажа

goods will not appear in the customs data, but may still appear in the procurement data.

To deal with these issues, we identify the sentences in the procurement data that are unlikely to have been correctly classified by $\hat{h}_C$ and instead group them into goods using the clustering procedure described in section A.3 below. We construct 2 measures of the likelihood that a sentence is correctly classified. First, the predicted value of the sentence's classification $\hat{g}_C(\mathbf{x}_i)$ as defined in (A.1). Second, the similarity between the sentence and the average sentence with the sentence's assigned classification in the *customs* data used to train the classifier.

To identify outlier sentences, we take the tokenized sentences that have been labeled as good g, $\mathbf{t}_g = \{\mathbf{t}_i : \hat{g}_C(\mathbf{x}_i) = g\}$ and transform them into vectors of indicators for the tokens $\mathbf{v}_{gi}$.[79] For each good, we then calculate the mean sentence vector in the customs data as $\mathbf{v}_g^C = \sum_{\mathbf{v}_{gi}, x_i \in \mathbf{X}^C} \mathbf{v}_{gi} / |\mathbf{t}_g|$.

[79]Note that these vectors differ from the inputs $\mathbf{x}_i$ to the classifier in two ways. First, they are specific to a certain good, and second, they omit bigrams of the tokens

Then, to identify outlier sentences in the procurement data, we calculate each sentence's normalized cosine similarity with the good's mean vector,

$$\theta_{gi} = \frac{\bar{s}_g - s(\mathbf{v}_{gi}, \mathbf{v}_g)}{\bar{s}_g} \tag{A.4}$$

where $s(\mathbf{v}_{gi}, \mathbf{v}_g) \equiv \cos(\mathbf{v}_{gi}, \mathbf{v}_g) = \frac{\mathbf{v}_{gi}\mathbf{v}_g}{\|\mathbf{v}_{gi}\|\|\mathbf{v}_g\|} = \frac{\sum_{k=1}^{K_g} t_{gik} t_{gk}}{\sqrt{\sum_{k=1}^{K_g} t_{gik}^2}\sqrt{\sum_{k=1}^{K_g} t_{gk}^2}}$ is the cosine similarity of the sentence vector $\mathbf{v}_{gi}$ with its good mean $\mathbf{v}_g$,[80] K_g is the number of tokens used in descriptions of good g, and $\bar{s}_g = \sum_{i=1}^{|\mathbf{t}_g|} s(\mathbf{v}_{gi}, \mathbf{v}_g)$ is the mean of good g's sentence cosine similarities. We deemed sentences to be correctly classified if their predicted value $\hat{g}_C(\mathbf{x}_i)$ was above the median and their normalized cosine similarity θ_{gi} was above the median. Figure D.1 panels A and B and Table E.4 show the robustness of our results to using the 45th or 55th percentile as thresholds.

A.3 Clustering

The third step of our procedure takes the misclassified sentences from the classification step and groups them into clusters of similar sentences. We will then use these clusters as our good classification for this group of purchases. To perform this clustering we use the popular K-means method. This method groups the tokenized sentences into k clusters by finding a centroid c_k for each cluster to minimize the sum of squared distances between the sentences and their group's centroid. That is, it solves

$$\min_{\mathbf{c}} \sum_{i=1}^{N} \|f(\mathbf{c}, \mathbf{t}_i) - \mathbf{t}_i\|^2 \tag{A.5}$$

where $f(\mathbf{c}, \mathbf{t}_i)$ returns the closest centroid to $\mathbf{t}_i$. To speed up the clustering on our large dataset we implemented the algorithm by mini-batch k-means. Mini-batch k means iterates over random subsamples (in our case of size 500) to minimize computation time. In each iteration, each sentence is assigned to it's closest centroid, and then the centroids are updated by taking a convex combination of the sentence and its centroid, with a weight on the sentence that converges to zero as the algorithm progresses (see Sculley (2010) for details).

The key parameter choice for the clustering exercise is k, the number of clusters to group the sentences into. As is common in the literature, we make this choice using the silhouette coefficient. For each sentence, its silhouette coefficient is given by

$$\eta(i) = \frac{b(i) - a(i)}{\max\{b(i), a(i)\}} \tag{A.6}$$

where $a(i)$ is the average distance between sentence i and the other sentences in the same cluster, and $b(i)$ is the average distance between sentence i and the sentences in the nearest cluster to

[80]Note that the cosine similarity ranges from 0 to 1, with 0 being orthogonal vectors and 1 indicating vectors pointing in the same direction.

sentence i's cluster. A high value of the silhouette coefficient indicates that the sentence is well clustered: it is close to the sentences in its cluster and far from the sentences in the nearest cluster. We start by using a k of 300 for each 2-digit product categories. For 2-digit product categories with an average silhouette coefficient larger than the overall average silhouette coefficient, we tried $k \in \{250, 200, 150, 100, 50, 25, 10, 7\}$ while for product categories with a lower than average silhouette coefficient we tried $k \in \{350,400,450,500,550, 600,650,700,750,800,850,900,950,1000\}$ until the average silhouette score was equalized across 2-digit product codes.

B Proofs of Propositions

B.1 Proof of Proposition 1

Proof. The suppliers choose their entry and bidding strategies to maximize expected profits. Working backwards from the second stage, when both firms enter, it is a dominant strategy for bidders to bid their fulfillment cost since bidder valuations are independent (see e.g. Milgrom, 2004). The winner is the bidder with the lowest fulfillment cost; she receives the contract at the other bidder's fulfillment cost. The expected profits from an auction in which firm i bids b_i are then $\mathbb{E}[\pi_i|b_i] = \mathbb{E}_{b_j}[b_j - b_i|b_j > b_i]\,\mathbb{P}(b_j > b_i)$ making the expected profits from the auction to bidder i, $\mathbb{E}[\pi_i] = \mathbb{E}_{b_i}[\mathbb{E}[\pi_i|b_i]]$.

Working back to the entry decisions, the two firms enter with probabilities q_F and q_L. If firm i pays the participation cost c_i and enters, with probability q_j firm j also enters and the auction takes place, yielding firm i expected profits of $\mathbb{E}[\pi_i]$, while with probability $1-q_j$, i is the only entrant and receives the contract at price $\bar{\theta}$ yielding expected profits of $\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i]$. If instead firm i chooses not to enter, her profits are zero but she does not have to pay the participation cost. The nature of the equilibrium depends on the size of the participation costs c_i. When participation costs are sufficiently small, both firms enter with certainty and the auction always takes place. For larger participation costs the equilibrium involves mixed strategies. In a mixed strategy equilibrium, the firms are indifferent between entering and not entering, pinning down the entry probabilities

$$q_j \mathbb{E}[\pi_i] + (1-q_j)(\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i]) = c_i \iff q_j = \frac{\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i] - c_i}{\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i] - \mathbb{E}[\pi_i]}, \tag{B.1}$$

where $i,j \in \{F,L\}$, $i \neq j$.

For the firms to be indifferent between entering and not entering, equation (B.1) must hold. Solving the equation requires us to derive expressions for $\mathbb{E}[b_i]$ and $\mathbb{E}[\pi_i]$. The distribution of the bids is given by the bidding functions $b_i = \bar{\theta}/\theta_i$ and the Pareto distributions of the productivities θ_i: $G_i(\theta_i) = 1 - \theta_i^{-\delta_i}$.

$$H_i(b) \equiv \mathbb{P}(b_i \leq b) = \mathbb{P}\left(\theta_i \geq \frac{\bar{\theta}}{b}\right) = \left(\frac{b}{\bar{\theta}}\right)^{\delta_i} \tag{B.2}$$

The expected bids are then simply $\mathbb{E}[b_i]=\int_0^{\bar{\theta}} b dH_i(b)=\frac{\delta_i}{1+\delta_i}\bar{\theta}$.

To derive expected profits from the auction $\mathbb{E}[\pi_i]$ we begin by considering expected profits conditional on a bidders fulfillment cost. Since the optimal bidding strategies are to bid the firm's true valuation, expected profits for a firm with valuation b_i are

$$\mathbb{E}[\pi_i|b_i]=\mathbb{E}_{b_j}[b_j-b_i|b_j>b_i]\mathbb{P}(b_j>b_i)=\int_{b_i}^{\bar{\theta}}(b_j-b_i)dH_j(b_j)$$
$$=\frac{\delta_j}{1+\delta_j}\bar{\theta}-b_i+b_i\left(\frac{b_i}{\bar{\theta}}\right)^{\delta_j}\frac{1}{1+\delta_j}, \tag{B.3}$$

where the final equality follows by inserting (B.2) and integrating. Now we can derive unconditional expected profits by the law of iterated expectations:

$$\mathbb{E}[\pi_i]=\mathbb{E}_{b_i}[\mathbb{E}[\pi_i|b_i]]=\int_0^{\bar{\theta}}\mathbb{E}[\pi_i|b_i]dH_i(b_i)=\left(\frac{1}{1+\delta_i}-\frac{1}{1+\delta_F+\delta_L}\right)\bar{\theta}. \tag{B.4}$$

Inserting these and the definition of the entry costs c_i into (B.1) and rearranging yields the statement in the proposition

$$q_i=\sqrt{\kappa(1-\alpha_c-\psi_c)}, \tag{B.5}$$

where $\kappa=\min\left\{[(1+\delta_F+\delta_L)/(1+\delta_L)]^2,1/(1-\alpha_c-\psi_c)\right\}$.

Turning to the expected prices, whenever neither or only one firm enters, the price is $\bar{\theta}$. When both enter, the price is the higher of the two bids.

$$\mathbb{P}(p\leq x)=\mathbb{P}(\max\{b_F,b_L\}\leq x)=H_F(x)H_L(x)=\left(\frac{x}{\bar{\theta}}\right)^{\delta_F+\delta_L} \tag{B.6}$$

As a result, the distribution and expectation of the log price when both firms enter is

$$\mathbb{P}(\log(p)\leq x)=\mathbb{P}(p\leq \mathrm{e}^x)=\left(\frac{\mathrm{e}^x}{\bar{\theta}}\right)^{\delta_F+\delta_L}$$
$$\mathbb{E}[\log(p)|\text{ both enter}]=\int_{-\infty}^{\log(\bar{\theta})}x\frac{\delta_F+\delta_L}{M^{\delta_F+\delta_L}}e^{(\delta_F+\delta_L)x}dx=\log(\bar{\theta})-\frac{1}{\delta_F+\delta_L} \tag{B.7}$$

The expected log price is then simply $\mathbb{E}[\log(p)]=q_Fq_L\mathbb{E}[\log(p)|\text{ both enter}]+(1-q_Fq_L)\log(\bar{\theta})$. Inserting (B.7) and the entry probabilities q_F and q_L yields expression (1) in the proposition.

The comparative statics on prices follow straightforwardly from equation (1). The comparative static on the number of bidders follows straightforwardly from noting that the expected number of entrants is q_F+q_L. $\square$

B.2 Proof of Proposition 2

Proof. In this setting it is optimal for bidder F to shade so that her bid net of the bid penalty is equal to her true fulfillment cost $b_F = \bar{\theta}/\gamma\theta_F$. However, when her shaded bid would have no chance of winning ($\theta_F < 1/\gamma$), she drops out and the contract is awarded to bidder L. This means that for any given bid, the preference regime lowers expected profits for foreign bidders and increases them for local bidders, as the policy intends. To see this, note that the expected profits of bids b_F and b_L are now

$$\begin{aligned}\mathbb{E}[\pi_F|b_F,\gamma] &= \mathbb{E}[\gamma(b_L - b_F)|b_L > b_F]\mathbb{P}(b_L > b_F) \\ \mathbb{E}[\pi_L|b_L,\gamma] &= \mathbb{E}[b_F - b_L|\bar{\theta} \geq b_F > b_L]\mathbb{P}(\bar{\theta} \geq b_F > b_L) + \mathbb{P}(\theta_F < 1/\gamma)(\bar{\theta} - b_L).\end{aligned} \tag{B.8}$$

For any particular bid, the profits to bidder F are shrunk by the penalty γ, forcing bidder F to bid more aggressively and lowering expected profits. For bidder L the probability of winning with any bid increases, and the bid penalty creates a discrete probability that bidder F drops out, both of which increase L's expected profits.

Consider the three cases in proposition 2 in turn.

Buyers with $\alpha_c + \psi_c \leq \underline{\mathbf{c}}$. In this case, both bidders enter the auction with certainty. Entering the auction is a best response to the other bidder entering whenever $\mathbb{E}[\pi_i|\gamma] - c_i > 0$. Expected profits are lower for bidder F and participation costs c_F are higher, so bidder F is the pivotal bidder for this case. Integrating bidder F's expected profits conditional on her bid (B.8) over all bids,

$$\mathbb{E}[\pi_F|\gamma < 1] = \int_0^M \mathbb{E}[\pi_F|b_F,\gamma < 1]dH_F(b_F|\gamma < 1) = \gamma^{1+\delta_F} M\left(\frac{1}{1+\delta_F} - \frac{1}{1+\delta_F+\delta_L}\right) \tag{B.9}$$

Setting (B.9) equal to c_F and rearranging yields the definition of $\underline{c}$ in the proposition. Since $\underline{c} < 1 - \left(\frac{1+\delta_L}{1+\delta_F+\delta_L}\right)^2$, both bidders enter the auction with or without the preferences and so participation is unchanged.

Since bidding behavior has changed, the expected price in the auction has changed. There are three possibilities:

$$p = \begin{cases} b_F & \text{if } b_L < b_F < \bar{\theta}, \\ \bar{\theta} & \text{if } b_L < M \leq b_F, \\ \gamma b_L & \text{if } b_F \leq b_L. \end{cases}$$

Combining these the distribution of prices is given by

$$\mathbb{P}(p \le x) = \begin{cases} H_F(x)H_L(x/\gamma) + \int_x^{x/\gamma}\int_{b_F}^{x/\gamma} h_L(b_L)db_L h_F(b_F)db_F & \text{if } 0 \le x \le \gamma\bar{\theta}, \\ H_F(x) + \int_x^{\bar{\theta}}\int_{b_F}^{\bar{\theta}} h_L(b_L)db_L h_F(b_F)db_F & \text{if } \gamma\bar{\theta} < x < \bar{\theta}, \\ 1 & \text{if } x = \bar{\theta} \end{cases}$$

$$= \begin{cases} \left(\frac{\delta_L}{\delta_F+\delta_L}\gamma^{-\delta_F-\delta_L} + \frac{\delta_F}{\delta_F+\delta_L}\right) H_F(x)H_L(x) & \text{if } 0 \le x \le \gamma\bar{\theta}, \\ \frac{\delta_L}{\delta_F+\delta_L}\gamma^{\delta_F} + \frac{\delta_F}{\delta_F+\delta_L} H_F(x)H_L(x) & \text{if } \gamma\bar{\theta} < x < \bar{\theta}, \\ 1 & \text{if } x = \bar{\theta} \end{cases}$$

In turn, the distribution of log prices is given by

$$\mathbb{P}(\log(p) \le x) = \mathbb{P}(p \le \mathrm{e}^x) = \begin{cases} \left(\frac{\delta_L}{\delta_F+\delta_L}\gamma^{-\delta_L} + \frac{\delta_F}{\delta_F+\delta_L}\gamma^{\delta_F}\right) \left(\frac{\mathrm{e}^x}{\bar{\theta}}\right)^{\delta_F+\delta_L} & \text{if } -\infty < x \le \log(\gamma\bar{\theta}), \\ \frac{\delta_L}{\delta_F+\delta_L}\gamma^{\delta_F} + \frac{\delta_F}{\delta_F+\delta_L}\gamma^{\delta_F}\left(\frac{\mathrm{e}^x}{\bar{\theta}}\right)^{\delta_F+\delta_L} & \text{if } \log(\gamma\bar{\theta}) < x < \log(\bar{\theta}), \\ 1 & \text{if } x = \log(\bar{\theta}) \end{cases}$$

making the expected log price in the auction

$$\begin{aligned} \mathbb{E}[\log(p)|\text{both enter}] = &\int_{-\infty}^{\log(\gamma\bar{\theta})} \frac{\delta_L\gamma^{-\delta_L} + \delta_F\gamma^{\delta_F}}{\bar{\theta}^{\delta_F+\delta_L}} x\mathrm{e}^{(\delta_F+\delta_L)x}\mathrm{d}x + \int_{\log(\gamma\bar{\theta})}^{\log(\bar{\theta})} \frac{\delta_F\gamma^{\delta_F}}{\bar{\theta}^{\delta_F+\delta_L}} x\mathrm{e}^{(\delta_F+\delta_L)x}\mathrm{d}x \\ &+ [1 - H_F(\bar{\theta})]\log(\bar{\theta}) \\ = &\log(\bar{\theta}) - \frac{\gamma^{\delta_F}\left(1 - \log\left(\gamma^{\delta_L}\right)\right)}{\delta_F+\delta_L}. \end{aligned} \tag{B.10}$$

Comparing (B.10) to the expected price without preferences (B.7), prices rise as long as $\gamma^{\delta_F}\left[1 - \log\left(\gamma^{\delta_L}\right)\right] < 1$.

Finally, the probability that the local bidder wins the auction when there are no preferences is

$$\mathbb{P}(L \text{ wins}) = \mathbb{P}(b_L < b_F) = \int_0^{\bar{\theta}} H_L(b_F|\gamma=1)dH_F(b_F|\gamma=1) = 1 - \frac{\delta_L}{\delta_F+\delta_L}, \tag{B.11}$$

while when there are preferences this increases to

$$\mathbb{P}(L \text{ wins}) = \mathbb{P}(b_L < b_F|\gamma<1) = \int_0^{\bar{\theta}} H_L(b_F|\gamma<1)dH_F(b_F|\gamma<1) = 1 - \gamma^{\delta_F}\frac{\delta_L}{\delta_F+\delta_L}. \tag{B.12}$$

Buyers with $\underline{c} < \alpha_c + \psi_c \le \bar{c}$. This case occurs when bidder L finds it worthwhile to enter the auction with certainty and bidder F's best response is to remain out of the auction with certainty. That is, when $\mathbb{E}[\pi_F|\gamma] - c_F < 0$ and $\mathbb{E}[\pi_L|\gamma] - c_L > 0$. In this case, since only L enters, the price is $\bar{\theta}$ with certainty, which is higher than in the absence of preferences since in the absence of preferences the auction always takes place with positive probability. Participation is therefore also lower, and since bidder L now wins with certainty, the probability that bidder L wins has increased.

The threshold $\underline{c}$ is defined in the previous case as the solution to $\mathbb{E}[\pi_L|\gamma]-c_L=0$. To find the upper threshold $\bar{c}$, we require an expression for $\mathbb{E}[\pi_L|\gamma]$:

$$\mathbb{E}[\pi_L|\gamma<1]=\int_0^{\bar{\theta}}\mathbb{E}[\pi_L|b_L,\gamma<1]dH_L(b_L|\gamma<1)=\bar{\theta}\left(\frac{1}{1+\delta_L}-\frac{\gamma^{\delta_F}}{1+\delta_F+\delta_L}\right). \tag{B.13}$$

Setting (B.13) equal to c_L and rearranging yields the definition of $\underline{c}$ in the proposition.

Buyers with $\bar{c}<\alpha_c+\psi_c$. This case occurs when neither bidder finds it optimal to enter with certainty: $\mathbb{E}[\pi_i|\gamma]-c_i<0\forall i$ and so the equilibrium is in mixed strategies. As in proposition 1, the entry probabilities are given by

$$q_i=\frac{\bar{\theta}-\mathbb{E}[\bar{\theta}/\theta_j]-c_j}{\bar{\theta}-\mathbb{E}[b_j]-\mathbb{E}[\pi_j|\gamma<1]}.$$

In this case the expected price is given by

$$\mathbb{E}[\log(p)]=\log(\bar{\theta})-q_Fq_L(\log(\bar{\theta})-\mathbb{E}[\log(p)|\text{both enter}])$$

Inserting the entry probabilities and the price equation (B.10) and rearranging, the expected price when there are preferences is lower whenever

$$\begin{aligned}
&q_F(\gamma<1)q_L(\gamma<1)(\log(\bar{\theta})-\mathbb{E}[\log(p)|\text{both enter, }\gamma<1])\\
&-q_F(\gamma=1)q_L(\gamma=1)(\log(\bar{\theta})-\mathbb{E}[\log(p)|\text{both enter, }\gamma=1])\geq 0\\
&\iff -\log\left(\gamma^{\delta_L}\right)-\frac{\delta_L}{1+\delta_F}\left(1-\gamma^{1+\delta_F}\right)\geq 0
\end{aligned} \tag{B.14}$$

Noting that (B.14) holds with equality when $\gamma=1$ and that the left hand side of (B.14) has slope $-\delta_L\left(\gamma^{-1}-\gamma^{\delta_F}\right)<0\;\forall\gamma<1$ shows that (B.14) holds for all $\gamma<1$. Participation in the auction is $\mathbb{E}[N]=q_F+q_L$. When there are no preferences

$$\mathbb{E}[N|\gamma=1]=q_F(\gamma=1)+q_L(\gamma=1)=2\frac{1+\delta_F+\delta_L}{1+\delta_L}\sqrt{1-\alpha_c-\psi_c}, \tag{B.15}$$

while with preferences participation is

$$\begin{aligned}
\mathbb{E}[N|\gamma<1]&=q_F(\gamma<1)+q_L(\gamma<1)\\
&=\left(\frac{1}{\gamma^{\delta_F}}+\frac{1}{\gamma^{1+\delta_F}+(1-\gamma^{1+\delta_F})\frac{1+\delta_F+\delta_L}{1+\delta_F}}\right)\frac{1+\delta_F+\delta_L}{1+\delta_L}\sqrt{1-\alpha_c-\psi_c}.
\end{aligned} \tag{B.16}$$

Comparing (B.15) to (B.16) shows that participation increases whenever

$$\frac{1}{\gamma^{\delta_F}} + \frac{1}{1+\frac{\delta_L}{\delta_F+\delta_L}(1-\gamma^{1+\delta_F})} > 2 \tag{B.17}$$

Equation (B.17) is implied by our assumption that we are in the case where $\gamma^{\delta_F}\left[1+\frac{\delta_L}{\delta_F+\delta_L}(1-\gamma^{1+\delta_F})\right] < 1$

Finally, to see that the probability that bidder L wins the contract at auction increases by more than in case 1 note that the probability that bidder L wins the contract is given by $q_F q_L \mathbb{P}(b_F < b_L)$. The probability that bidder L wins will increase by more if $q_F(\gamma=1)q_L(\gamma=1) < q_F(\gamma<1)q_L(\gamma<1)$. Computing the components of this

$$\frac{q_F(\gamma=1)}{q_F(\gamma<1)} = \frac{\bar{\theta}-\mathbb{E}[\bar{\theta}/\theta_L]-\mathbb{E}[\pi_L|\gamma<1]}{\bar{\theta}-\mathbb{E}[\bar{\theta}/\theta_L]-\mathbb{E}[\pi_L|\gamma=1]} = \gamma^{\delta_F}$$
$$\frac{q_L(\gamma=1)}{q_L(\gamma<1)} = \frac{\bar{\theta}-\mathbb{E}[\bar{\theta}/\theta_F]-\mathbb{E}[\pi_F|\gamma<1]}{\bar{\theta}-\mathbb{E}[\bar{\theta}/\theta_F]-\mathbb{E}[\pi_F|\gamma=1]} = 1+\frac{\delta_L}{1+\delta_F}\left(1-\gamma^{1+\delta_F}\right)$$

Combining these two components shows that the statement is correct as long as

$$\gamma^{-\delta_F} > \left[1+\frac{\delta_L}{\delta_F+\delta_L}\left(1-\gamma^{1+\delta_F}\right)\right] \tag{B.18}$$

Condition is implied by the condition stated at the top of the proposition that $\gamma^{-\delta_F} > 1-\log(\gamma^{\delta_L})$. To see this, note that both conditions are decreasing in γ and that their limits are the same as γ approaches 1 from below. Then, note that the slope of the right-hand-side of the condition in the proposition is steeper than the slope of condition (B.18): The slope of the condition in the proposition is $-\delta_L\gamma^{-1}$ while the slope of condition (B.18) is $-\frac{\delta_F}{\delta_L+\delta_F}(1+\delta_F)\gamma^{\delta_F}$ which is flatter since rearranging

$$-\delta_L\gamma^{-1} < -\frac{\delta_F}{\delta_L+\delta_F}(1+\delta_F)\gamma^{\delta_F} \iff \frac{\delta_L}{\delta_F}\frac{\delta_L+\delta_F}{1+\delta_F} > \gamma^{1+\delta_F} \tag{B.19}$$

and both terms on the left are larger than 1 while the term on the right is smaller than 1. Hence, the condition in the proposition implies condition (B.18). □

C Identification of Bureaucrat and Organization Effects with Multiple Connected Sets

As shown in Abowd *et al.* (2002), it isn't possible to identify all the bureaucrat and organization effects. In particular, they show that (a) the effects are identified only within connected sets of bureaucrats and organizations; and (b) within each connected set s containing $N_{b,s}$ bureaucrats and $N_{o,s}$ organizations, only the group mean of the lhs variable, and $N_{b,s}-1+N_{o,s}-1$ of the bureaucrat and organization effects are identified. More generally, within each connected set, we can identify $N_{b,s}+N_{o,s}-1$ linear combinations of the bureaucrat and organization effects.

To see this explicitly, write the model as

$$\mathbf{p}=\mathbf{X}\boldsymbol{\beta}+\mathbf{B}\boldsymbol{\alpha}+\mathbf{F}\boldsymbol{\psi} \tag{C.1}$$

where $\mathbf{p}$ is the $N\times 1$ vector of item prices; $\mathbf{X}$ is an $N\times k$ matrix of control variables, $\mathbf{B}$ is the $N\times N_b$ design matrix indicating the bureaucrat responsible for each purchase; $\boldsymbol{\alpha}$ is the $N_b\times 1$ vector of bureaucrat effects; $\mathbf{F}$ is the $N\times N_o$ design matrix indicating the organization responsible for each purchase; and $\boldsymbol{\psi}$ is the $N_o\times 1$ vector of organization effects.

Suppressing $\mathbf{X}\boldsymbol{\beta}$ for simplicity, the OLS normal equations for this model are

$$\begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\alpha}}_{OLS} \\ \hat{\boldsymbol{\psi}}_{OLS} \end{bmatrix} = \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \mathbf{p} \tag{C.2}$$

As Abowd *et al.* (2002) show, these equations do not have a unique solution because $[\mathbf{B}\,\mathbf{F}]'[\mathbf{B}\,\mathbf{F}]$ only has rank $N_b+N_o-N_s$, where N_s is the number of connected sets. As a result, to identify a particular solution to the normal equations, we need N_s additional restrictions on the αs and ψs.

Abowd *et al.* (2002) add N_s restrictions setting the mean of the person effects to 0 in each connected set. They also set the grand mean of the firm effects to 0. However, this makes it difficult to compare across connected sets since all the firm effects are interpreted as deviations from the grand mean, which is a mean across connected sets. Instead, we will add $2N_s$ restrictions setting the mean of the bureaucrat and organization effects to 0 within each connected set. These N_s additional constraints also allow us to identify S connected set means $\gamma_s=\bar{\alpha}_s+\bar{\psi}_s$ which facilitate comparison across connected sets and allow us to interpret the variances of the estimated bureaucrat and organization effects as lower bounds on the true variances of the bureaucrat and organization effects.

Specifically, we augment the model to be

$$\mathbf{p}=\mathbf{B}\tilde{\boldsymbol{\alpha}}+\mathbf{F}\tilde{\boldsymbol{\psi}}+\mathbf{S}\boldsymbol{\gamma} \tag{C.3}$$

where $\mathbf{S}$ is the $N\times N_s$ design matrix indicating which connected set each item belongs to; $\boldsymbol{\gamma}$ is the $N_s\times 1$ vector of connected set effects; and we add the restriction that $\tilde{\alpha}$ and $\tilde{\psi}$ have mean zero in each connected set. Our fixed effects estimates thus solve the normal equations of this augmented model, plus $2N_s$ zero-mean restrictions:

$$\begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} & \mathbf{S} \end{bmatrix} \\ \begin{bmatrix} \mathbf{S_b} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S_o} & \mathbf{0} \end{bmatrix} \end{bmatrix} \begin{bmatrix} \hat{\alpha} \\ \hat{\psi} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \mathbf{p} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \tag{C.4}$$

where $\mathbf{S}_b$ is the $N_s \times N_b$ design matrix indicating which connected set each bureaucrat belongs to, and $\mathbf{S}_o$ is the $N_s \times N_o$ design matrix indicating which connected set each organization belongs to.

The following proposition describes the relationship between these estimators and the bureaucrat and organization effects.

Proposition 3 (**Identification**). *If the true model is given by (C.1), then $\hat{\alpha}$, $\hat{\psi}$, and $\hat{\gamma}$, the estimators of $\tilde{\alpha}$, $\tilde{\psi}$ and γ in the augmented model (C.3) that solve the augmented normal equations (C.4) (i) are uniquely identified, and (ii) are related to the true bureaucrat and organization effects $\boldsymbol{\alpha}$ and ψ by*

$$\begin{bmatrix} \hat{\alpha} \\ \hat{\psi} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\alpha} - \mathbf{S_b}'\overline{\boldsymbol{\alpha}} \\ \psi - \mathbf{S_o}'\overline{\psi} \\ \overline{\boldsymbol{\alpha}} + \overline{\psi} \end{bmatrix} \tag{C.5}$$

where $\overline{\boldsymbol{\alpha}}$ is the $N_s \times 1$ vector of connected-set bureaucrat effect means, and $\overline{\psi}$ is the $N_s \times 1$ vector of connected-set organization effect means.

Proof. We will prove each part of the result separately. To see uniqueness, first note that the standard normal equations for (C.3) only has rank $N_b + N_o - N_s$. To see this, we note that $\mathbf{B}\mathbf{S_b}' = \mathbf{F}\mathbf{S_o}' = \mathbf{S}$ and so $2N_s$ columns of the $N \times (N_b + N_o + N_s)$ matrix $[\mathbf{B}\,\mathbf{F}\,\mathbf{S}]$ are collinear. However, the $2N_s$ restrictions $\mathbf{S_b}\hat{\alpha} = 0$ and $\mathbf{S_o}\hat{\psi} = 0$ are independent of the standard normal equations, so the first matrix in (C.4) has rank $N_b + N_o + N_s$ and hence the solution to (C.4) is unique.

To see the second part, it suffices to show that (C.5) solves (C.4). First, substitute the estimators out of (C.4) using (C.5) and substitute in the true model using (C.1) to rewrite (C.4) as

$$\begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \left[\mathbf{B}(\boldsymbol{\alpha} - \mathbf{S_b}'\overline{\boldsymbol{\alpha}}) + \mathbf{F}\left(\psi - \mathbf{S_o}'\overline{\psi}\right) + \mathbf{S}\left(\overline{\boldsymbol{\alpha}} + \overline{\psi}\right)\right] \\ \mathbf{S_b}(\boldsymbol{\alpha} - \mathbf{S_b}'\overline{\boldsymbol{\alpha}}) \\ \mathbf{S_o}\left(\psi - \mathbf{S_o}'\overline{\psi}\right) \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} [\mathbf{B}\boldsymbol{\alpha} + \mathbf{F}\psi] \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}$$

From here, noting again that $\mathbf{B}\mathbf{S_b}' = \mathbf{F}\mathbf{S_o}' = \mathbf{S}$; that $\mathbf{S_b}\boldsymbol{\alpha}$ is an $N_s \times 1$ vector in which each entry is the sum of the bureaucrat effects; and that $\mathbf{S_o}\psi$ is an $N_s \times 1$ vector in which each entry is the sum of the organization effects, shows that the two sides are equal, yielding the result. □

The above analysis focuses on the simple case in which there are no other covariates in the model. In the more general model with covariates it is not always possible to separately identify the connected set intercepts γ, particularly when the covariates $\mathbf{X}$ include categorical variables. Nevertheless, the identification of the bureaucrat effects $\tilde{\alpha}$ and organization effects $\tilde{\psi}$ remains as above. In our empirical application we have categorical covariates and so we focus on the bureaucrat- and organization- effects and do not results on the connected set intercepts γ.

D Additional Results on Event Studies to Identify the Effectiveness of Individuals and Organizations

In Sub-section 5.1 we argue that using event studies around the time that organizations change the bureaucrat they work with can identify their effectiveness. In this appendix, we show that this argument is robust to changing a series of choices made in constructing the event studies.

In figure D.1 we consider the choices made in how the sample was built for the analysis. As described in Appendix A, we deemed contract descriptions to be correctly classified whenever their predicted value and their normalized cosine similarity with their labeled good's mean vector were both above the median. In Panel A we instead classify them as correct whenever they are above the 45th percentile, and in panel B we use the 55th percentile as our threshold. The results are essentially unchanged. Our baseline classifier uses the logistic function as its objective function, which performs very well. Nevertheless, in Panel C, we instead use a support vector machine (SVM) objective function. And in Panel D we train a sequence of hierarchical classifiers exploiting the hierarchical structure of the HS product codes (details are in Appendix A). In both cases, the results are unchanged. Finally, in panel E, we trim the top and bottom 2.5% of each product rather than the 5% we use in our baseline data. Again, the results are unaffected.

In Figures D.2 and D.3 we change a series of the choices made in constructing the event studies. In figure D.2 we vary the units of time we use to define the spells that we combine to create events. In Figure 1 we define a spell as a sequence of two *weeks*, separated by fewer than 400 days. In Panel A, rather than weeks, we use days. In Panel B we use fotnights. In panel C we use months. And in Panel D we define a spell as a sequence of *three* weeks instead of two. The results are very similar in all cases. In figure D.3 we consider four more design choices. In Panel A we use a coarser categorization of the effectiveness of the bureaucrats in each event: We use terciles instead of quartiles. In Panel B we use a global ranking of bureaucrats (instead of a separate ranking for each semester as in Figure 1). Finally, we consider spells in which the weeks are separated by up to 350 days (Panel C) or 450 days (Panel D) rather than 400 days. In all cases, the results are very similar.

Our main event study studies the prices paid by organizations around the time they switch the bureaucrat they work with. Figure D.4 considers two other such changes: Bureaucrats switching which good they buy (Panel A) and organizations switching which good they buy (Panel B). Again, the results strongly support the use of switches to identify the effectiveness of bureaucrats and organizations. Table D.1 displays the data underlying our main event study in Figure 1 along with some additional summary statistics on the event study (the sample sizes in columns (1) and (2) and the time gaps between event time periods in columns (7)–(9)). Table D.2 compares the sample used in the event study to the analysis sample, showing that the two samples are comparable.

Figure D.1: Robustness of Event Studies to Alternative Text Classifiers, Classification Accuracy Thresholds, and Outlier Trimming

Panel A: Classifier Accuracy Threshold 45th Percentile

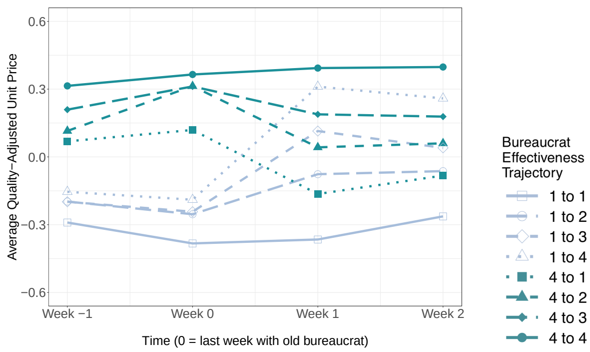

Panel B: Classifier Accuracy Threshold 55th Percentile

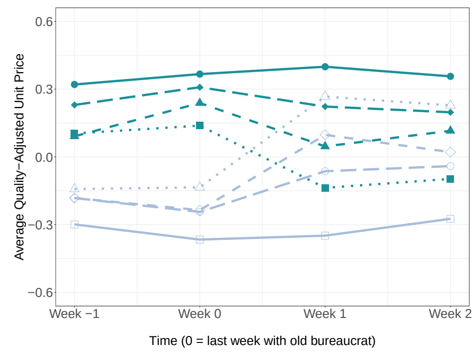

Panel C: Support Vector Machine Classifier

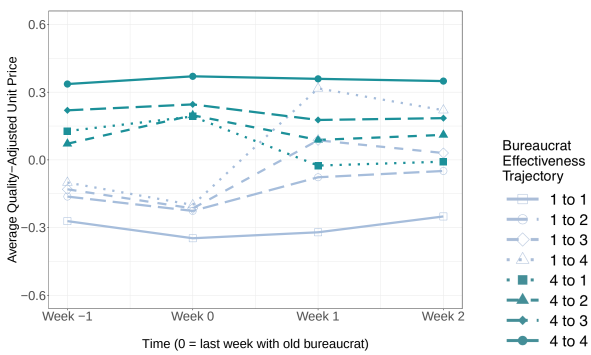

Panel D: Hierarchical Model Classifier

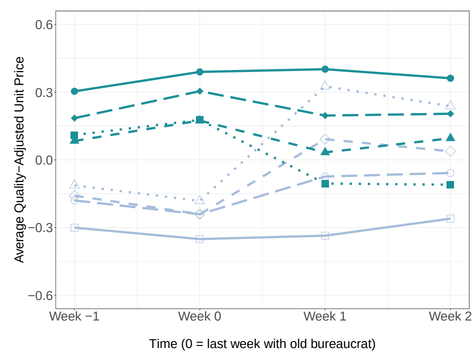

Panel E: Dropping Top and Bottom 2.5% Outliers

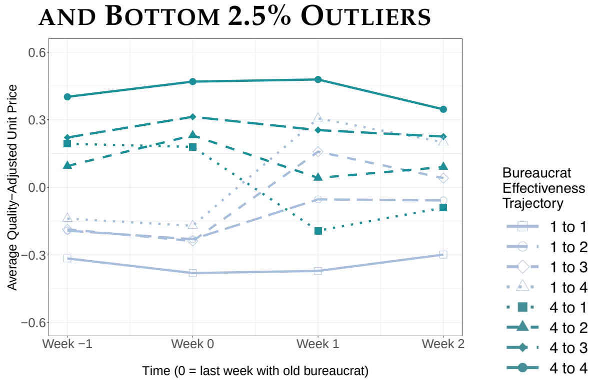

Notes: Each panel in the figure is analogous to Figure 1 (see notes to that figure for details of construction), with the following changes. Rather than requiring our classifier's predicted value and normalized cosine similarity to be above the median, we require them to be above the 45th percentile (Panel A) or the 55th percentile (Panel B). The classifier in Panel C uses a support vector machine objective function rather than a logistic function. The classifier in Panel D is a hierarchical series of classifiers exploiting the hierarchical structure of HS codes. See appendix A for details. Finally, in Panel E, we trim the top and bottom 2.5% of each product rather than the 5% we use in our baseline data.

FIGURE D.2: ROBUSTNESS OF EVENT STUDIES TO DESIGN CHOICES (1)

PANEL A: DAYS AS EVENT TIME

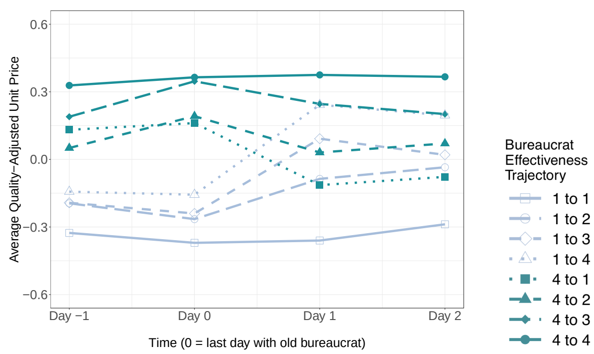

PANEL B: FORTNIGHTS AS EVENT TIME

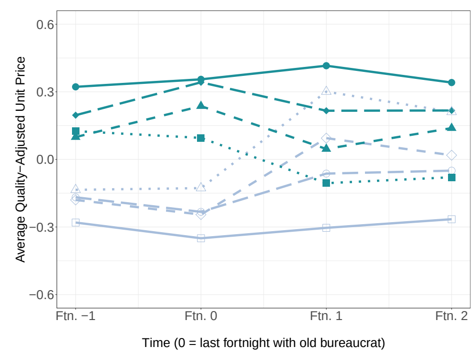

PANEL C: MONTHS AS EVENT TIME

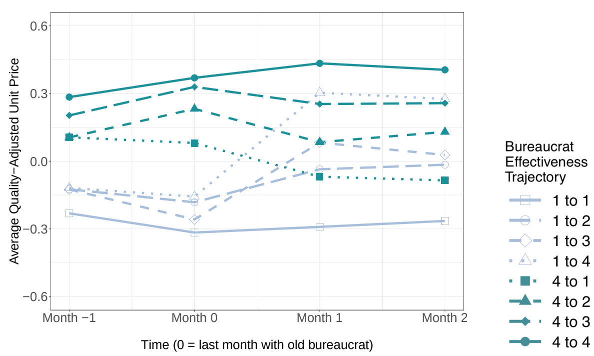

PANEL D: BALANCED PANEL $\pm$ 3 WEEKS

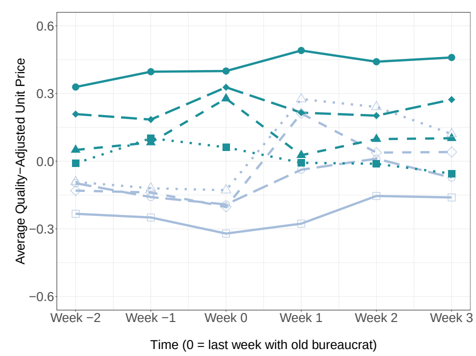

Notes: Each panel in the figure is analogous to Figure 1 (see notes to that figure for details of construction), with the following changes. In Panel A, rather than requiring the bureaucrat-organization pair to work together in two separate weeks, we require the pair to work together on two separate days. In Panel B, two separate fortnights; and in Panel C, two separate months. In Panel D we require bureaucrat-organization pairs to work together in three separate weeks.

Figure D.3: Robustness of Event Studies to Design Choices (2)

Panel A: Classifying Bureaucrats into Terciles

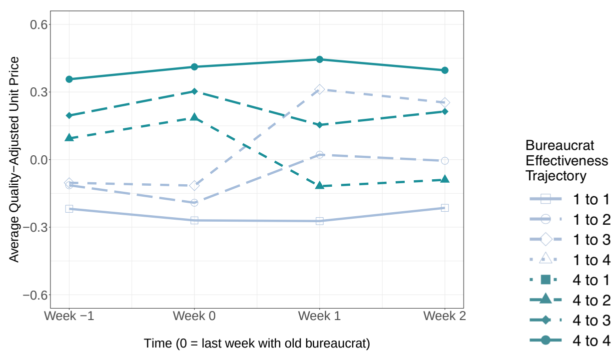

Panel B: Global Ranking of Bureaucrats

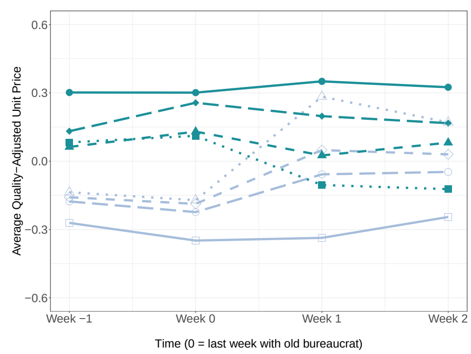

Panel C: 350 Day Spell Length

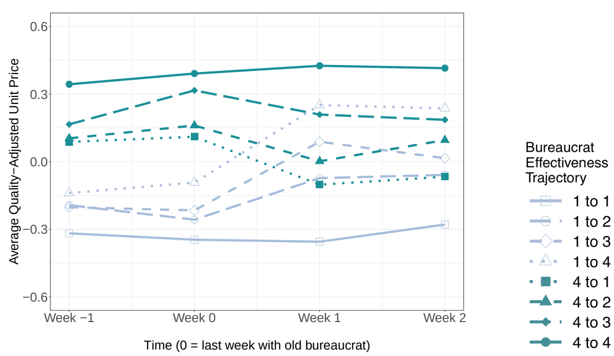

Panel D: 450 Day Spell Length

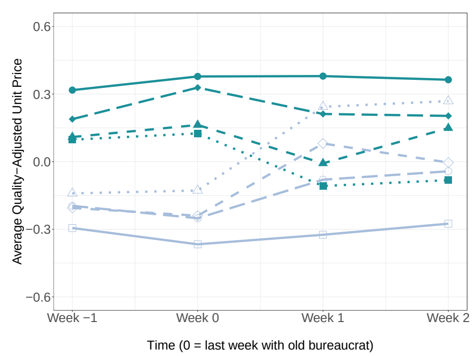

Notes: Each panel in the figure is analogous to Figure 1 (see notes to that figure for details of construction), with the following changes. In Panel A, we categorize bureaucrats by terciles rather than quartiles. In panel B, we construct quartiles by ranking bureaucrats based on the entire sample period rather than each semester separately. Rather than defining spells as weeks separated by fewer than 400 days as in Figure 1, we require them to be separated by 350 days (Panel C) or 450 days (Panel D).

FIGURE D.4: EVENT STUDY OF PROCUREMENT PRICES AROUND TIMES BUREAUCRATS AND ORGANIZATIONS SWITCH GOODS

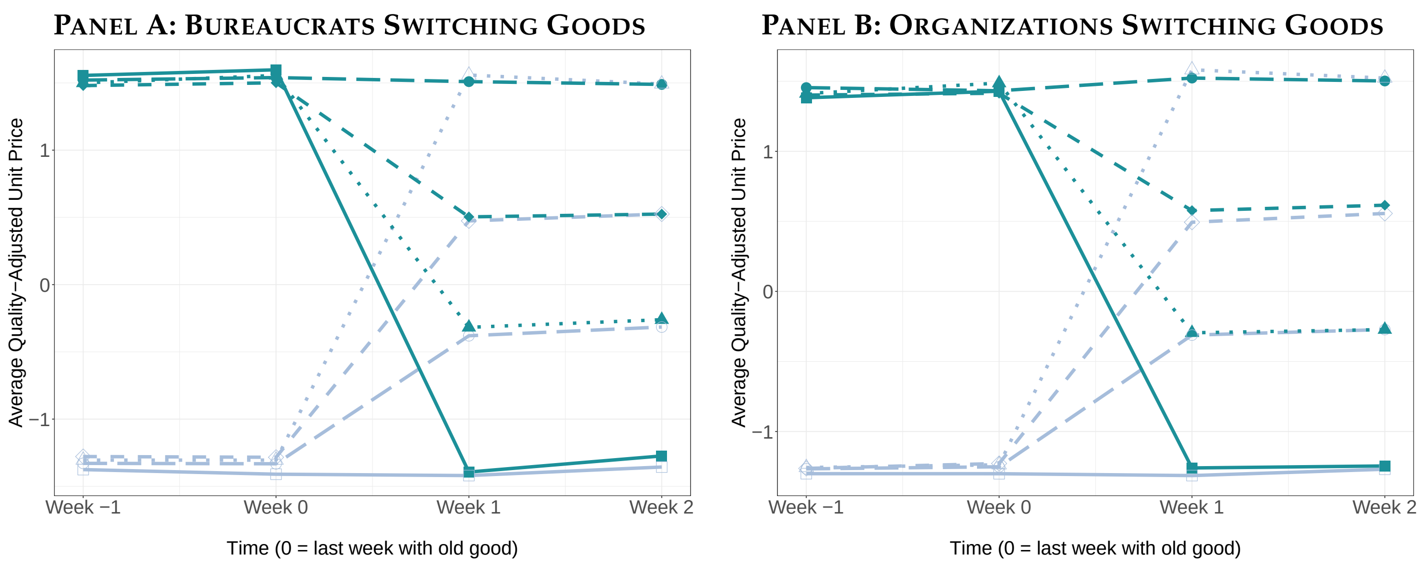

Notes: Each panel in the figure is analogous to Figure 1 that studies price changes around the time that organizations switch the bureaucrat making their purchases (see notes to that figure for details of construction). Panel A shows price changes around the time that bureaucrats switch the good they are purchasing. Panel B shows price changes around the time that organizations switch the good they are purchasing.

Table D.1: Event Studies Summary Statistics

Origin/destination Quartile*	Number of Moves (1)	Number of Observations (2)	Mean Log Residuals of Bureaucrat Movers: Week -1 (3)	Week 0 (4)	Week 1 (5)	Week 2 (6)	Mean Weeks Betw. Cols: (3)-(4) (7)	(4)-(5) (8)	(5)-(6) (9)
1 to 1	5,685	231,239	-0.303	-0.370	-0.354	-0.278	12.183	28.461	11.632
1 to 2	5,111	219,899	-0.181	-0.243	-0.101	-0.041	12.203	24.218	12.635
1 to 3	3,305	145,585	-0.196	-0.204	0.050	0.010	13.531	28.926	13.021
1 to 4	1,729	69,805	-0.152	-0.138	0.239	0.247	13.711	36.492	16.547
2 to 1	5,531	221,186	-0.052	-0.089	-0.207	-0.182	13.019	26.553	12.925
2 to 2	8,166	414,882	-0.042	-0.060	-0.024	-0.027	12.050	26.899	12.773
2 to 3	6,127	275,958	-0.020	-0.037	0.088	0.039	12.437	27.809	14.576
2 to 4	2,309	88,508	0.031	0.008	0.254	0.185	12.948	37.759	16.010
3 to 1	3,593	139,741	0.066	0.016	-0.113	-0.167	15.783	24.878	11.246
3 to 2	5,889	259,733	-0.006	0.050	0.019	-0.003	13.304	24.527	12.343
3 to 3	5,726	255,334	0.021	0.088	0.126	0.133	15.624	25.689	13.841
3 to 4	2,870	117,135	0.196	0.179	0.290	0.228	13.305	30.488	16.974
4 to 1	1,415	58,270	0.096	0.115	-0.102	-0.064	15.498	31.873	12.212
4 to 2	1,666	73,327	0.100	0.139	-0.001	0.121	15.569	29.774	12.139
4 to 3	2,254	93,073	0.204	0.334	0.249	0.206	15.395	30.908	13.196
4 to 4	2,614	117,670	0.321	0.380	0.391	0.366	15.673	27.332	14.974
Totals	63,990	2,781,345							

The table shows information on events in which organizations switch bureaucrats used in Figure 1. The sample used is the All Products-Analysis Sample summarized in Table 1. Events are defined using the procedure described in detail in Sub-section 5.1. We define an employment spell as a sequence of at least two weeks a bureaucrat-organization pair conducts purchases together, with the weeks less than 400 days apart. Wherever possible, we then match an employment spell (event time $\leq$ 0) with the earliest future spell (event time > 0) involving the same organization but a different bureaucrat. This change of bureaucrats then constitutes an event (event time = 0). We classify the two bureaucrats involved in the event using the average quality-adjusted price they achieve in purchases they make for *other* organizations during the half-year that the spell ends (for the earlier spell) or starts (for the later spell). We run equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. This regression regresses the price achieved in an auction on log quantity, good fixed effects, month fixed effects, interactions between 2-digit HS product categories, years, regions, and lot size, as explained in detail in Sub-section 5.2. Using the price residuals, we then classify bureaucrats by the average they achieve in purchases they make for other organizations. We assign this bureaucrat-average quality-adjusted price to the relevant quartile of the distribution of the average quality-adjusted prices of all bureaucrats that themselves are part of an event in the same half-year as the bureaucrat in question.

Table D.2: Comparing Event Study Data

	Full Sample	Event Study Data
(1) # of Bureaucrats	37,722	6,345
(2) # of Organizations	44,560	17,248
(3) # of Connected Sets	616	289
(4) # of Bureaucrats with >1 Org.	11,320	4,806
(5) # of Organizations with >1 Bur.	37,536	16,968
(6) Mean # of Bureaucrats per Org.	6.02	4.81
(7) Mean # of Organizations per Bur.	7.12	13.1
(8) # of Federal Organizations	1,583	147
(9) # of Regional Organizations	15,530	6,918
(10) # of Municipal Organizations	27,447	10,183
(11) # of Health Organizations	7,231	4,215
(12) # of Education Organizations	25,271	9,273
(13) # of Internal Affairs Organizations	668	136
(14) # of Agr/Environ Organizations	255	92
(15) # of Other Organizations	11,135	3,532
(16) # of Goods	15,649	12,964
(17) Mean # of Goods Per Bur.	93.2	124
(18) # of Regions	86	86
(19) Mean # of Regions per Bur.	1	1
(20) # of Auction Requests	1,871,717	378,297
(21) Mean # of Requests per Bur.	49.6	59.6
(22) Mean # of Applicants	3.46	3.62
(23) Mean # of Bidders	2.07	2.07
(24) Mean Reservation Price	0.134	0.122
(25) Quantity Mean	1,124	1,022
Median	27	35.1
SD	174,951	115,045
(26) Total Price Mean (bil. USD)	81.2	64
Median	4.74	3.3
SD	482	422
(27) Unit Price Mean (bil. USD)	55.6	38.4
Median	0.18	0.086
SD	19,168	1,214
(28) Mean # of Contract Renegotiations (log)	0.133	0.117
(29) Mean Size of Cost Over-run	-0.002	-0.002
(30) Mean Length of Delay in Days (log)	0.057	0.082
(31) Mean 1[End User Complained about Contract]	0.001	0.001
(32) Mean 1[Contract Cancelled]	0.009	0.016
(33) Mean 1[Product is of Substandard Quality]	0.009	0.006
(34) # of Observations	16,348,332	4,042,144
(35) Total Procurement Volume (bil. USD)	629	122

The table reports summary statistics for two samples. Organizations working in Internal Affairs include police, emergency services, local administration, taxes, and transportation. Organizations working in Agriculture or the Environment include environmental protection funds, agricultural departments and nature promotion agencies. The Other category includes funds, monitoring agencies, and land cadasters, among many others. All sums are measured in billions of US dollars at an exchange rate of 43 rubles to 1 US dollar.

E Additional Results on Variance Decomposition

This appendix presents additional results on the variance decomposition discussed in sections 5.2 and 5.3. Sub-section E.1 presents further evidence in support of the log-linear specification (3) used in the variance decomposition. Sub-section E.2 presents additional results showing the robustness of the findings to various design choices.

E.1 Misspecification

The model we have estimated assumes that the price achieved is approximately log-linear in the bureaucrat and organization effects. Three pieces of evidence suggest that match-based forms of endogenous mobility that would violate the identifying assumptions underlying our interpretation of the results from our empirical model rarely occur in Russian public procurement. First, the event studies in Sub-section 5.1 provide direct visual evidence that the price paid is approximately log-linear in the bureaucrat and organization effects. We saw no evidence of sorting on match effects in Figure 1.

Second, a direct piece of evidence in support of the log-linearity assumption comes from studying the distribution of the residuals across bureaucrat and organization effect deciles. If the log-linear specification was substantially incorrect, we would expect to see systematic patterns in the residuals. For example, positive match effects would lead the residuals to be large when the bureaucrat and organization are both in the top deciles of effectiveness. Panel A of Figure E.1 shows a heat map of residuals. The map reveals no clear patterns in the residuals. Panel B shows an analogous heat map of residuals from running (3) in levels rather than logs. The figure provides clear evidence that such a model is mis-specified, leading to systematically large residuals especially in the top right of the figure, where both the bureaucrat and organization are in the top deciles of effectiveness.

Third, we reestimate equation (3) but include fixed effects for each bureaucrat-organization pair, allowing for arbitrary patterns of complementarity between bureaucrats and organizations (see also Card *et al.*, 2013). If there are indeed strong or moderate match effects that our model omits, then we expect this pair effect model to fit significantly better. The pair effect model does not fit the data much better than our baseline model: adding pair effects decreases the RMSE of the residuals from 1.147 to 1.121 and increases the adjusted R^2 from 0.963 to 0.964, and the pair effects have a much smaller variance than the procurer effects from the log-linear model (results available from the authors upon request).

Overall, we do not find evidence supporting a rejection of our log-linearity assumption.

E.2 Robustness to design choices

In figure 2 discussed in section 5.4, we argue that our results are robust to focusing on more homogeneous subsets of goods in our sample. We use the measure of the scope for quality differentiation developed by Sutton (1998). As an alternative, we repeat the exercise using the measure developed by Khandelwal (2010) in figure E.2. The results are extremely similar. In particular, the share of the variation in prices explained by the bureaucrats and organizations remains constant as we increase the the degree of good homogeneity moving from right to left.

As we discuss in section 5.4, prices are the most important outcome in procurement, but not the only one, and so we also study the impact of bureaucrats and organizations on the spending quality measures described in section 3.2. We argue that these outcomes are endogenous to the bureaucrats and organizations in charge of procurement, and hence do not belong as controls in the variance decomposition. Nevertheless, in Appendix Table E.1 we re-estimate the variance decomposition including the spending quality outcomes as controls, and show that the results are essentially unchanged from our baseline specification in table 2 (for example, the standard deviation of the joint effect of the buyers goes from 0.499 down only to 0.484).

As discussed in section 5.2, bureaucrat- and organization- effects can only be estimated within sets of organizations connected by bureaucrats switching between them — connected sets. In our main analysis we pool the connected sets. As a robustness check, here we present results using only the largest connected set in the data. Table E.2 presents summary statistics of this largest connected set. The sample is broadly comparable to the main sample. Table E.3 shows the results of the variance decomposition in the largest connected set. The results are very similar to the main sample. The fixed effects, split-sample and shrinkage methods all attribute roughly the same share of the variation to the bureaucrats and organizations as in the full sample. The covariance shrinkage method attributes a bit less, 30%, slightly less than in the full sample. This gives us confidence that our results apply well beyond the lasrgest connected set.

Section 5.2 and appendix A describe the steps we took to build our analysis sample. Table E.4 shows the robustness of our estimates to the three main design choices. Column (1) replicates the findings in column (1) of table 2. Columns (2) and (3) use lower (45th percentile) and higher (55th percentile) thresholds of confidence to identify correctly classified items, respectively. Columns (4) and (5) trim fewer (top and bottom 2.5%) and more (top and bottom 10%) outlier observations for each good. Column (6) uses the Support Vector Machine classifier and column (7) uses the hierarchical classifier. All details are described in section A. As the table reveals, the results are remarkably stable across samples, reassuring us that our results are not driven by our sample building strategy.

FIGURE E.1: CORRELATION OF RESIDUALS WITH ESTIMATED BUREAUCRAT AND ORGANIZATION EFFECTS

PANEL A: PRICES IN LOGS (MAIN SPECIFICATION)

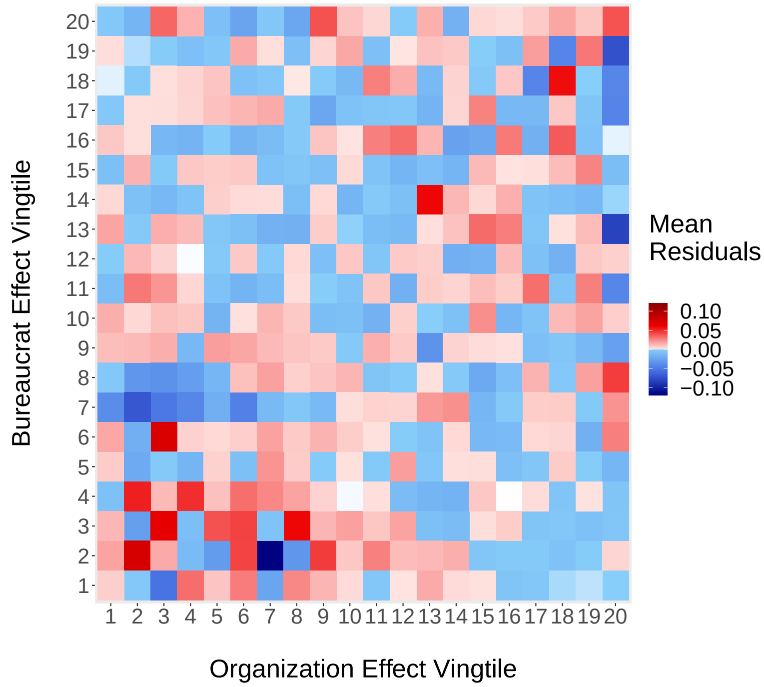

PANEL B: PRICES IN LEVELS (ILLUSTRATING MISSPECIFICATION)

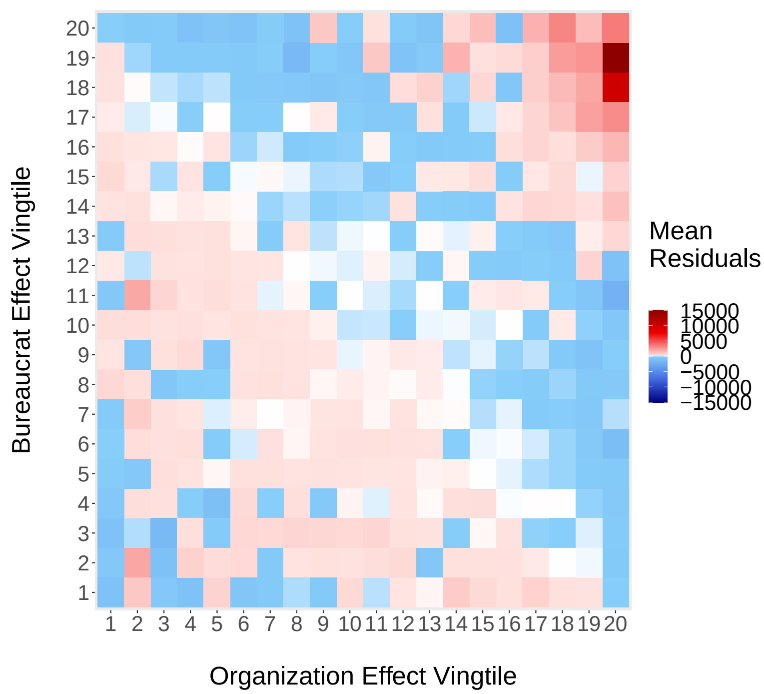

Notes: The figure presents heatmaps of averages of the residuals from the estimation of equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ — in logs (Panel A) and in levels (Panel B). The residuals are binned by vingtiles of the estimated bureaucrat effect $\hat{\alpha}_b$ and organization effect $\hat{\psi}_j$ within each connected set. The sample used is the Analysis Sample (All Products) summarized in Table 1.

FIGURE E.2: ROBUSTNESS TO USING SUBSAMPLES OF INCREASINGLY HETEROGENEOUS GOODS (KHANDELWAL (2010) MEASURE)

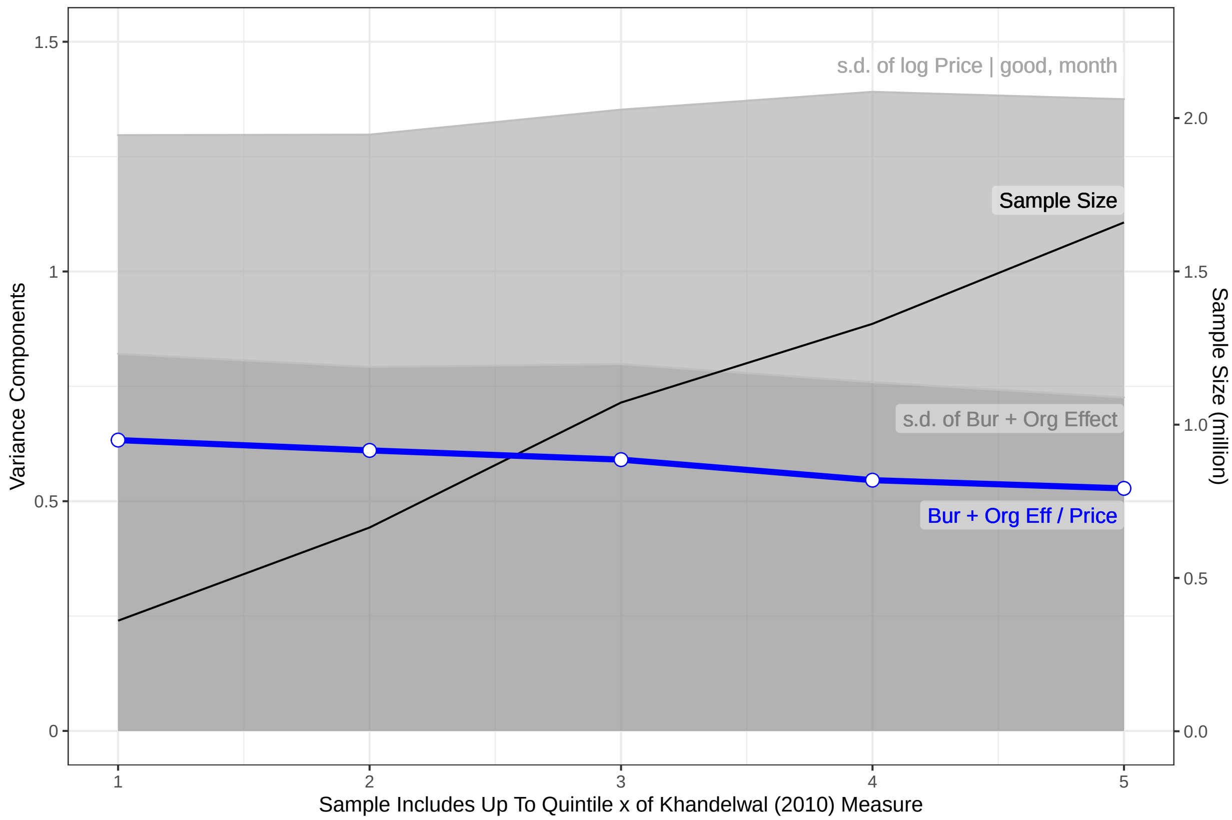

Notes: The figure shows the components of the variance of prices due to bureaucrats and organizations estimated by implementing the variance decomposition in equation (4) (see notes to Table 2 for details). The figure uses the sub-set of the sample that we can match to the scope-for-quality-differentiation ladder developed by Khandelwal (2010). Moving from right to left we remove quintiles of the data with the highest scope for quality differentiation, as shown by the black line, which indicates the sample size used. The dark shaded region is the variance of prices attributable to the bureaucrats and organizations. The dark and light shaded regions show the total variance of prices. The blue line shows the fraction of the overall variance attributable to bureaucrats and organization, highlighting that it remains roughly constant as we add more heterogeneous goods to the sample.

TABLE E.1: ROBUSTNESS OF VARIANCE DECOMPOSITION OF PRICES TO INCLUDING SPENDING QUALITY CONTROLS

	Fixed Effects (1)	(s.e.) (2)	Split Sample (3)	(s.e.) (4)	Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	1.192	(0.037)	1.257	(0.0349)	0.821	0.429
(2) s.d. of Organization Effects (across orgs)	1.118	(0.043)	1.185	(0.0491)	0.775	0.365
(3) s.d. of Bureaucrat Effects (across items)	0.780	(0.036)	0.828	(0.0447)	0.594	0.270
(4) s.d. of Organization Effects (across items)	0.914	(0.0456)	0.970	(0.0543)	0.702	0.329
(5) Bur-Org Effect Correlation (across items)	-0.718	(0.0153)	-0.518	(0.0419)	-0.664	0.301
(6) s.d. of Bur + Org Effects Within CS (across items)	0.648	(0.0183)	0.657	(0.0189)	0.540	0.484
(7) s.d. of log unit price	2.188		2.188		2.188	2.188
(8) s.d. of log unit price | good, month	1.280		1.280		1.280	1.280
(9) Adjusted R-squared	0.963		0.963		0.963	0.963
(10) Number of Bureaucrats	37,722		37,722		37,722	37,722
(11) Number of Organizations	44,560		44,560		44,560	44,560
(12) Number of Bureaucrat-Organization Pairs	248,898		248,898		248,898	248,898
(13) Number of Connected Sets	616		616		616	616
(14) Number of Observations	11,339,188		11,339,188		11,339,188	11,339,188

Notes: The table shows the components of the variance due to bureaucrats, organizations, and controls sets estimated by implementing the variance decomposition in equation (4) extended to include our spending quality measures as controls. The sample used is the All Products-Analysis Sample summarized in Table 1. Rows 1 & 2 show the s.d. of the bureaucrat, organization and connected set effects. Rows 3–6 show the components of the variance of prices across purchases, effectively weighting the estimates in rows 1 & 2 by the number of purchases they conduct. Column 1 uses the fixed effect estimates from equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates. Columns 2 and 4 show standard errors of the estimates in columns 1 and 3, respectively, estimated by bootstrapping 100 times. Column 5 uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2, and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2/(\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column 1, and analogously for the organization effects. Column 6 shows our preferred estimates, which form predictions of the bureaucrat and organization effects that minimize the expected sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates. Formally, the covariance shrinkage predictors solve $\min_{\boldsymbol{\Lambda}} \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}})'(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}})\right]$ where $\hat{\boldsymbol{\theta}}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 5.2.

TABLE E.2: SUMMARY STATISTICS - LARGEST CONNECTED SET

	All Products		Pharmaceuticals Subsample	
	(1) No Preferences Analysis Sample	(2) Largest Connected Set	(3) No Preferences Analysis Sample	(4) Largest Connected Set
(1) # of Bureaucrats	37,722	6,083	2,473	3,088
(2) # of Organizations	44,560	9,001	1,866	1,900
(3) # of Connected Sets	616	1	129	1
(4) # of Bureaucrats with >1 Org.	11,063	1,792	926	31
(5) # of Organizations with >1 Bur.	37,306	7,213	1,449	637
(6) Mean # of Bureaucrats per Org.	5.59	6.42	4.32	1.65
(7) Mean # of Organizations per Bur.	6.6	9.51	3.26	1.02
(8) # of Federal Organizations	1,583	166	26	478
(9) # of Regional Organizations	15,530	3,513	1,599	1,271
(10) # of Municipal Organizations	27,447	5,322	241	151
(11) # of Health Organizations	7,231	1,604	1,705	1,580
(12) # of Education Organizations	25,271	4,892	61	36
(13) # of Internal Affairs Organizations	668	98	3	102
(14) # of Agr/Environ Organizations	255	59	1	25
(15) # of Other Organizations	11,135	2,348	96	157
(16) # of Goods	14,875	12,048	3,861	3,713
(17) Mean # of Goods Per Bur.	72.5	83.8	42.5	22.8
(18) # of Regions	86	28	85	85
(19) Mean # of Regions per Bur.	1	1	1	1
(20) # of Auction Requests	1,199,363	248,999	42,874	19,818
(21) Mean # of Requests per Bur.	31.8	40.9	17.3	6.42
(22) Mean # of Applicants	3.6	3.63	3.03	2.85
(23) Mean # of Bidders	2.07	2.08	1.94	1.88
(24) Mean Reservation Price	25,140	18,675	0.062	0.17
(25) Quantity Mean	1,053	951	1,719	333
Median	25	30	45	35
SD	90,917	40,257	172,145	2,972
(26) Total Price Mean (bil. USD)	80.1	70.8	91.1	189
Median	4.32	3.72	6.7	5.69
SD	493	460	493	259
(27) Unit Price Mean (bil. USD)	61.3	48.8	25.4	11.6
Median	0.167	0.132	0.18	0.169
SD	23,015	2,076	265	138
(28) Mean # of Contract Renegotiations (log)	0.121	0.12	0.141	0.168
(29) Mean Size of Cost Over-run	-0.002	-0.002	-0.003	-0.004
(30) Mean Length of Delay in Days (log)	0.064	0.07	0.076	0.078
(31) Mean 1[End User Complained about Contract]	0.001	0.001	0	0.001
(32) Mean 1[Contract Cancelled]	0.012	0.013	0.016	0.015
(33) Mean 1[Product is of Substandard Quality]	0.005	0.003	0.058	0.112
(34) # of Observations	11,339,188	2,258,081	181,961	108,378
(35) Total Procurement Volume (bil. USD)	395	54.1	9.38	5.14

Notes: The table reports summary statistics for four samples. The All Products columns show statistics for purchases of all off-the-shelf goods, while the Pharmaceuticals Subsample columns restrict attention to purchases of medicines. Analysis Sample denotes all unpreferenced auctions in connected sets that fulfill three restrictions: singleton bureaucrat-organization, bureaucrat-good, and organization-good pairs are removed; each procurer (bureaucrats and organizations) implements a minimum of five purchases; and connected sets have at least three bureaucrats and organizations. Largest Connected Set is the largest connected set from the Analysis Sample (as measured by the number of organizations). Organizations working in Education include schools, universities, pre-schools, and youth organizations. Organizations working in Internal Affairs include police, emergency services, local administration, taxes, and transportation. Organizations working in Agriculture or the Environment include environmental protection funds, agricultural departments and nature promotion agencies. The Other category includes funds, monitoring agencies, and land cadasters, among many others. All sums are measured in billions of US dollars at an exchange rate of 43 rubles to 1 US dollar.

TABLE E.3: SHARE OF VARIANCE OF PROCUREMENT PRICES EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS: LARGEST CONNECTED SET

	Fixed Effects (1)	(s.e.) (2)	Split Sample (3)	(s.e.) (4)	Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	0.971	(0.0196)	1.001	(0.0241)	0.695	0.380
(2) s.d. of Organization Effects (across orgs)	1.031	(0.0386)	1.056	(0.0425)	0.753	0.342
(3) s.d. of Bureaucrat Effects (across items)	0.525	(0.0165)	0.540	(0.0222)	0.427	0.200
(4) s.d. of Organization Effects (across items)	0.692	(0.0254)	0.683	(0.0289)	0.602	0.232
(5) Bur-Org Effect Correlation (across items)	-0.624	(0.0177)	-0.432	(0.031)	-0.561	0.352
(6) s.d. of Bur + Org Effects Within CS (across items)	0.549	(0.00904)	0.538	(0.00966)	0.507	0.355
(7) s.d. of log unit price	2.165		2.165		2.165	2.165
(8) s.d. of log unit price | good, month	1.207		1.207		1.207	1.207
(9) Adjusted R-squared	0.964		0.964		0.964	0.964
(10) Number of Bureaucrats	6,083		6,083		6,083	6,083
(11) Number of Organizations	9,001		9,001		9,001	9,001
(12) Number of Bureaucrat-Organization Pairs	57,822		57,822		57,822	57,822
(13) Number of Connected Sets	1		1		1	1
(14) Number of Observations	2,258,081		2,258,081		2,258,081	2,258,081

Notes: The table shows the components of the variance due to bureaucrats, organizations, and connected sets estimated by implementing the variance decomposition in equation (4). The sample used is the All Products-Largest Connected Set Sample summarized in Table E.2. Rows 1 & 2 show the s.d. of the bureaucrat, organization and connected set effects. Rows 3–6 show the components of the variance of prices across purchases, effectively weighting the estimates in rows 1 & 2 by the number of purchases they conduct. Column 1 uses the fixed effect estimates from equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates. Columns 2 and 4 show standard errors of the estimates in columns 1 and 3, respectively, estimated by bootstrapping 100 times. Column 5 uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2, and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2/(\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column 1, and analogously for the organization effects. Column 6 shows our preferred estimates, which form predictions of the bureaucrat and organization effects that minimize the expected sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates. Formally, the covariance shrinkage predictors solve $\min_{\boldsymbol{\Lambda}} \mathbb{E}\left[(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}})'(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}})\right]$ where $\hat{\boldsymbol{\theta}}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 5.2.

TABLE E.4: VARIANCE DECOMPOSITION RESULTS: ROBUSTNESS TO SAMPLE DEFINITION

Machine Learning Method	LR	LR	LR	LR	LR	SVM	HM
Classification Confidence Threshold	50	45	55	50	50	50	50
Outlier Trimming	5	5	5	2.5	10	5	5
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) s.d. of Bureaucrat Effects (across burs)	1.199	1.154	1.123	1.336	0.857	1.108	1.104
(2) s.d. of Organization Effects (across orgs)	1.133	1.048	1.002	1.185	0.716	0.989	0.984
(3) s.d. of Bureaucrat Effects (across items)	0.795	0.722	0.686	0.834	0.525	0.674	0.680
(4) s.d. of Organization Effects (across items)	0.931	0.821	0.790	0.947	0.604	0.766	0.770
(5) Bur-Org Effect Correlation (across items)	-0.726	-0.682	-0.660	-0.688	-0.661	-0.663	-0.653
(6) s.d. of Bur + Org Effects Within CS (across items)	0.651	0.622	0.616	0.711	0.470	0.597	0.610
(7) s.d. of log unit price	2.188	2.214	2.197	2.417	1.854	2.194	2.188
(8) s.d. of log unit price | good, month	1.280	1.302	1.282	1.411	1.094	1.250	1.282
(9) Adjusted R-squared	0.963	0.963	0.964	0.958	0.970	0.965	0.963
(10) Number of Bureaucrats	37,722	38,154	37,722	40,892	34,393	37,893	37,563
(11) Number of Organizations	44,560	44,736	44,560	46,719	41,866	44,759	44,506
(12) Number of Bureaucrat-Organization Pairs	248,898	250,475	249,012	265,269	231,656	250,394	248,787
(13) Number of Connected Sets	616	614	616	619	604	618	618
(14) Number of Observations	11,339,188	11,364,608	11,341,098	12,081,256	10,012,706	11,365,756	11,343,316

The table shows the components of the variance due to bureaucrats, organizations, and connected sets estimated by implementing the variance decomposition in equation (4) in different samples. The decomposition uses the fixed effect estimates from equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Column (1) replicates the findings in column (1) of table 2. Columns (2) and (3) use lower (45th percentile) and higher (55th percentile) thresholds of confidence to identify correctly classified items, respectively. Columns (4) and (5) trim fewer (top and bottom 2.5%) and more (top and bottom 10%) outlier observations for each good. Column (6) uses the Support Vector Machine classifier and column (7) uses the hierarchical classifier. All details are described in Appendix A.

E.3 Crude Counterfactuals and Comparison to Existing Estimates of Individuals' and Organizations' Effects on Output

Our large estimates of the share of variation in performance attributable to bureaucrats and organizations have correspondingly dramatic implications for the scope of potential savings from improving the effectiveness of the bureaucracy. To illustrate the magnitude, we can consider simple counterfactual bureaucracies in which bureaucrats and/or organizations with low effectiveness are improved, for example through changes in recruiting, training of existing bureaucrats, or improved organizational management. Figure E.1 shows two such counterfactuals. Panel A shows the shift in the distribution of bureaucrat effects that would occur if the lowest quartile of bureaucrats were able to be improved to the 75th percentile. This would save the Russian government 4.5 percent of annual procurement expenses. In Panel B we consider moving all bureaucrats *and* organizations below 25th percentile-effectiveness to 75th percentile-effectiveness. The panel shows the distribution of pair (bureaucrat plus organization) effects that would result. The government would save 12.1 percent of procurement expenditures.[81] Annual procurement expenses are USD 86 billion, so this implies savings of USD 10 billion each year, or 0.7 percent of non-resource GDP (see Table H.2)—roughly one fifth, for example, of the total amount spent on health care in 2013 and 2014.[82]

How do our results compare to existing estimates of the extent to which individuals and organizations affect output in other settings? While we are not aware of comparable estimates of the causal effects of workers and organizations on output in a low or middle-income country government context, several studies are indirectly comparable. First, studying front-line service providers in rich countries, Chetty *et al.* (2014) find that increasing the performance of 5th percentile American grade 3–8 teachers to 50th percentile would increase the present value of their students' lifetime incomes by 2.76 percent, and Silver (2016) finds that improving the performance of American emergency room doctors by one standard deviation would decrease time-of-care by 11 percent. We find that the same (relative) improvement in performance among Russian procurement officers would lower prices paid by 32.6 and 30.5 percent respectively.[83] However, teachers and doctors may differ from procurement officers in the complexity of the job performed, motivations, and many other dimensions.

Second, in studies of workers in the private sector performing a simpler task, Mas & Moretti (2009) and Lacetera *et al.* (2016) find, respectively, that increasing performance by one standard deviation would decrease cashier processing times in a U.S. supermarket chain and increase the probability of cars being sold in U.S. used-car auctions by 11 and 4.3 percent, while in our case the improvement is 36.5 percent. Of course, in the public sector, output is less easily measured and monitored, and so we expect greater scope for differences between bureaucrats. Bertrand & Schoar

[81] Figure E.1 shows how these counterfactuals affect the distributions of effectiveness.

[82] Online Appendix E.3 compares these magnitudes to other studies of individuals' and organizations' effects on output in other settings.

[83] We perform these calculations separately in each connected set and report the average, weighting by the number of items.

(2003) find that CEOs in the top quartile of performance achieve a return-on-assets that is about 200 percent higher than CEOs in the bottom quartile. In our context, bureaucrats in the bottom quartile save 54.1 percent relative to the top quartile due solely to the bureaucrat effects.

FIGURE E.1: CRUDE COUNTERFACTUALS

Panel A: Moving Least Effective 25% of Bureaucrats to 75th Percentile Effectiveness

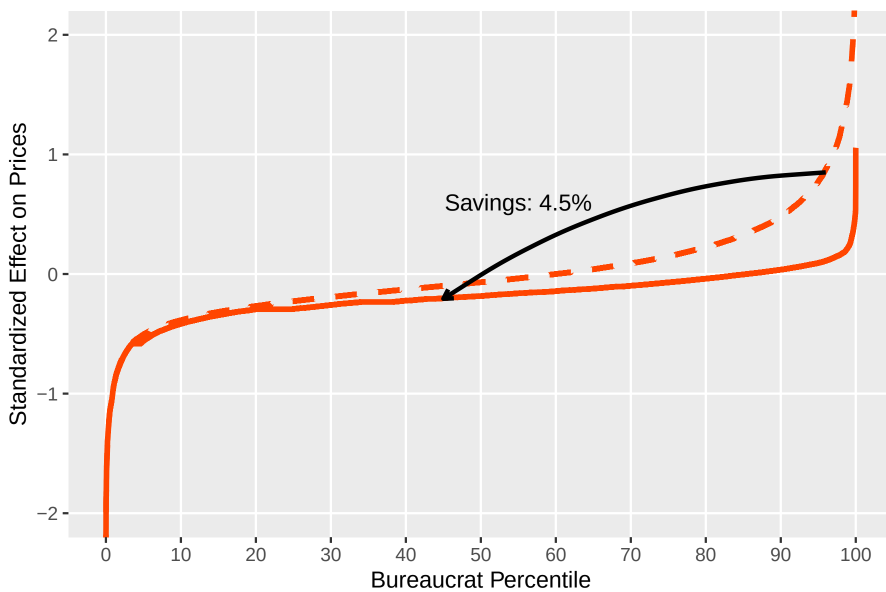

Panel B: Moving Least Effective 25% of Bureaucrats and Organizations to 75th Percentile Effectiveness

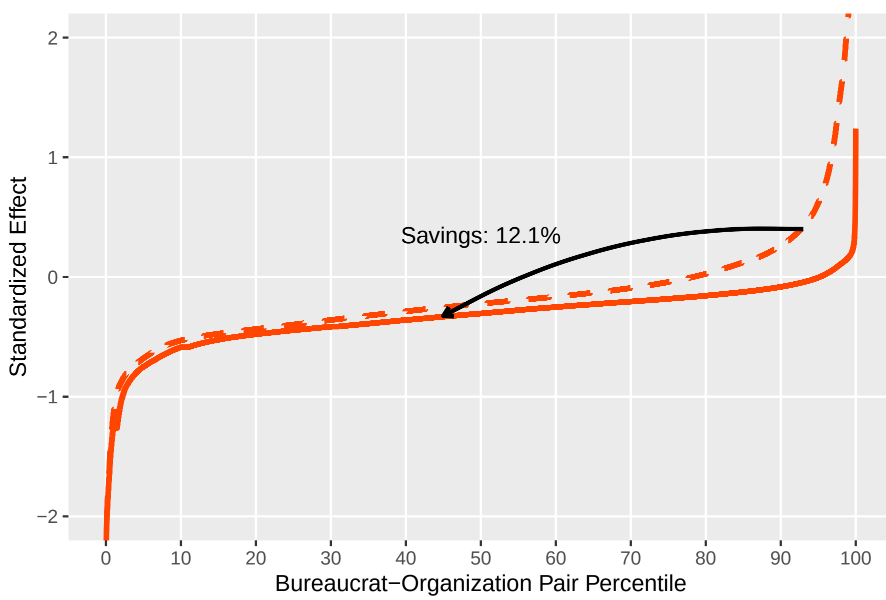

Notes: The figure shows the impact of two counterfactual scenarios on the distribution of our estimated price effects. Panel A considers moving all bureaucrats above the 75th percentile of their connected set's distribution of covariance shrunken price effects down to their connected set's 25th percentile. The dashed line shows the distribution of our covariance shrunken estimates of the bureaucrat effects, while the solid line shows the distribution that would result from implementing the counterfactual. Panel B considers moving both all bureaucrats and all organizations above the 75th percentile of their connected set's distribution of covariance shrunken price effects down to their connected set's 25th percentile. The dashed line shows the distribution of bureaucrat-organization pair effects we estimate, while the solid line shows the distribution that would occur in the counterfactual scenario. Overlaid on both panels are the implied aggregate savings.

F Additional Results on What Effective Bureaucracies do Differently

This appendix presents additional results on the variance decomposition discussed in section 5.5. In section 5.5 we exploit the richness of our data to analyze the correlates of bureaucratic and organizational effectiveness. To avoid overfitting and for the sake of parsimony, we use a LASSO procedure to first select 30 predictor variables.[84] We then regress each purchase's covariance-shrunk bureaucrat/organization effect on these variables, the purchase's organization effect, and the controls in (3). In Figures 4, 5, and F.1–F.6, the left panels show regression coefficients from a series of bivariate regressions of the bureaucrat/organization price/spending quality effect on each of the selected observables. The right panels show the LASSO coefficients (as crrosses) and the coefficients from the multivariate regression of the procurer effects on all of the selected variables (as circles). To facilitate comparison, all variables are standardized to have unit standard deviation. The coefficients can thus be interpreted as the association between a one-standard deviation change in the measure of procurer behavior and the causal impact of the procurer.

In the main paper, we present results on correlates of bureaucrats' price (Figure 4) and spending quality (Figure 5) effects. Figures (F.1) and (F.2) present the analogs for organizations. For parsimony we selected 30 predictor variables, but Figures (F.3) – (F.6) extend Figures (4), (5), (F.1) and (F.2) to pick 60 variables instead of 30. To account for small firms not being covered by the *Ruslana* data and the strong correlation between some of our variables, we also use an elastic net regularizer (a weighted average of LASSO and Ridge regression). Figures F.7 and F.8 show that the results are not sensitive to placing more weight on the Ridge regression. Finally, Table F.1 summarizes the data used in this exercise.

[84] The procedure selects the smallest model with *at least* 30 variables so the actual number varies slightly from figure to figure. Table F.1 shows pairwise coefficients from regressing price-effectiveness on each of the 160 potential explanatory variables we start out with. Tables F.3 and F.4 instead show results from using the LASSO procedure to select 60 instead of 30 predictors. The patterns in the findings are very similar to those described below.

FIGURE F.1: CORRELATES OF ORGANIZATION EFFECTIVENESS (PRICE, 30 VARIABLES)

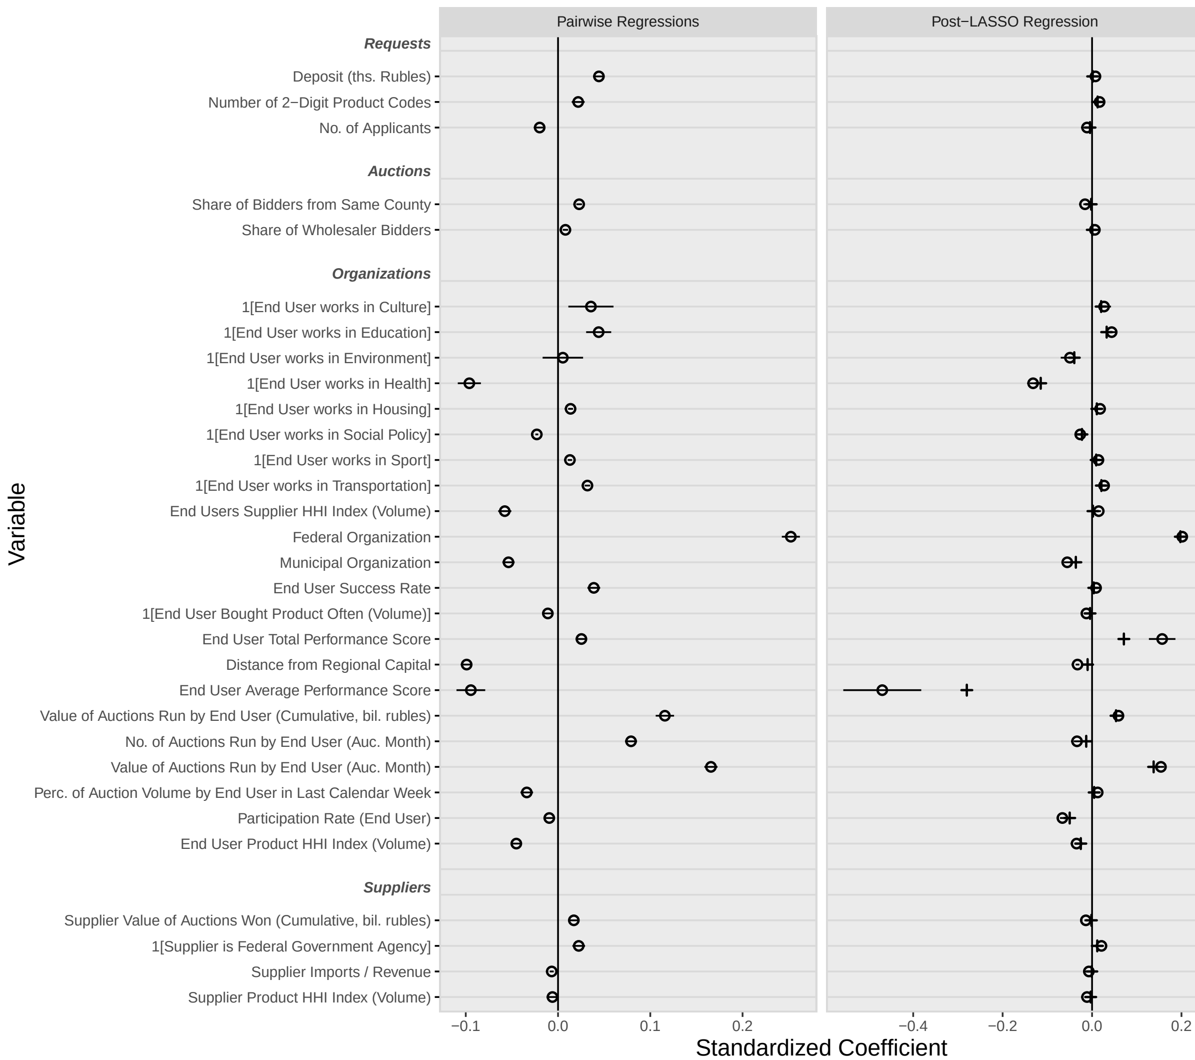

Notes: The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (3): $p_i = \mathbf{X}_i \boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ with prices as the outcome on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 30 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients from the multivariate regression of the effects on all of the selected variables. All variables are standardized to have unit standard deviation.

FIGURE F.2: CORRELATES OF ORGANIZATION EFFECTIVENESS (QUALITY, 30 VARIABLES)

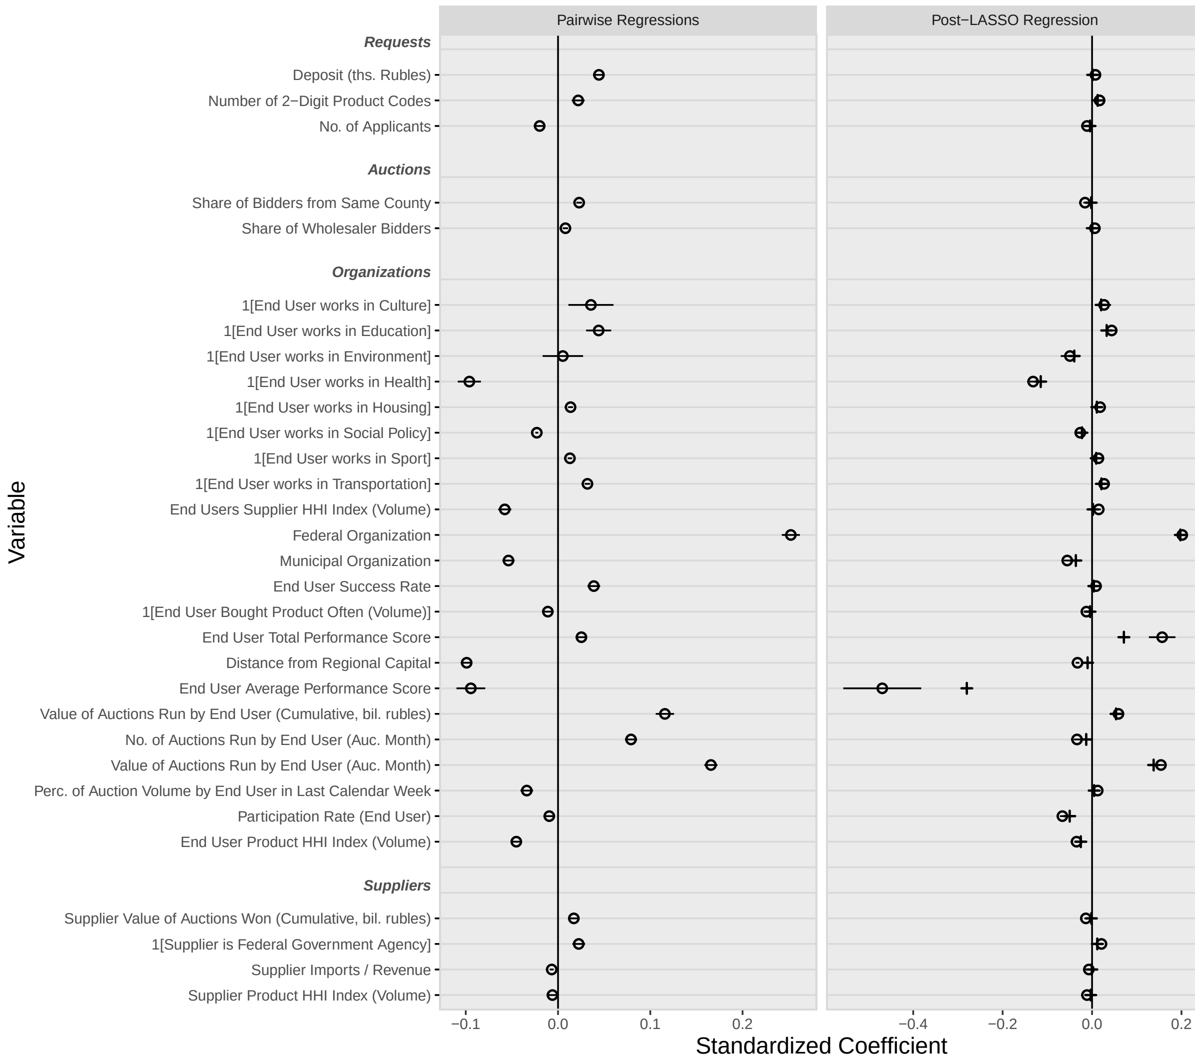

Notes: The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (3): $q_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ with spending quality as the outcome on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 30 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients from the multivariate regression of the effects on all of the selected variables. All variables are standardized to have unit standard deviation.

FIGURE F.3: CORRELATES OF BUREAUCRAT EFFECTIVENESS (PRICE, 60 VARIABLES)

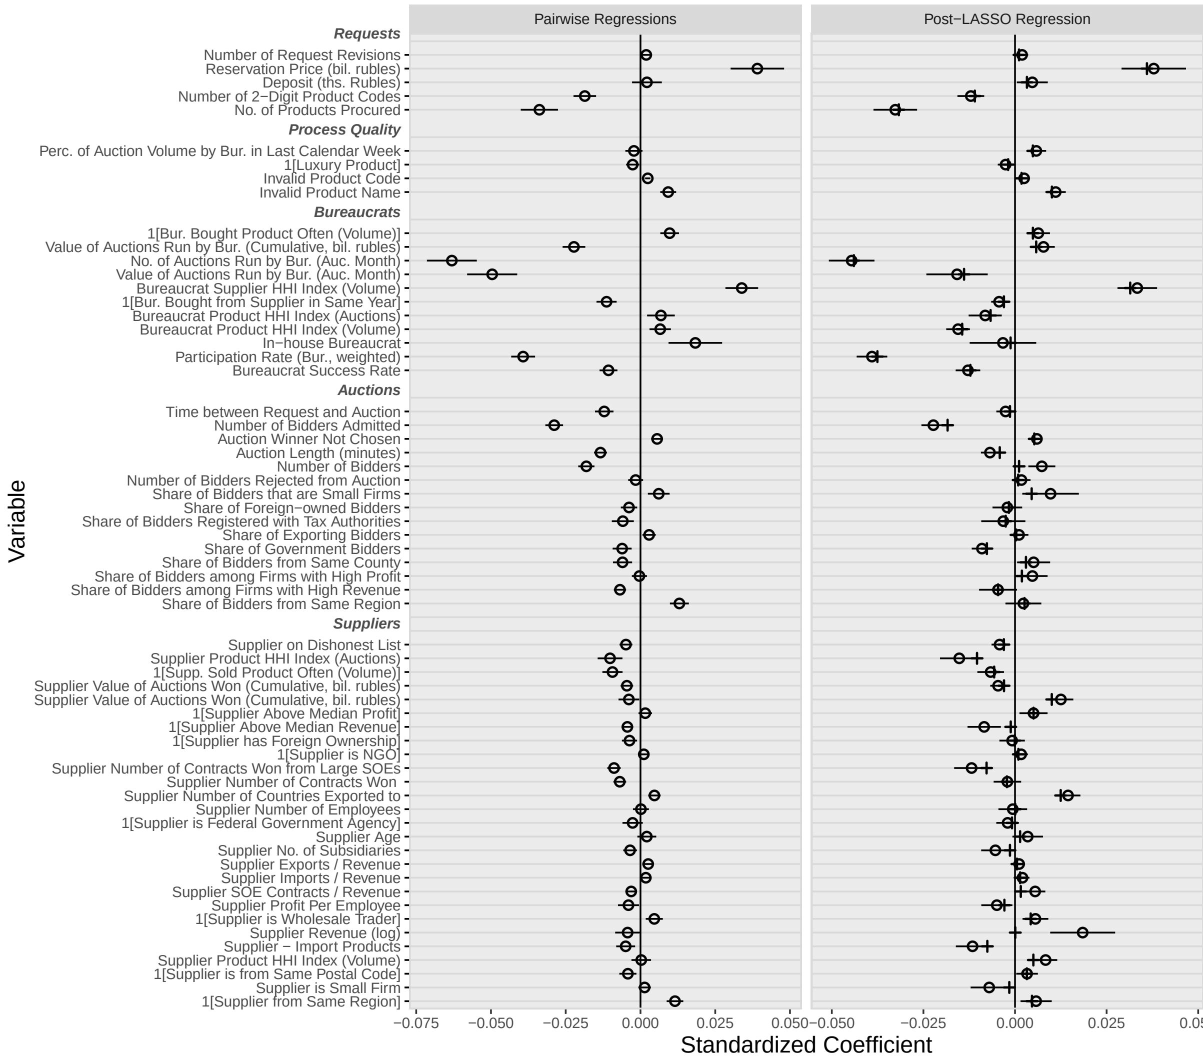

Notes: The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ with price as the outcome on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 60 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients from the multivariate regression of the effects on all of the selected variables. All variables are standardized to have unit standard deviation.

FIGURE F.4: CORRELATES OF BUREAUCRAT EFFECTIVENESS (QUALITY, 60 VARIABLES)

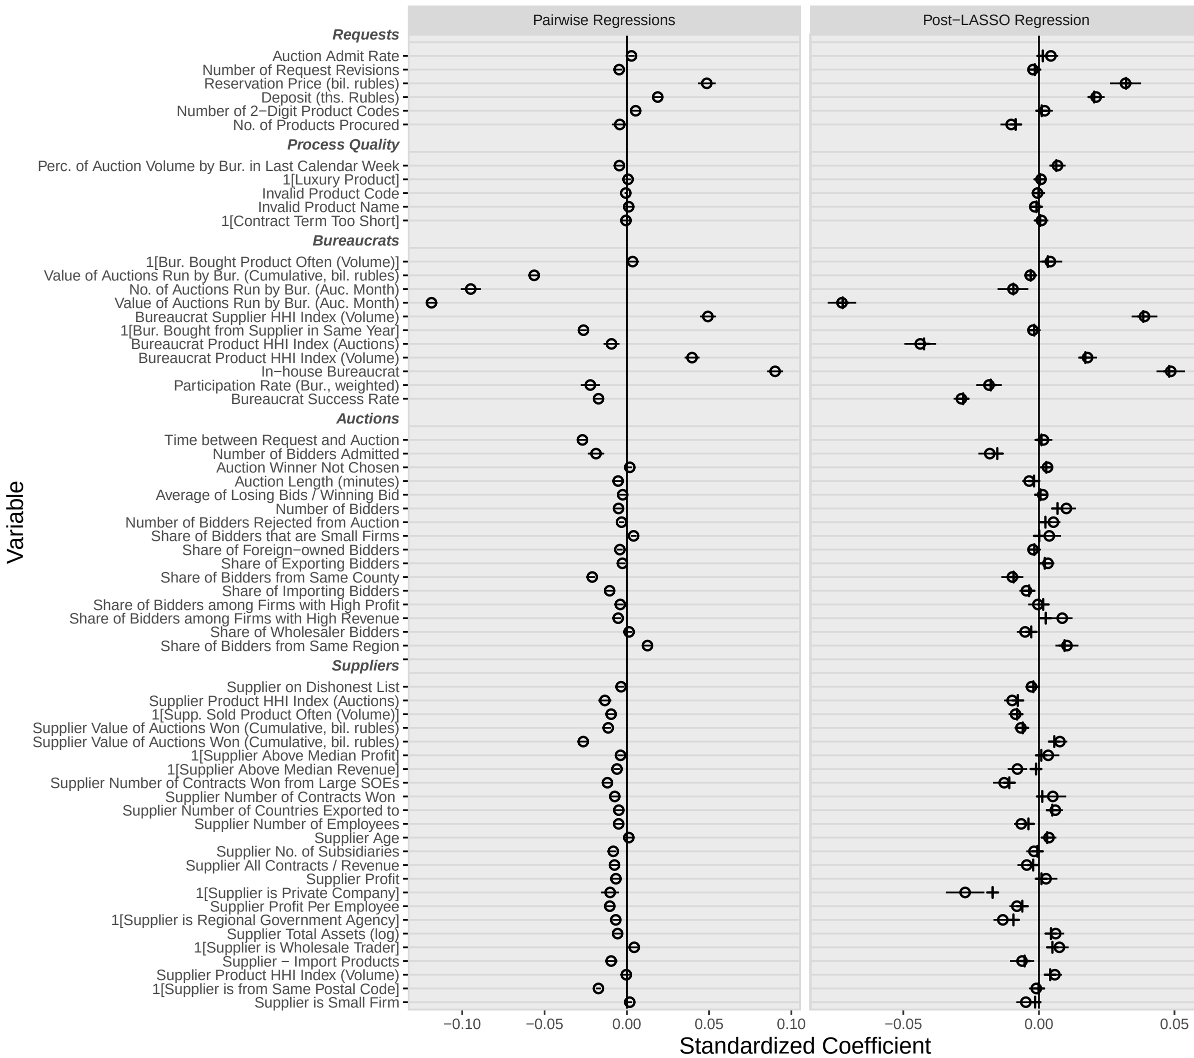

Notes: The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (3): $q_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ with spending quality as the outcome on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 60 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients from the multivariate regression of the effects on all of the selected variables. All variables are standardized to have unit standard deviation.

FIGURE F.5: CORRELATES OF ORGANIZATION EFFECTIVENESS (PRICE, 60 VARIABLES)

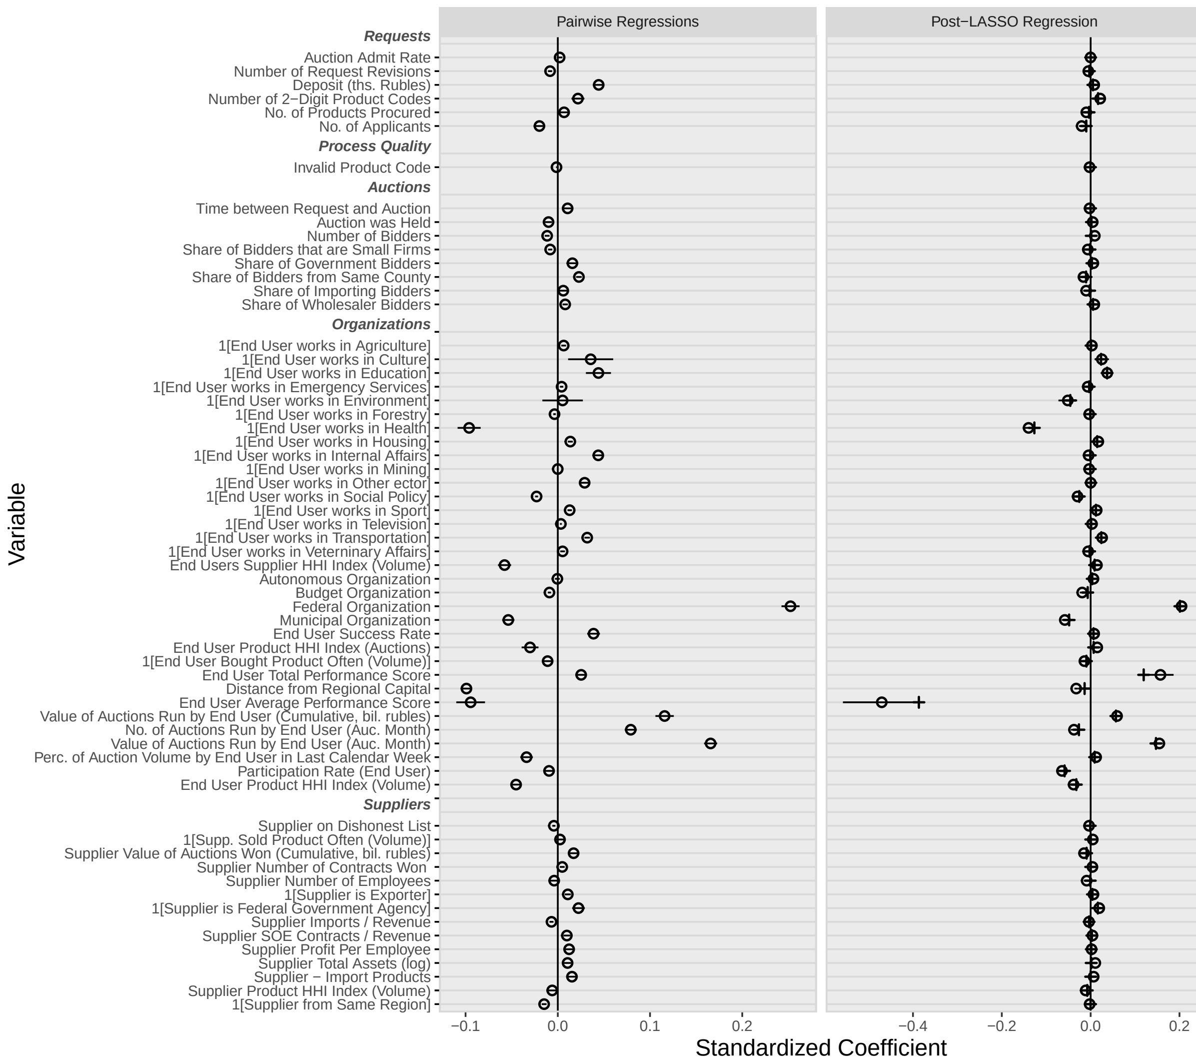

Notes: The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (3): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ with price as the outcome on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 30 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients from the multivariate regression of the effects on all of the selected variables. All variables are standardized to have unit standard deviation.

FIGURE F.6: CORRELATES OF ORGANIZATION EFFECTIVENESS (QUALITY, 60 VARIABLES)

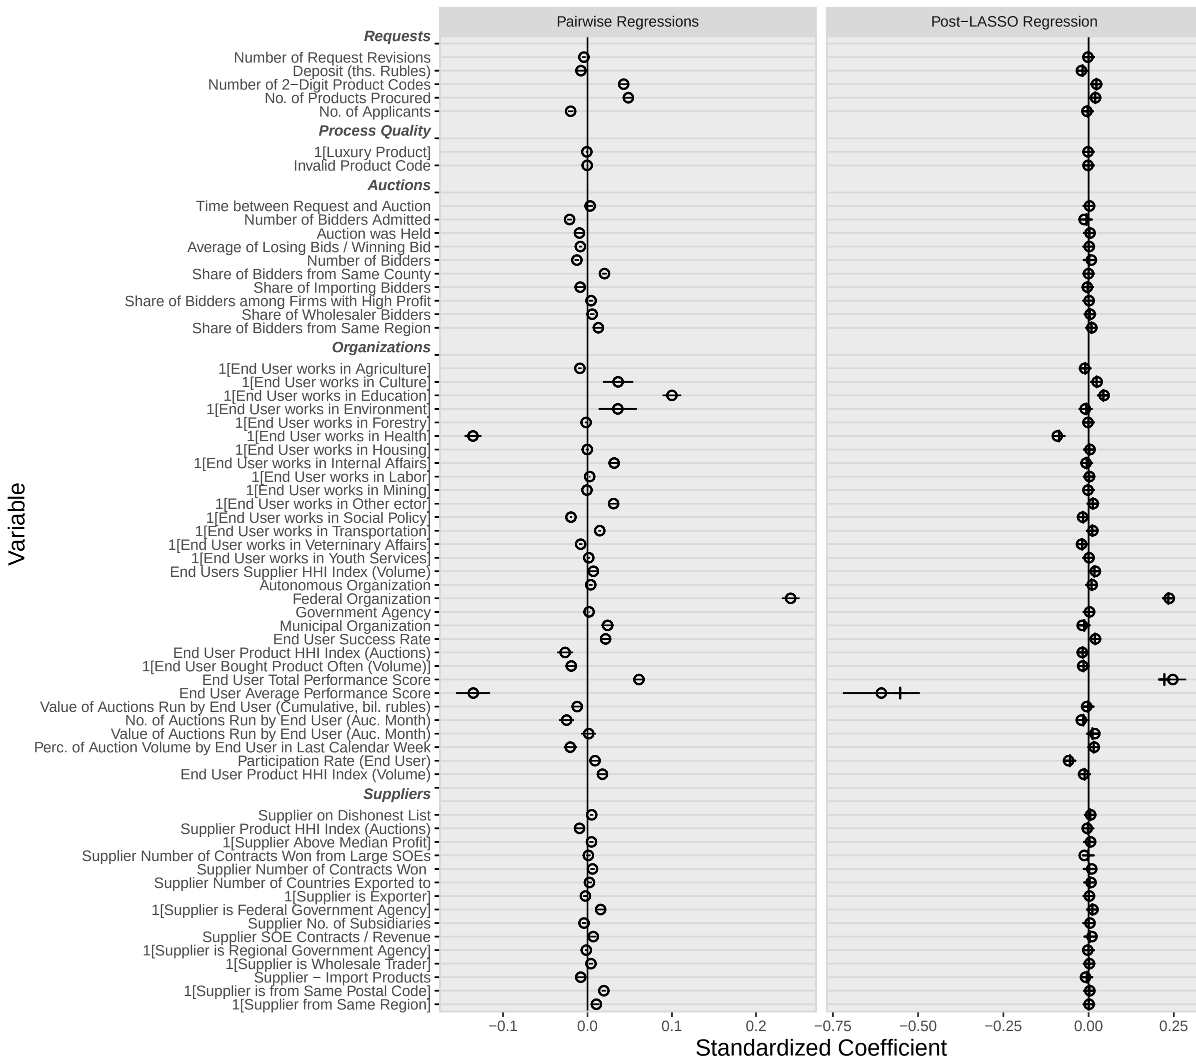

Notes: The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (3): $q_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ with spending quality as the outcome on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 30 predictor variables and regress each purchase's covariance-shrunk bureaucrat effect on these variables, the purchase's organization effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect on each of the selected observables. The right panels show the coefficients from the multivariate regression of the effects on all of the selected variables. All variables are standardized to have unit standard deviation.

FIGURE F.7: CORRELATES OF BUREAUCRAT EFFECTIVENESS (PRICE): ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS

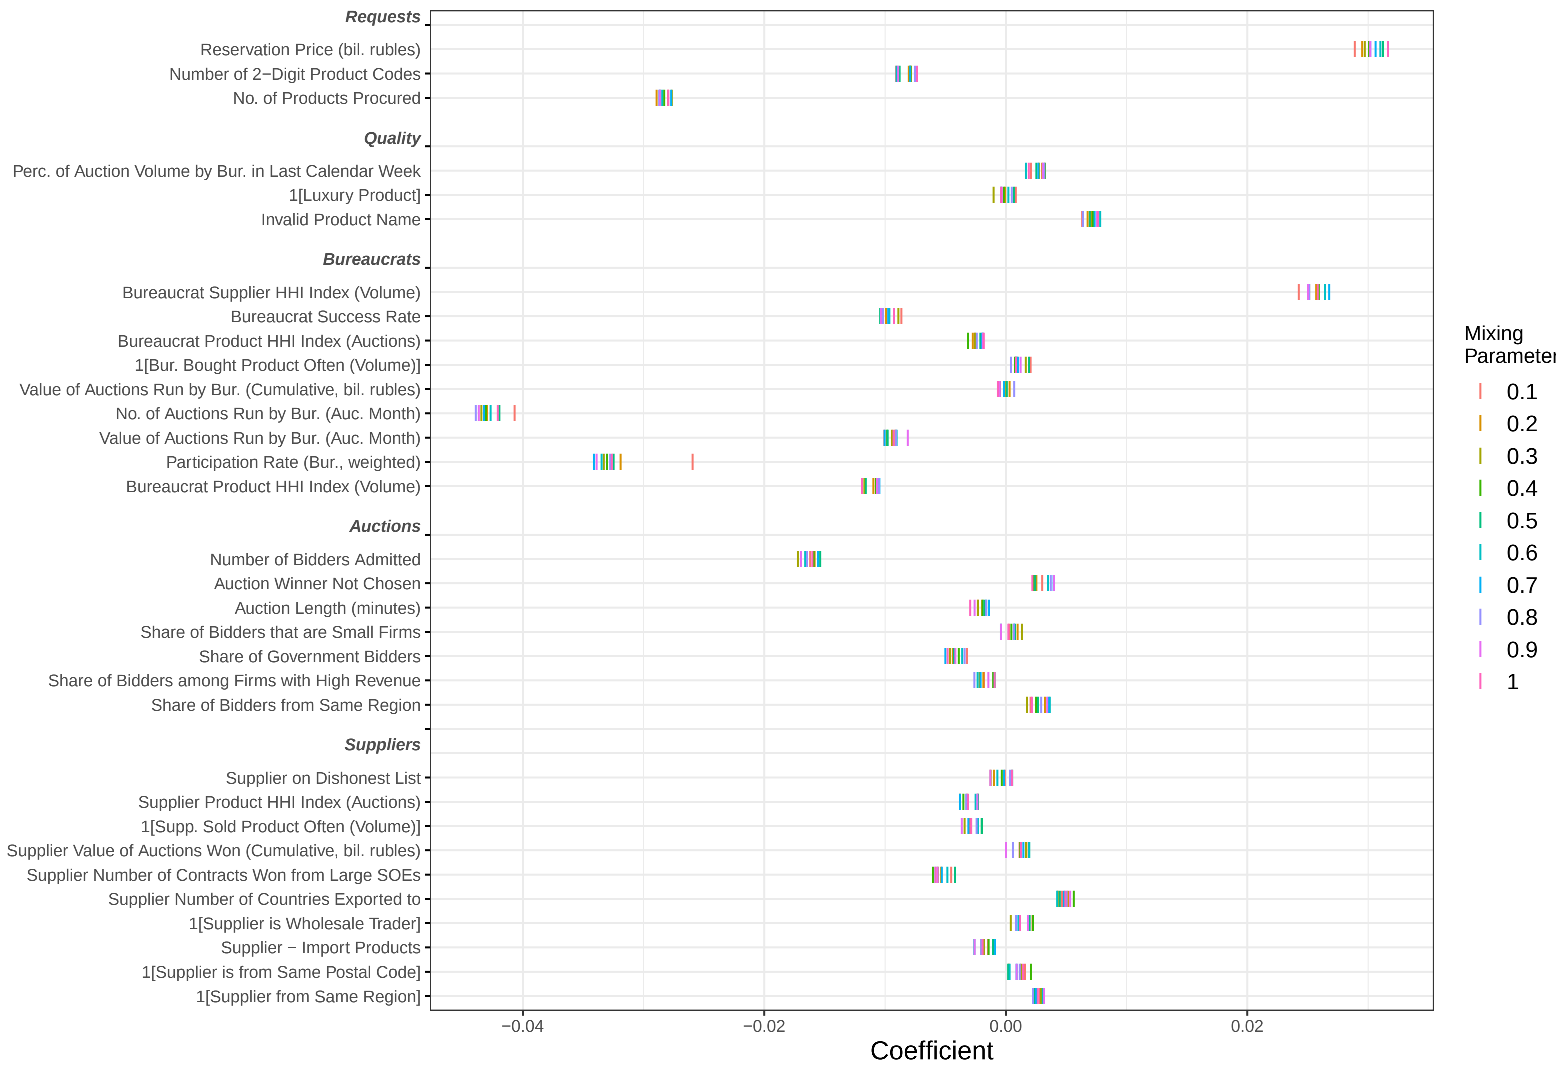

Notes: The figure shows the coefficients from the elastic net regularization procedure on the estimated bureaucrat effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. A mixing parameter of 1 represents LASSO, our baseline model. The variables shown are from the base model shown in Table 4 where the values of the regularization penalty lambda λ are chosen to return 30 predictor variables.

Figure F.8: Correlates of Bureaucrat Effectiveness (Quality): Elastic Net Regularization Coefficients Across Different Mixing Parameters

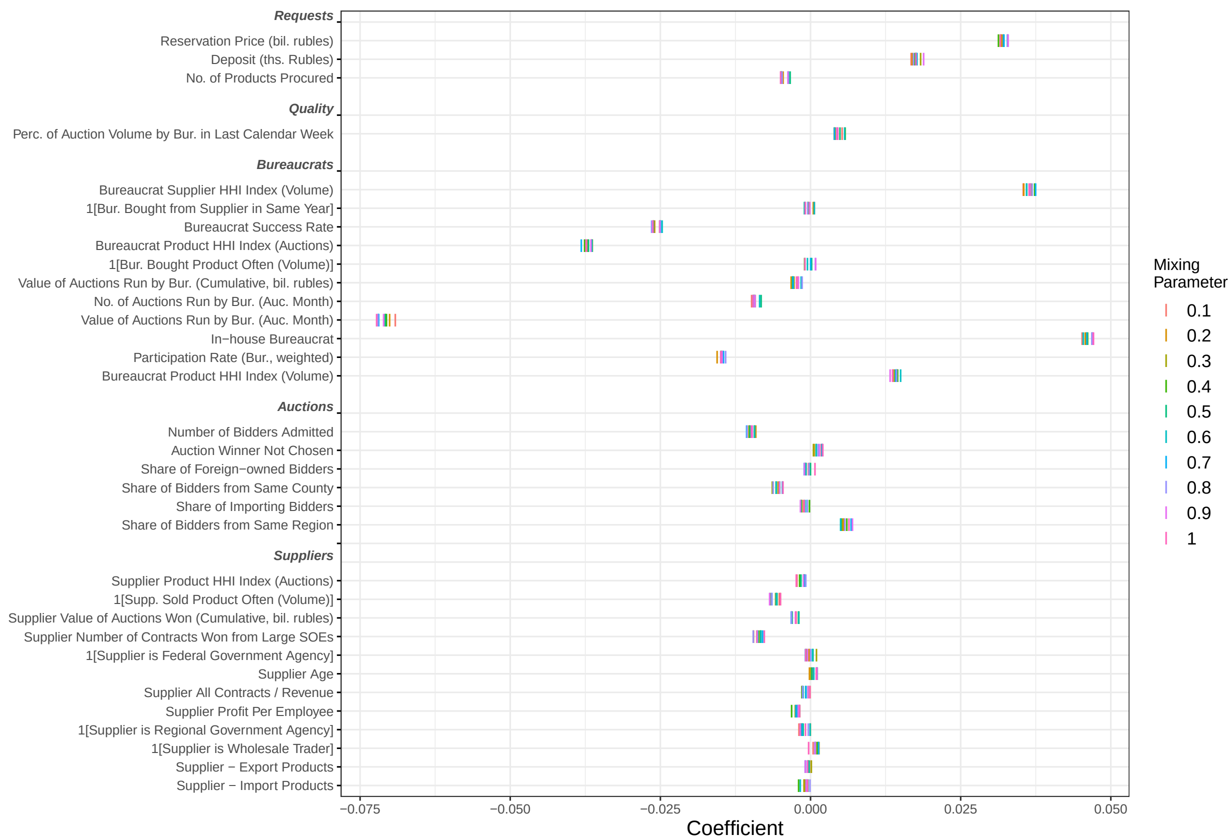

Notes: The figure shows the coefficients from the elastic net regularization procedure on the estimated bureaucrat effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. A mixing parameter of 1 represents LASSO, our baseline model. The variables shown are from the base model shown in Table 5 where the values of the regularization penalty lambda λ are chosen to return 30 predictor variables.

Table F.1: Correlatives of Bureaucrat and Organization Effectiveness: Variable Descriptions

Auctions	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description
Auction Length (minutes)	-0.01338 (0.00111)	-0.00767 (0.00108)	0.05	-0.52	3.21	Length of the auction in minutes
Auction Winner Not Chosen	0.00557 (0.0007)	0.00203 (0.00108)	0.02	-0.19	5.33	Indicator if the winner of the auction was ultimately not the suppler listed on the contract
Auction was Held	-0.01492 (0.0013)	-0.01003 (0.00181)	0	-1.26	0.79	Indicator if the auction was held (i.e. more than one supplier was admitted to the auction)
Average of Losing Bids / Winning Bid	-0.01001 (0.00105)	-0.00677 (0.00099)	0.01	-0.32	26.49	Ratio of the average of all losing bids over the final winning bid
Number of Bidders	-0.01811 (0.00138)	-0.01154 (0.00132)	0.02	-0.77	13.38	Number of bidders that entered bids
Number of Bidders Admitted	-0.02885 (0.00151)	-0.01959 (0.00182)	0.04	-0.73	29.3	Number of bidders admitted to participate in the auctino
Number of Bidders Rejected from Auction	-0.00161 (0.00127)	-0.00394 (0.00135)	0.04	-0.35	50.11	Number of bidders who were not allowed to participate in the auction
Share of Bidders Registered with Tax Authorities	-0.0059 (0.00188)	0.00649 (0.00105)	-0.07	-2.59	0.55	Share of bidders that participated in the auction that were registered with federal tax authorities
Share of Bidders among Firms with High Profit	-0.00036 (0.00129)	0.00485 (0.00148)	-0.1	-1.6	0.89	Share of bidders that participated in the auction that had above-median profits (relative to full sample of suppliers)
Share of Bidders among Firms with High Revenue	-0.00687 (0.00104)	0.00819 (0.00147)	-0.11	-1.76	0.79	Share of bidders that participated in the auction that had above-median revenue (relative to full sample of suppliers)
Share of Bidders from Same County	-0.006 (0.00163)	0.02292 (0.00154)	-0.02	-0.63	1.95	Share of bidders that participated in the auction that were located in the same county as the End User
Share of Bidders from Same Region	0.01301 (0.00162)	-0.01778 (0.00156)	0.07	-1.53	0.85	Share of bidders that participated in the auction that were located in the same region as the End User
Share of Bidders that are Small Firms	0.00611 (0.00185)	-0.00825 (0.00115)	0.1	-0.41	3.52	Share of bidders that participated in the auction that were registered as small firms
Share of Exporting Bidders	0.00289 (0.00112)	0.00958 (0.00186)	-0.07	-0.2	6.87	Share of bidders that participated in the auction that had exporting activities
Share of Foreign-owned Bidders	-0.00383	0.00351	-0.03	-0.16	9.19	Share of bidders that participated in the auction that were foreign-owned

	(0.0014)	(0.00125)				
Share of Government Bidders	-0.00613 (0.00162)	0.01585 (0.00326)	-0.06	-0.17	7.14	Share of bidders that participated in the auction owned by federal, regional, or municipal governments
Share of Importing Bidders	-0.00298 (0.00167)	0.00595 (0.0021)	-0.08	-0.48	2.71	Share of bidders that participated in the auction that had importing activities
Share of Wholesaler Bidders	0.00171 (0.00126)	0.00804 (0.00145)	0.03	-0.48	2.93	Share of bidders that participated in the auction that operated primarily as wholesale traders
Time between Request and Auction	-0.01212 (0.00158)	0.01067 (0.00286)	-0.04	-2.48	2.19	Number of days elapsed between the day the request was posted and the day the auction was held
Bureaucrats	**PwCorr - BurFE**	**PwCorr - OrgFE**	**Mean**	**Min**	**Max**	**Description**
1[Bur. Bought Product Often (Volume)]	0.00975 (0.00159)		-0.11	-0.51	1.97	Indicator if the main product was also the most common product purchased overall by the Bureaucrat (volume)
1[Bur. Sold to Supplier in Same Year]	-0.01134 (0.00172)		-0.1	-1.22	0.82	Indicator if Supplier won an auction in the previous calendar year with the same bureaucrat
Bureaucrat Product HHI Index (Auctions)	0.00684 (0.00236)		-0.11	-1.34	3.56	HHI measuring the distribution of auctions (count) by each bureaucrat across two-digit product types
Bureaucrat Product HHI Index (Volume)	0.0066 (0.00182)		0.02	-1.59	3.33	HHI measuring total sales volume of all auctions by each bureaucrat across two-digit product types
Bureaucrat Success Rate	-0.0107 (0.00152)		-0.02	-8.41	1.72	Percentage of requests administered by the Bureaucrat that led to a successful contract
Bureaucrat Supplier HHI Index (Volume)	0.03387 (0.00279)		0.02	-0.82	5.37	HHI measuring total volume of all auctions won by supplier per bureaucrat across two-digit product types
In-house Bureaucrat	0.01834 (0.00457)		0.02	-0.81	1.23	Indicator if the Bureaucrat worked directly at the End User
No. of Auctions Run by Bur. (Auc. Month)	-0.06308 (0.00426)		-0.11	-0.78	3.03	Number of auctions the Bureaucrat was running simultaneously in the same month as the auction
Participation Rate (Bur.)	-0.03833 (0.00205)		0.07	-0.99	52.59	Fraction of the relevant pool of suppliers that Bureaucrat is able to attract to their auction
Participation Rate (Bur., weighted)	-0.03924 (0.00204)		0.07	-1	53.9	Fraction of relevant pool of suppliers that Bureaucrat is able to attract to their auction, weighted by auction volume
Value of Auctions Run by Bur. (Auc. Month)	-0.0496 (0.00427)		-0.11	-5.33	3.77	Total sales volume of the auctions the Bureaucrat was running simultaneously in the same month as the auction
Value of Auctions Run by Bur. (Cumulative, bil. rubles)	-0.02225		-0.1	-3.76	1.63	Total sales volume of the auctions the Bureaucrat had run cumulatively to the date of the auction

	(0.00193)					
End Users	**PwCorr - BurFE**	**PwCorr - OrgFE**	**Mean**	**Min**	**Max**	**Description**
1[End User Bought Product Often (Volume)]		-0.01103 (0.00226)	-0.11	-0.56	1.77	Indicator if the main product was also the most common product purchased overall by the End User (volume)
1[End User works in Agriculture]		0.00639 (0.0013)	0.01	-0.03	30.68	End User works in the agricultural sector
1[End User works in Culture]		0.03567 (0.01248)	0.03	-0.09	11.4	End User works on cultural affairs
1[End User works in Education]		0.04405 (0.00693)	0.15	-0.44	2.25	End User works in education
1[End User works in Emergency Services]		0.00413 (0.00091)	0.02	-0.07	14.06	End User works in emergency services
1[End User works in Environment]		0.00522 (0.01119)	-0.01	-0.1	10.2	End User works in the environmental sector
1[End User works in Forestry]		-0.00355 (0.00069)	0.01	-0.02	41.5	End User works in the forestry sector
1[End User works in Health]		-0.09612 (0.00639)	-0.21	-1.44	0.69	End User works in the health care sector
1[End User works in Housing]		0.01359 (0.00135)	0.02	-0.07	14.77	End User works in the housing sector
1[End User works in Internal Affairs]		0.0439 (0.00253)	0.03	-0.11	9.09	End User works in internal affairs (police, justice, etc.)
1[End User works in Labor]		-0.0034 (0.00051)	0.01	-0.03	33.44	End User works in the labor sector (re-training, unemployment assistance, etc.)
1[End User works in Mining]		-0.00005 (0.00012)	0	-0.01	76.8	End User works in the mining sector
1[End User works in Natural Resources]		-0.00002 (0.00036)	0	-0.01	182.04	End User works in the natural resources sector
1[End User works in News]		0.00066 (0.00067)	0	-0.01	186.57	End User works in news and journalism
1[End User works in Other ector]		0.02911 (0.00219)	0.08	-0.28	3.63	End User works in other sector
1[End User works in Social Policy]		-0.02299 (0.00079)	0.05	-0.19	5.18	End User works on social policy (welfare, pensions, etc.)

1[End User works in Sport]		0.01289 (0.00121)	0.02	-0.05	19.03	End User works in the sport and recreational sector
1[End User works in Television]		0.00336 (0.00045)	0.01	-0.02	54.13	End User works in television and mass communications
1[End User works in Transportation]		0.03189 (0.00151)	0.02	-0.07	13.5	End User works in the transportation sector
1[End User works in Veterninary Affairs]		0.00521 (0.00129)	0	-0.04	26.76	End User works in veterinary affairs
1[End User works in Youth Services]		0.00127 (0.00073)	0.02	-0.05	20.38	End User works in youth services
1[**End Users** Sold to Supplier in Same Year]		0.02846 (0.00171)	-0.13	-1.23	0.81	Indicator if Supplier won an auction in the previous calendar year with the same End User
Autonomous Organization		-0.00046 (0.00132)	0.01	-0.18	5.58	End User is a non-commercial organization created by the government that enjoys more financial autonomy
Budget Organization		-0.00905 (0.00086)	0.01	-0.11	9.41	Non-commercial organization with less financial autonomy and stricter budget control from government owner
Distance from Regional Capital		-0.09904 (0.00276)	0.04	-1	1.45	Distance between the End User and the capital of the region where it is located (log kilometers)
End User Average Performance Score		-0.09447 (0.00795)	-0.09	-1.2	2.38	Average performance score across categories for the End User from evaluations by the Federal Treasury
End User Product HHI Index (Auctions)		-0.03009 (0.00466)	-0.07	-1.6	4.33	HHI measuring the distribution of auctions (count) by each End User across two-digit product types
End User Product HHI Index (Volume)		-0.04521 (0.00221)	0.1	-1.47	3.41	HHI measuring total sales volume of all auctions by each end user across two-digit product types
End User Success Rate		0.03877 (0.00352)	-0.01	-7.79	1.66	Percentage of requests administered for the End User that led to a successful contract
End User Total Performance Score		0.02535 (0.00261)	-0.08	-1.04	1.81	Total performance score for the End User from independent surveys and evaluations by the Federal Treasury
End Users Supplier HHI Index (Volume)		-0.05773 (0.00363)	0.04	-0.85	5.68	HHI measuring total volume of auctions won by supplier per End User across two-digit product types
Federal Organization		0.25232 (0.005)	0.04	-0.35	2.87	End User receives funds from the federal government and operates on the federal level
Government Agency		0.00054 (0.00044)	0	-0.03	31.44	End User is classified as a separate government agency, operating more independent of government oversight

Municipal Organization		-0.05369 (0.00338)	0.11	-0.51	1.94	End User receives funds from the municipal government and operates on the municipal level
No. of Auctions Run by End User (Auc. Month)		0.07912 (0.00334)	-0.19	-0.89	2.73	Number of auctions the End User was running simultaneously in the same month as the auction
Other Government Body		0.00389 (0.00121)	-0.02	-4.67	0.21	End User has a much less common legal classification, such as a natural monopoly, audit agency, etc.
Participation Rate (End User)		-0.00948 (0.00305)	0.09	-1.03	38.91	Fraction of the relevant pool of suppliers that End User is able to attract to their auction
Participation Rate (End User, weighted)		-0.00841 (0.00301)	0.09	-1.03	36.42	Fraction of relevant pool of suppliers that End User is able to attract to their auction, weighted by auction volume
Perc. of Auction Volume by End User in Last Calendar Week		-0.03387 (0.00349)	-0.01	-0.88	27.02	Percentage of all auctions (by volume) that End User ran in the last calendar week of the year
Regional Organization		-0.10366 (0.00585)	-0.13	-1.47	0.68	End User receives funds from the regional government and operates on the regional level
Value of Auctions Run by End User (Auc. Month)		0.16571 (0.00363)	-0.14	-6.28	4.17	Total sales volume of the auctions the End User was running simultaneously in the same month as the auction
Value of Auctions Run by End User (Cumulative, bil. rubles)		0.11581 (0.00507)	-0.14	-5.05	2.53	Total sales volume of the auctions the End User had run cumulatively to the date of the auction
Quality	**PwCorr - BurFE**	**PwCorr - OrgFE**	**Mean**	**Min**	**Max**	**Description**
1[Contract Term Too Short]	0.00106 (0.00113)	0.00238 (0.00101)	-0.01	-0.03	30.47	Indicator if amount of time to execute the contract is too short
1[Luxury Product]	-0.00261 (0.00112)	0.00131 (0.00063)	0	-0.01	103.24	Product purchased is considered to be luxury, per data from ClearSpending
Invalid Product Code	0.00244 (0.00045)	-0.0014 (0.00025)	0	-0.01	173.16	Request had an invalid product code per analysis by ClearSpending.Ru
Invalid Product Name	0.00927 (0.00135)	-0.01232 (0.00181)	0.08	-0.44	2.29	Request had an invalid product name per analysis by ClearSpending.Ru
Perc. of Auction Volume by Bur. in Last Calendar Week	-0.00219 (0.00146)		-0.02	-0.73	26.13	Percentage of all auctions (by volume) that Bureaucrat ran in the last calendar week of the year

Regions	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description
Public Perceptions of Corruption	-3.90612 (5.68683)	11.88589 (4.22785)	0	-2.09	2.43	Public perception of the severity of corruption as measured by popular survets
Regional Number of Corruption Cases	-7.72611 (7.64794)	1.44157 (7.1892)	0.01	-1.09	3.44	Number of corruption cases filed by officials in the region in which the auction was held
Regional Number of Corruption Convictions	-3.41518 (1.49156)	1.50883 (1.30614)	-0.01	-1.59	3.06	Number of corruption convictions secured by officials in the region in which the auction was held
Regional Number of Major Corruption Convictions	-2.39339 (1.09834)	0.99031 (0.94917)	0	-1.63	2.54	Number of major corruption convictions secured by officials in the region in which the auction was held
Regional Number of Officials Found Guilty	-3.50688 (1.17276)	1.99023 (1.15954)	0.02	-1.11	2.27	Number of corruption cases where officials were found guilty in the region in which the auction was held
Requests	**PwCorr - BurFE**	**PwCorr - OrgFE**	**Mean**	**Min**	**Max**	**Description**
Auction Admit Rate	-0.0005 (0.00126)	0.00199 (0.00152)	-0.02	-5.91	0.4	Percentage of supplier applicanst admitted to auction
Deposit (ths. Rubles)	0.00216 (0.00254)	0.04434 (0.00234)	0.01	-1.07	1.85	Amount bidders are required to deposit before entering auction
No. of Applicants	-0.02734 (0.00151)	-0.01991 (0.00205)	0.05	-0.77	28.49	Number of suppliers that submitted applications to participate in the auction
No. of Products Procured	-0.03381 (0.00318)	0.00682 (0.00275)	0.09	-0.98	4.1	Number of products overall
Number of 2-Digit Product Codes	-0.01862 (0.00192)	0.02185 (0.00357)	0.17	-0.49	4.61	Number of unique products (as measured by their two-digit codes)
Number of Request Revisions	0.00193 (0.00087)	-0.00838 (0.00114)	0.02	-0.16	21.29	Number of revisions that the Bureaucrat made to the contract before it was finalized
Reservation Price (bil. rubles)	0.03909 (0.00456)	-0.09431 (0.00252)	0.04	-1	1.45	Amount of Reservation price in billions of rubles
Winners	**PwCorr - BurFE**	**PwCorr - OrgFE**	**Mean**	**Min**	**Max**	**Description**
1[Supp. Sold Product Often (Volume)]	-0.00936 (0.00172)	0.0024 (0.00219)	-0.05	-0.89	1.12	Indicator if the main product was also the most common product supplied overall by the Supplier (volume)
1[Supplier Above Median Profit]	0.0016 (0.00115)	0.00307 (0.00139)	-0.09	-1.38	0.73	Indicator if the Supplier has above-median profit relative to the other suppliers in the dataset

1[Supplier Above Median Revenue]	-0.00438 (0.00098)	0.00523 (0.00141)	-0.09	-1.53	0.65	Indicator if the Supplier has above-median revenue relative to the other suppliers in the dataset
1[Supplier from Same Region]	0.01154 (0.00142)	-0.01482 (0.00133)	0.06	-1.4	0.71	Indicator if the Supplier is located in the same region as the End User
1[Supplier has Foreign Ownership]	-0.00367 (0.00129)	0.00403 (0.00123)	-0.03	-0.13	7.56	Indicator if the Supplier has foreign ownership
1[Supplier is Exporter]	0.00263 (0.00103)	0.01086 (0.00169)	-0.06	-0.17	5.76	Indicator if the Supplier has exporting activities
1[Supplier is Federal Government Agency]	-0.00263 (0.00173)	0.02244 (0.0035)	0.01	-0.06	15.59	Indicator if the Supplier is registered as a federal government agency
1[Supplier is Importer]	-0.00304 (0.0014)	0.00563 (0.00191)	-0.08	-0.43	2.32	Indicator if the Supplier has importing activities
1[Supplier is NGO]	0.00112 (0.00095)	-0.00012 (0.001)	0.01	-0.02	47.09	Indicator if the Supplier is a nongovernmental organization
1[Supplier is New Firm]	0.00407 (0.01055)	0.02785 (0.01359)	0.06	-0.45	2.22	Indicator if Supplier is a very new firm
1[Supplier is Private Company]	0.0059 (0.00355)	-0.03675 (0.00607)	-0.04	-2	0.5	Indiciator if Supplier is a Private Company
1[Supplier is Regional Government Agency]	-0.00379 (0.0009)	-0.00189 (0.00112)	-0.07	-0.15	6.89	Indicator if the Supplier is registered as a regional government agency
1[Supplier is Wholesale Trader]	0.00462 (0.00146)	0.00532 (0.00127)	0.02	-0.4	2.51	Indicator if Supplier is a wholesale trader
1[Supplier is from Same Postal Code]	-0.00421 (0.00146)	0.02333 (0.00159)	0	-0.59	1.71	Indicator if the Supplier is located in the same postal code as the End User
Supplier - Export Products	0.00453 (0.00114)	0.01001 (0.00193)	-0.08	-0.28	3.97	Number of unique products the Supplier exports
Supplier - Import Products	-0.00495 (0.00161)	0.01536 (0.00224)	-0.07	-0.33	3.67	Number of unique products the Supplier imports
Supplier Age	0.00212 (0.00161)	-0.00127 (0.00162)	-0.08	-1.2	1.88	Age of supplier in years
Supplier All Contracts / Revenue	-0.00294 (0.00123)	0.00919 (0.00152)	0.01	-0.57	2.62	Ratio of Supplier's total contract volume to revenue
Supplier Exports / Revenue	0.00258 (0.00071)	0.00584 (0.00096)	-0.02	-0.07	32.69	Ratio of Supplier's total export volume to revenue

Supplier Imports / Revenue	0.00179 (0.00097)	-0.00684 (0.00109)	-0.01	-0.15	13.97	Ratio of Supplier's total import volume to revenue
Supplier No. of Subsidiaries	-0.00348 (0.00117)	0.00138 (0.0015)	-0.08	-0.3	12.6	Number of subsidiaries owned by Supplier
Supplier Number of Contracts Won	-0.00692 (0.00109)	0.00481 (0.00131)	-0.01	-0.93	2.73	Cumulative number of contracts won by the supplier
Supplier Number of Contracts Won from Large SOEs	-0.00884 (0.00114)	0.00481 (0.00149)	0.01	-0.55	4.24	Cumulative number of contracts won by the supplier under FZ-223 regulating contracts with government agencies
Supplier Number of Countries Exported to	0.00466 (0.00106)	0.01216 (0.00182)	-0.07	-0.3	3.37	Number of unique countries that the Supplier exported to
Supplier Number of Employees	0.00015 (0.00136)	-0.0039 (0.00197)	-0.03	-1.98	4.15	Number of employees working for Supplier
Supplier Product HHI Index (Auctions)	-0.01018 (0.00212)	-0.00861 (0.00385)	-0.12	-2.11	1.61	HHI measuring number of auctions (count) won by supplier across two-digit product types
Supplier Product HHI Index (Volume)	0.00025 (0.00167)	-0.00605 (0.00267)	-0.03	-2.24	1.55	HHI measuring sales volume of all auctions won by supplier across two-digit product types
Supplier Profit	0.00068 (0.00148)	0.00888 (0.00186)	-0.11	-4.92	4.13	Supplier net profit
Supplier Profit Per Employee	-0.00401 (0.00179)	0.01231 (0.00218)	-0.08	-0.7	3.09	Ratio of Supplier profits to number of employees
Supplier Revenue (log)	-0.0043 (0.00212)	0.01738 (0.00291)	-0.13	-1.59	3.04	Supplier revenue (log)
Supplier SOE Contracts / Revenue	-0.00309 (0.00089)	0.00983 (0.00185)	0.02	-0.37	3.93	Ratio of Supplier's total volume of contracts with state-owned enterprises to revenue
Supplier Total Assets (log)	-0.00402 (0.00133)	0.01063 (0.00195)	-0.13	-1.12	3.61	Supplier total assets (log)
Supplier Value of Auctions Won (Cumulative)	-0.00388 (0.00176)	0.01712 (0.00218)	-0.18	-7.8	3.9	Total sales volume of auctions the Supplier was participating in was running simultaneously in the same month
Supplier Value of Auctions Won (Cumulative, bil. rubles)	-0.00455 (0.00107)	0.00514 (0.00103)	-0.15	-3.94	2.11	Total sales volume of the auctions the Supplier had participated in cumulatively to the date of the auction
Supplier is Registered with Tax Authorities	-0.0012 (0.00092)	0.00345 (0.00117)	-0.06	-1.92	0.52	Indicator if the Supplier is registered with the tax authorities

Supplier is Small Firm	0.00139 (0.00107)	-0.00298 (0.0009)	0.09	-0.33	2.99	Indicator if the Supplier is registered as a small firm
Supplier on Dishonest List	-0.00486 (0.00109)	-0.00427 (0.00085)	0.01	-0.32	3.12	Indicator if Supplier is on the official list of dishonest suppliers

Notes: The table describes the full set of variables included in the analysis of bureaucrat and organization effectiveness. The columns 'PwCorr-BurFE' and 'PwCorr-OrgFE' give the pairwise coefficient and standard error between each variable and the estimated bureaucrat and organization effects. Bureaucrat pairwise coefficients are blank for variables not included in the models examining organization effectiveness, while organization pairwise coefficients are blank for variables not included in the models examining bureaucrat effectiveness. Basic summary statistics for each variable are also given, as well as a description of how each was calculated. Firms with less than 100 workers and less than 25 percent ownership by a larger firm do not have to register with the Russian statistical authorities, and are thus not covered by the *Ruslana* data. This includes microenterprises and individual entrepreneurs who participate in procurement and will have missing data. To account for the missing data, we include dummy variables indicating missing data and require the regularization procedure to include them in the final model.

G Additional Results on Policy Design with a Heterogeneous Bureaucracy

FIGURE G.1: END USERS DO NOT CHANGE THE TIMING OF THEIR PROCUREMENT IN ANTICIPATION OF PREFERENCE LAWS

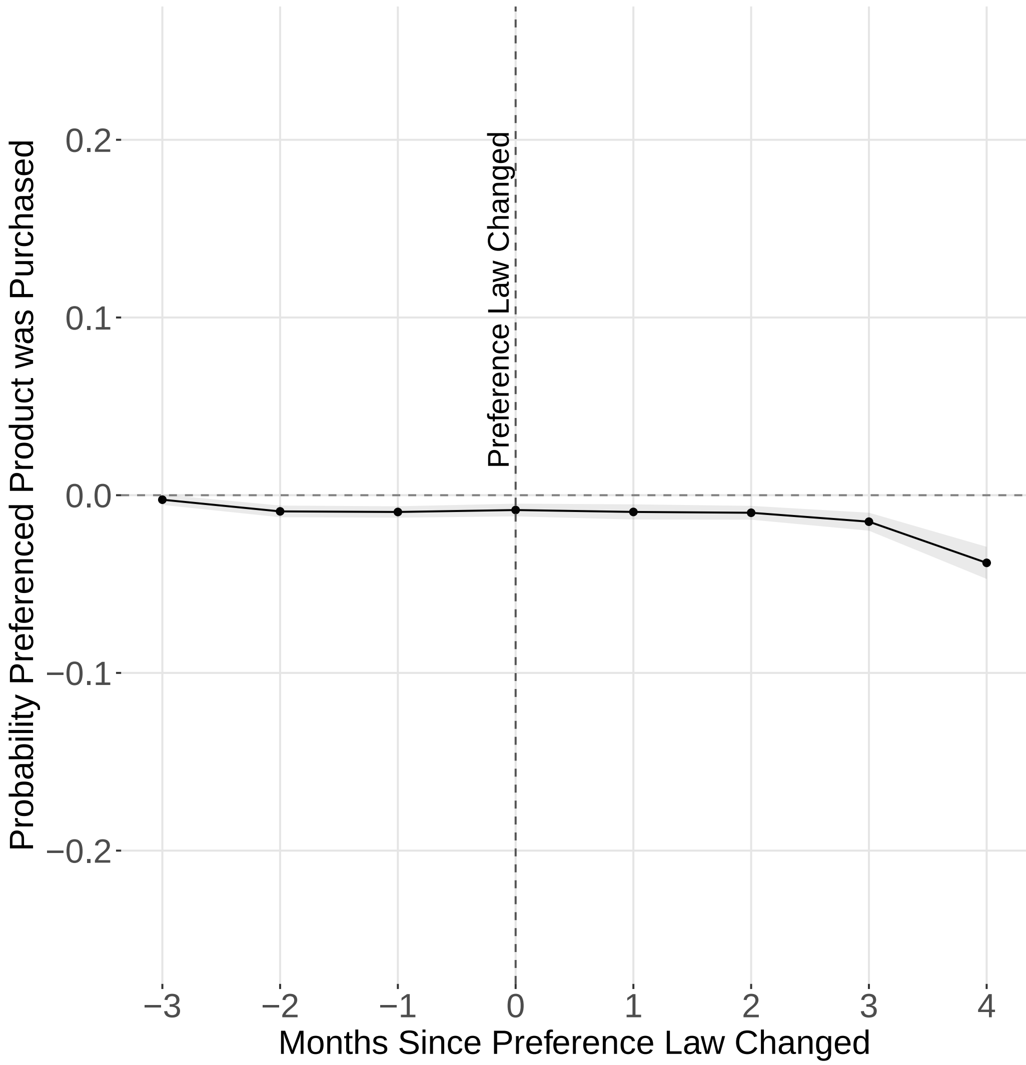

Notes: The figure shows the results of an event study analysis of the timing of procurement around the time the preference list is published each year. The x-axis is measured in the number of months preceding or following the activation of the annual preferences laws in 2011, 2012, 2013, and 2014. The dotted vertical lines indicates when the policy was became active. The y-axis in each plot shows the month-specific coefficients from estimation of equation: $Preferenced_{gt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \mathbf{1}\{t - \text{ListMonth}_t = s\} + \varepsilon_{igt}$, where Preferenced_{gt} is a dummy indicating that g is on the preferences list in the year month t falls within and ListMonth_t is the month closest to month t in which a preference list is published. $\mathbf{X}_{igt}$ are the same controls we use in Section 5, but we remove the month fixed effects. ε_{igt} is an error term we allow to be clustered by month and good.

FIGURE G.2: HETEROGENEITY OF BID PREFERENCES' EFFECT BY ORGANIZATION EFFECTIVENESS

PANEL A: DIFFERENCE IN DIFFERENCES BY ORGANIZATION EFFECTIVENESS DECILE

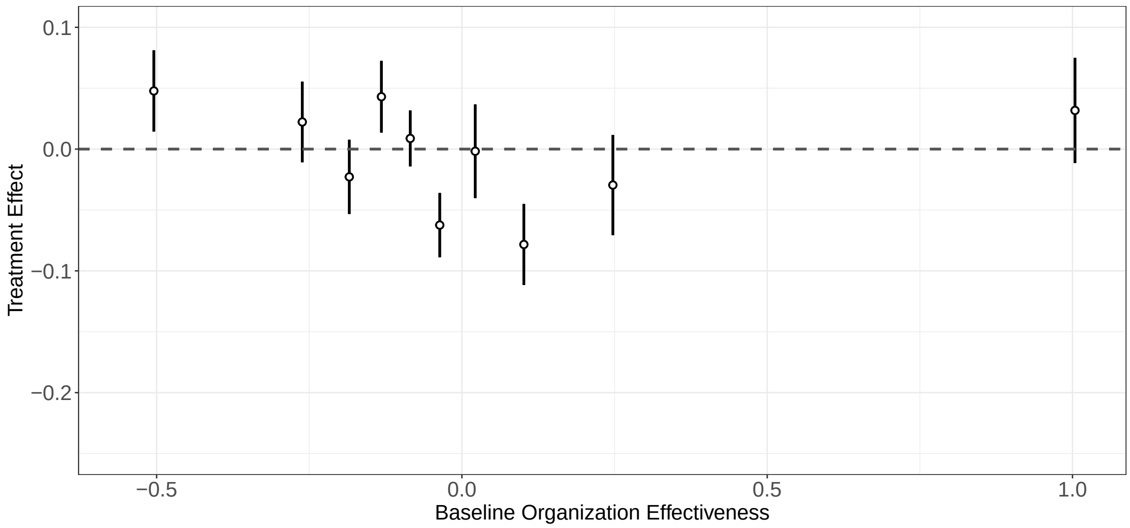

PANEL B: EVENT STUDY BY ORGANIZATION EFFECTIVENESS

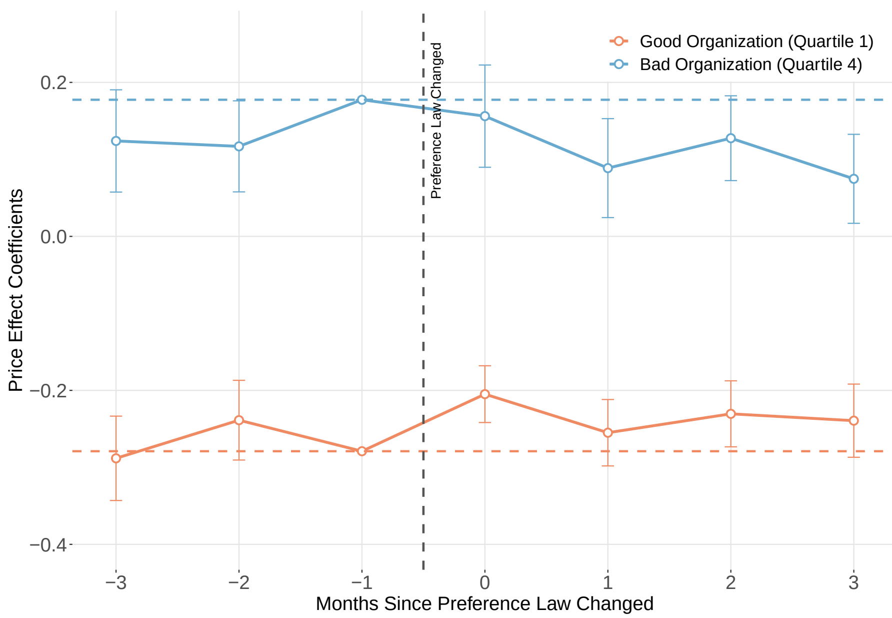

Notes: The figure shows how the impacts of the introduction of bid preferences varies by the effectiveness of the implementing organization. Panel A shows estimates from implementing the triple difference model (8) to estimate separate effects for each decile of organization effectiveness: $y_{igt} = \sum_{k=1}^{10}\left\{D_{kb}+D_{kj}\times\left(\rho_k \text{Preferenced}_{gt}+\eta_k \text{PolicyActive}_t+\pi_k \text{Preferenced}_{gt}\times\text{PolicyActive}_t\right)\right\}+\mathbf{X}_{igt}\beta+\mu_g+\mu_t+\varepsilon_{igt}$ where D_{kj} and D_{kb} are indicators for organization j and bureaucrat b belonging to decile k of their respective distributions of effectiveness. The horizontal axis plots the average effectiveness within the relevant decile, while the vertical axis plots the estimated treatment effects π_k with their 95% confidence intervals. Panel B extends the event study (6) shown in figure 6 (see notes to figure 6 for details) to estimate separate effects for the top and bottom quartile of organizations. Rather than normalizing the reference month (the month before the preference list is published) to zero, we normalized it to the baseline performance in each group to better highlight how different their performance was before the preferences were introduced, and how their performance converges as a result of the preferences.

Figure G.3: Heterogeneity of Effect of Bid Preferences on Number of Bidders

Panel A: Difference in Differences by Bureaucrat Effectiveness

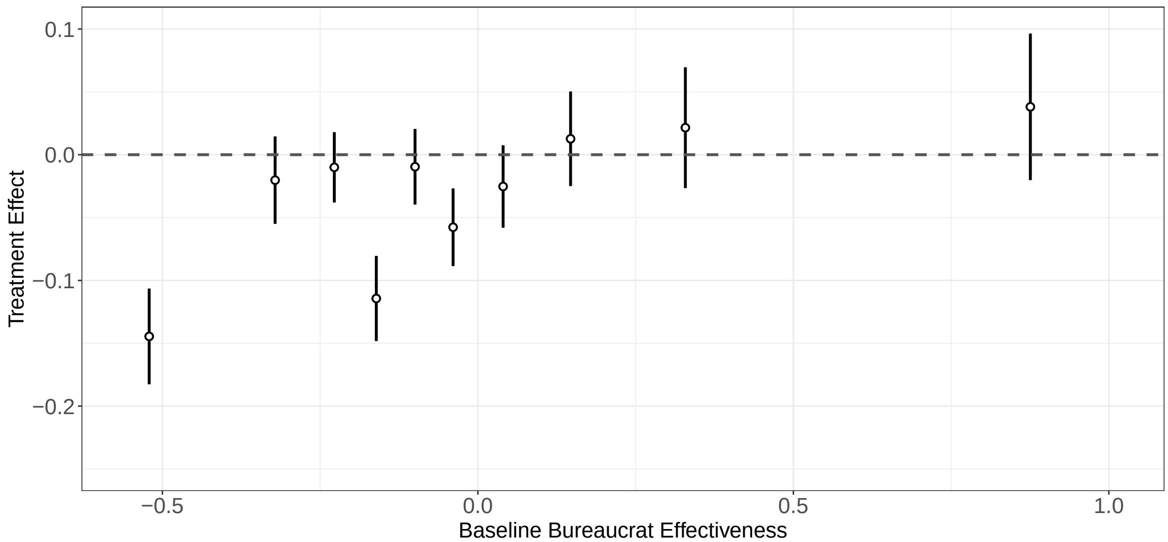

Panel B: Difference in Differences by Organization Effectiveness

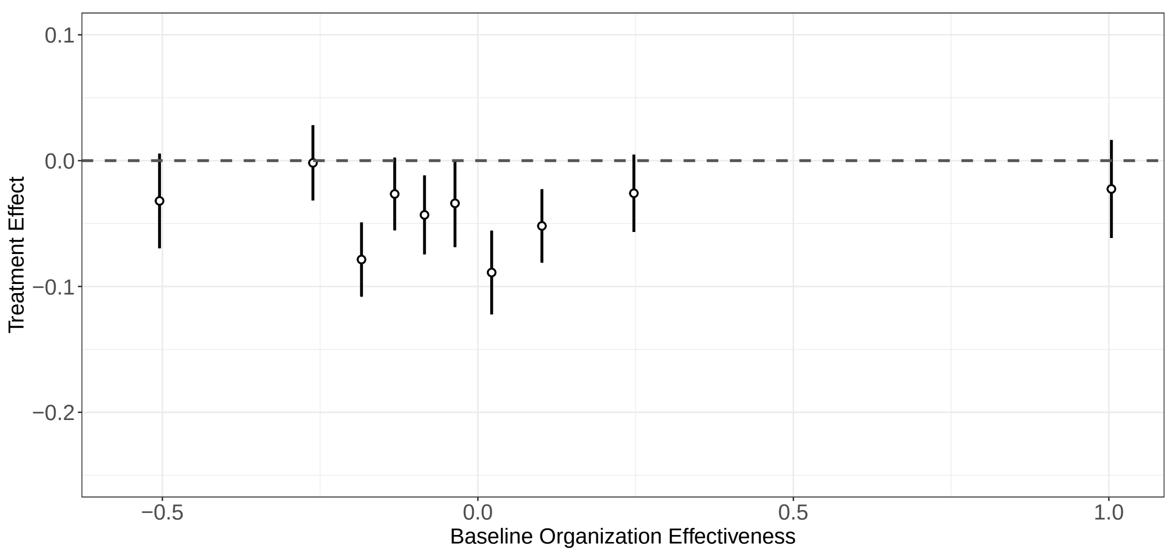

Notes: The figure shows how the impacts of the introduction of bid preferences on the number of bidders varies by the effectiveness of the implementing buyer. Panel A shows estimates from implementing the triple difference model (8) to estimate separate effects for each decile of bureaucrat effectiveness: $y_{igt} = \sum_{k=1}^{10} \left\{ D_{kj} + D_{kb} \times \left(\rho_k \text{Preferenced}_{gt} + \eta_k \text{PolicyActive}_t + \pi_k \text{Preferenced}_{gt} \times \text{PolicyActive}_t \right) \right\} + \mathbf{X}_{igt}\beta + \mu_g + \mu_t + \varepsilon_{igt}$ while Panel B estimates separate effects for each decile of organization effectiveness: $y_{igt} = \sum_{k=1}^{10} \left\{ D_{kb} + D_{kj} \times \left(\rho_k \text{Preferenced}_{gt} + \eta_k \text{PolicyActive}_t + \pi_k \text{Preferenced}_{gt} \times \text{PolicyActive}_t \right) \right\} + \mathbf{X}_{igt}\beta + \mu_g + \mu_t + \varepsilon_{igt}$ where D_{kj} and D_{kb} are indicators for organization j and bureaucrat b belonging to decile k of their respective distributions of effectiveness. The horizontal axis plots the average effectiveness within the relevant decile, while the vertical axis plots the estimated treatment effects π_k with their 95% confidence intervals.

FIGURE G.4: HETEROGENEITY OF EFFECT OF BID PREFERENCES IN PHARMACEUTICALS SUBSAMPLE BY BUREAUCRAT EFFECTIVENESS

PANEL A: HETEROGENEITY OF EFFECT ON PRICES

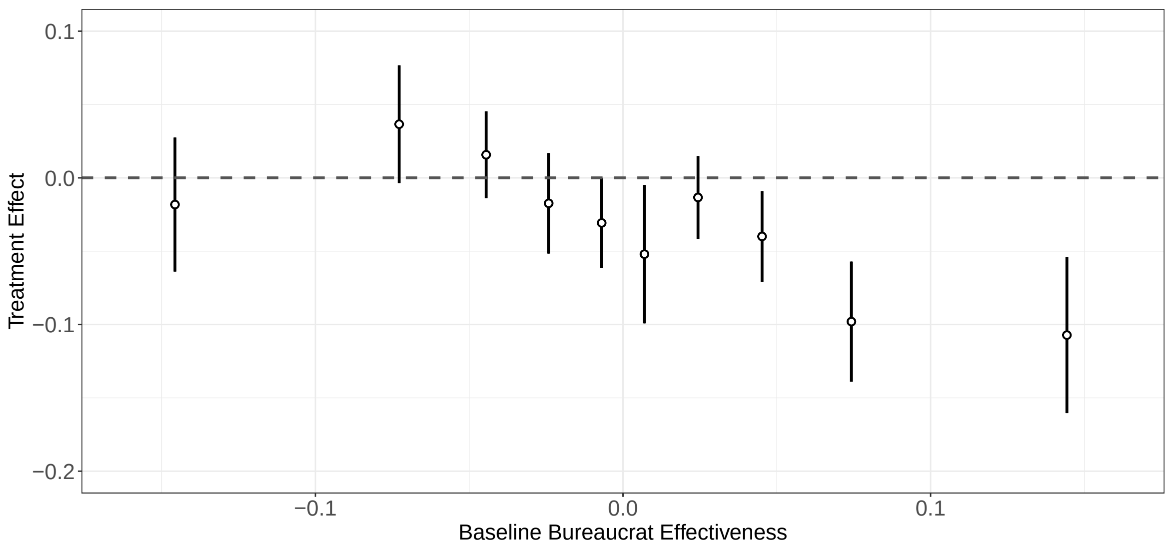

PANEL B: HETEROGENEITY OF EFFECT ON PROBABILITY OF A DOMESTIC WINNER

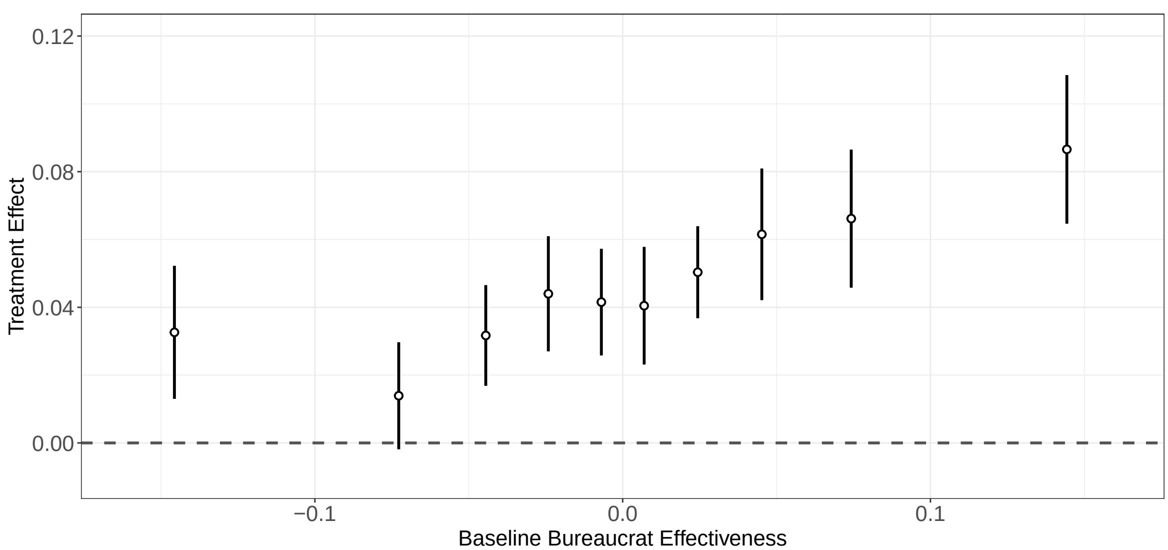

Notes: The figure shows how the impacts of the introduction of bid preferences varies by the effectiveness of the implementing buyer in the pharmaceuticals subsample. We estimate the triple difference model (8) to estimate separate effects for each decile of bureaucrat effectiveness: $y_{igt} = \sum_{k=1}^{10} \left\{D_{kj} + D_{kb} \times \left(\rho_k \text{Preferenced}_{gt} + \eta_k \text{PolicyActive}_t + \pi_k \text{Preferenced}_{gt} \times \text{PolicyActive}_t\right)\right\} + \mathbf{X}_{igt}\beta + \mu_g + \mu_t + \varepsilon_{igt}$ where D_{kj} and D_{kb} are indicators for organization j and bureaucrat b belonging to decile k of their respective distributions of effectiveness. The horizontal axis plots the average effectiveness within the relevant decile, while the vertical axis plots the estimated treatment effects π_k with their 95% confidence intervals. Panel A shows effects on prices, while Panel B shows effects on the probability the winning bid offers domestically manufactured pharmaceuticals.

FIGURE G.5: PREDICTORS OF HETEROGENEITY OF EFFECT OF BID PREFERENCES ON SPENDING QUALITY

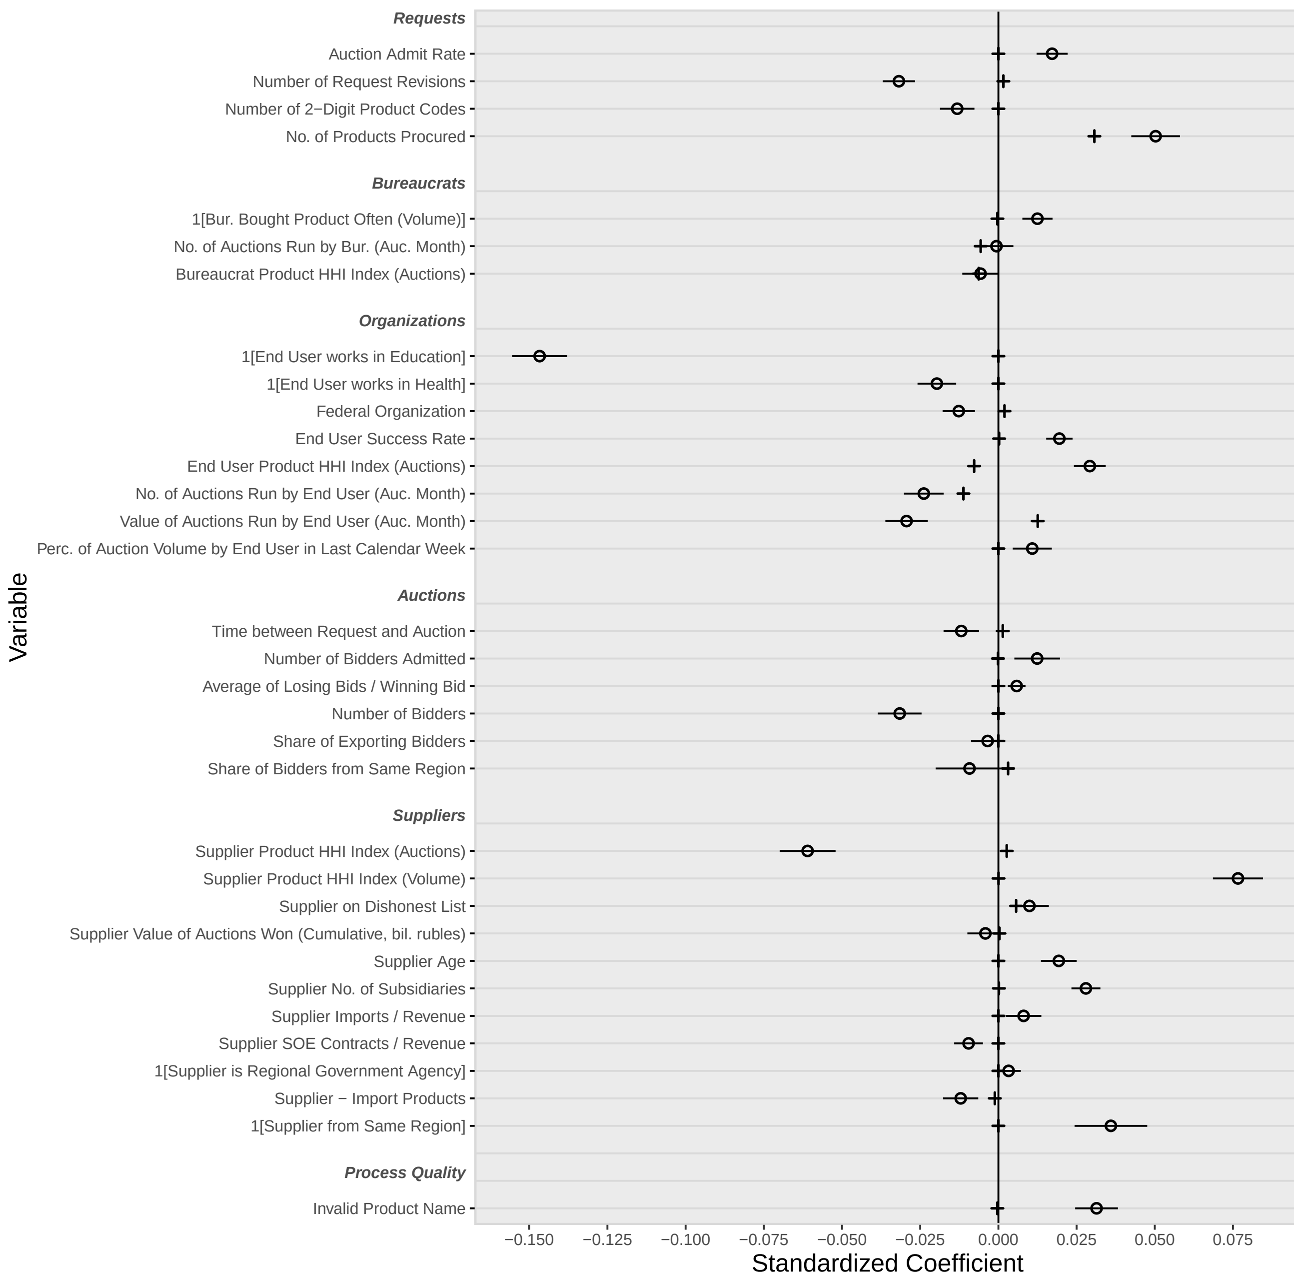

Notes: The figure shows the results of estimating our triple-differences specification for heterogeneity of the effect of bid preferences (9): $y_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \mathbf{Z}_{igt}\boldsymbol{\theta} + \text{Preferenced}_{gt} \times \mathbf{Z}_{igt}\boldsymbol{\gamma} + \text{PolicyActive}_t \times \mathbf{Z}_{igt}\boldsymbol{\eta} + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \mathbf{Z}_{igt}\boldsymbol{\pi} + \varepsilon_{igt}$ where the elements of the vector of observables $\mathbf{Z}_{igt}$ are picked by LASSO using the largest regularization penalty that returns 30 non-zero coefficients. The coefficients from the LASSO are shown as crosses, while the circles show the coefficients and 95% confidence intervals of a multivariate regression including the 30 observables.

FIGURE G.6: CORRELATES OF PRICE DID: ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS

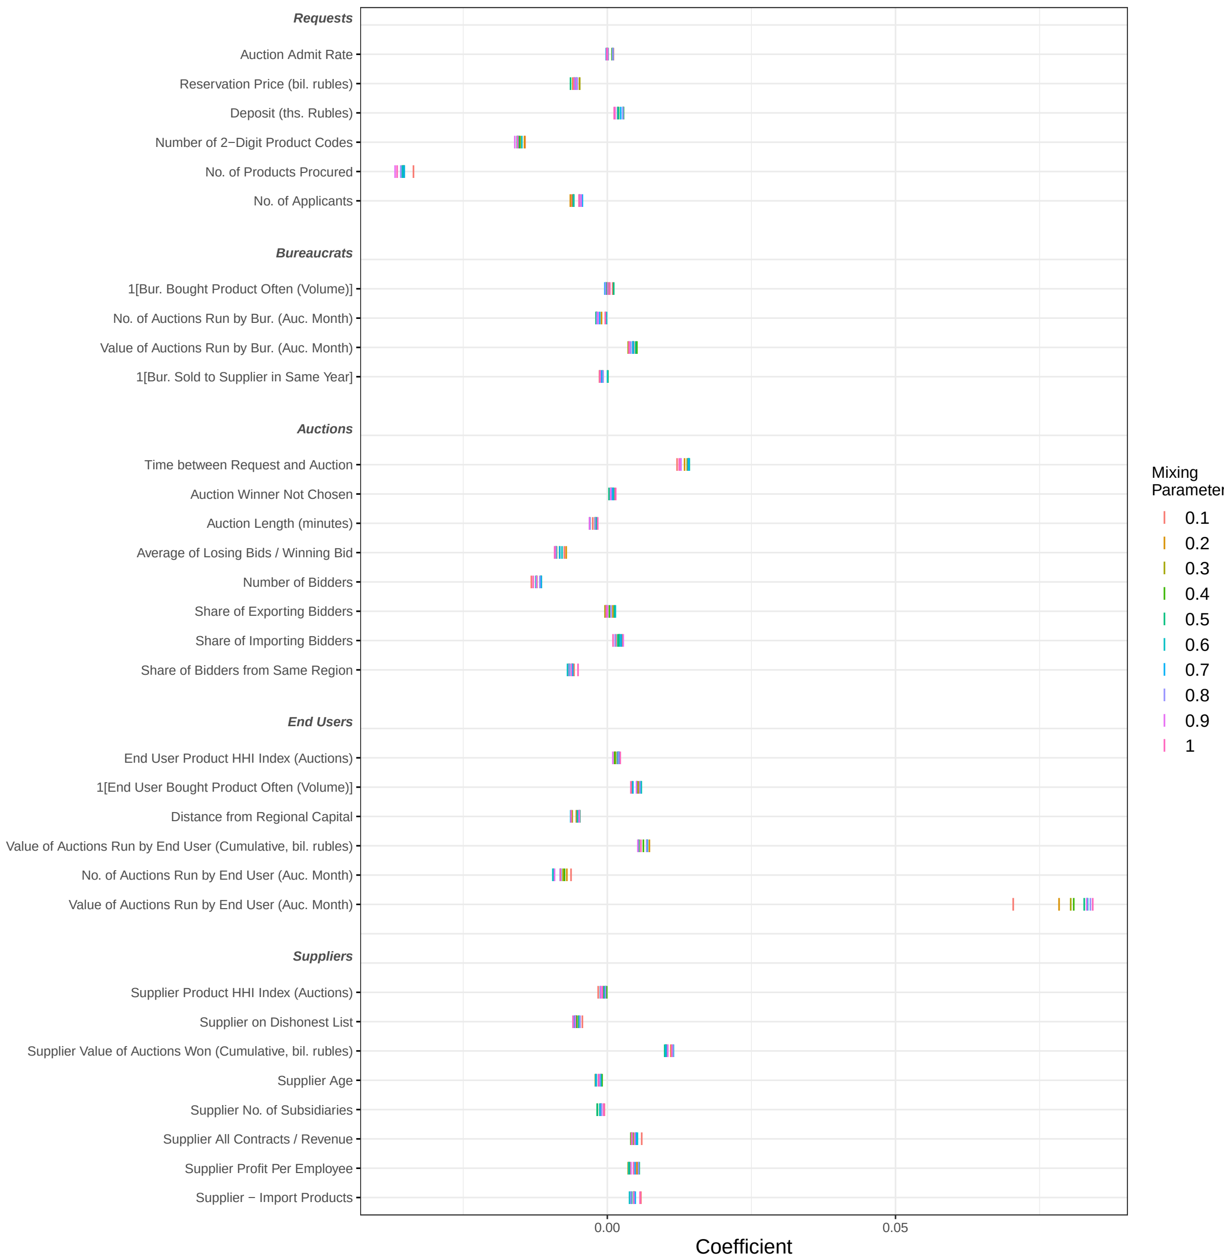

Notes: The figure shows the coefficients from the elastic net regularization procedure on the estimated difference-in-differences effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. A mixing parameter of 1 represents LASSO, our baseline model. The variables shown are from the base model shown in Figure 8 where the values of the regularization penalty lambda λ is chosen to return 30 variables.

FIGURE G.7: CORRELATES OF QUALITY DID: ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS

Notes: The figure shows the coefficients from the elastic net regularization procedure on the estimated difference-in-differences effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. A mixing parameter of 1 represents LASSO, our baseline model. The variables shown are from the base model shown in Figure G.5 where the values of the regularization penalty lambda λ are chosen to return 30 variables.

FIGURE G.8: CONSTANT SEMI-ELASTICITY APPROXIMATION FOR EQUIVALENT PREFERENCE POLICY EXERCISE

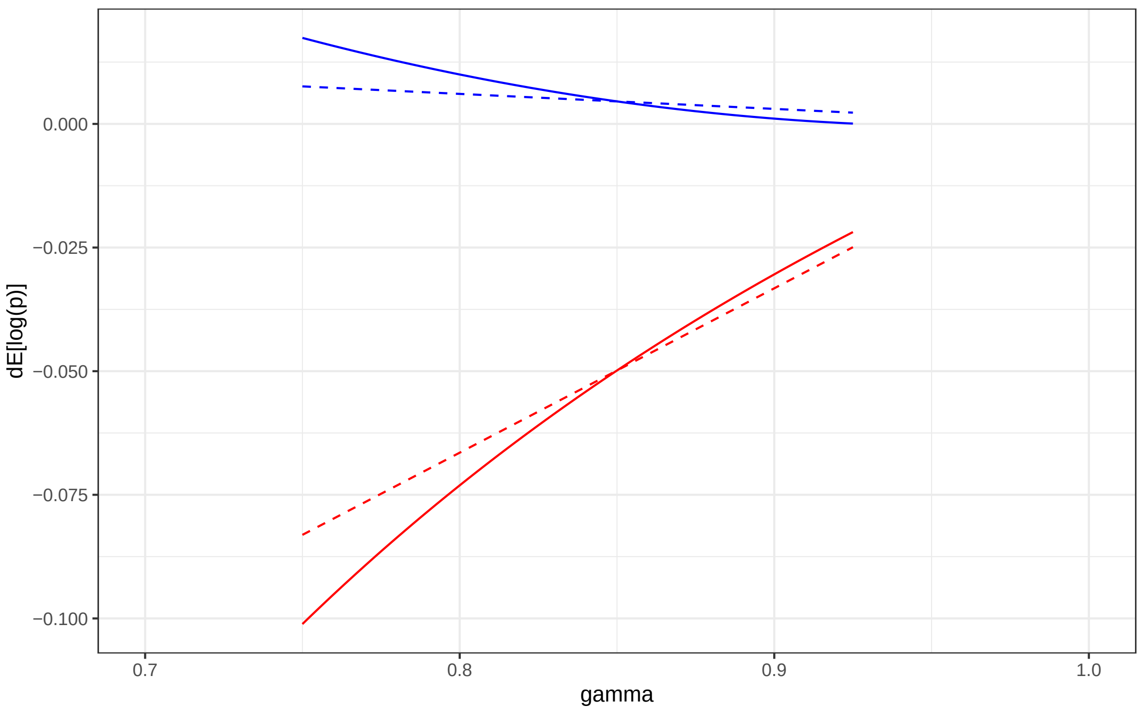

Notes: The figure shows a calibration of the model in section 4 in solid lines, together with the constant semi-elasticity approximation discussed in section 6.4 in dashed lines. To calibrate the model we set $\log(\bar{\theta}) = 1$; the pareto parameter of the productivity distribution of the foreign bidders to be $\delta_F = 1.5$; and the pareto parameter for the local bidders such that the mean productivity is 10% higher for foreign bidders: $\delta_L = 1.588$. We show how the expected log price changes as γ, the fraction of the final bid that a foreign winner receives, changes, as described in proposition 2. The blue lines show this for a high-effectiveness buyer in case 1 of the proposition (specifically, we set $\alpha_c + \psi_c = 0.25$). The red lines show this for a low-effectiveness buyer in case 3 of the proposition (specifically, we set $\alpha_c + \psi_c = 0.85$). The solid and dashed lines are not substantially different from each other.

H Additional Figures and Tables

FIGURE H.1: PROCUREMENT PROCESS FLOW-CHART

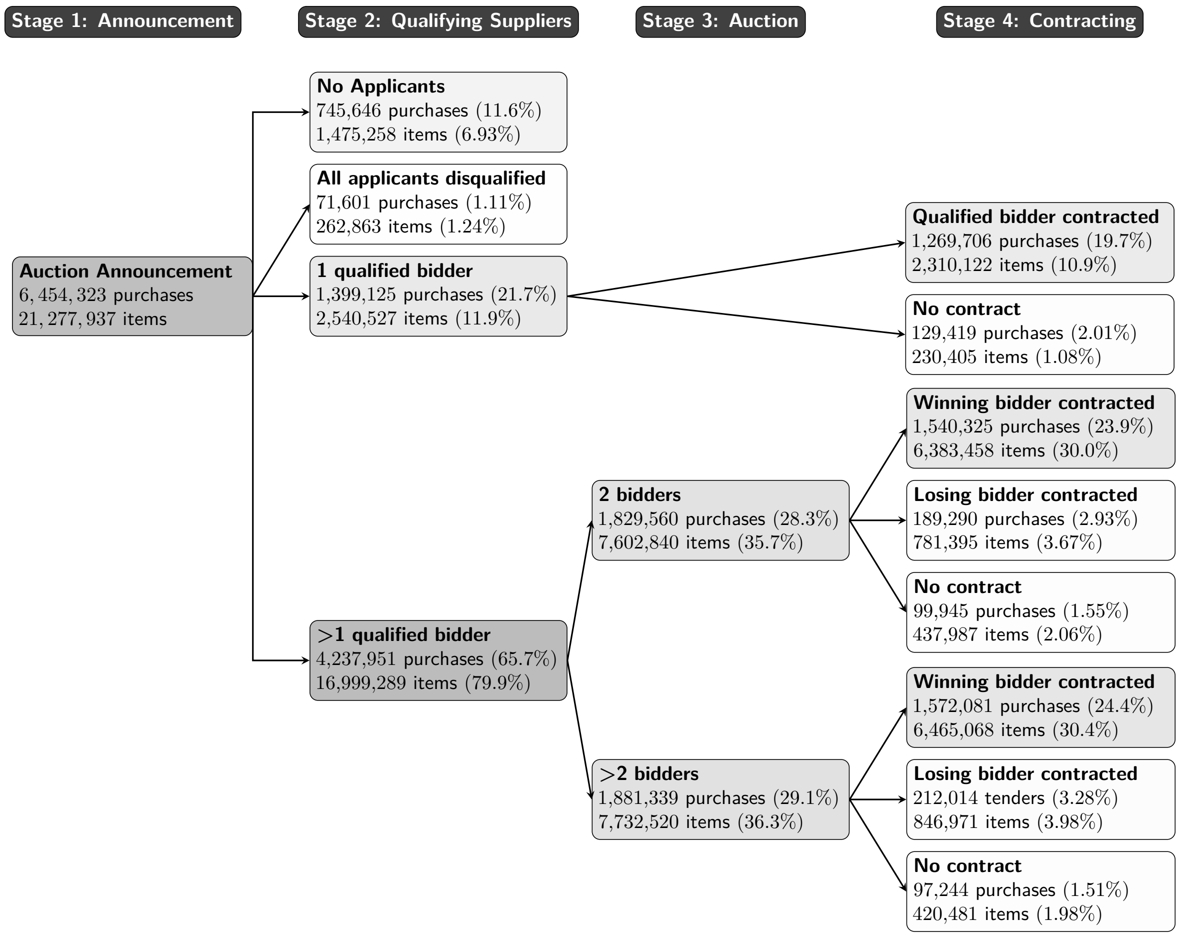

Notes: This figure lays out the stages of the process public procurement purchases of off-the-shelf goods through electronic auctions follow in Russia. Numbers are based on all purchases made under laws 94 and 44 in 2011-2016. The stages are described in detail in Sub-section 2.1.

FIGURE H.2: EXAMPLE OF BUREAUCRATS DENYING APPLICANTS

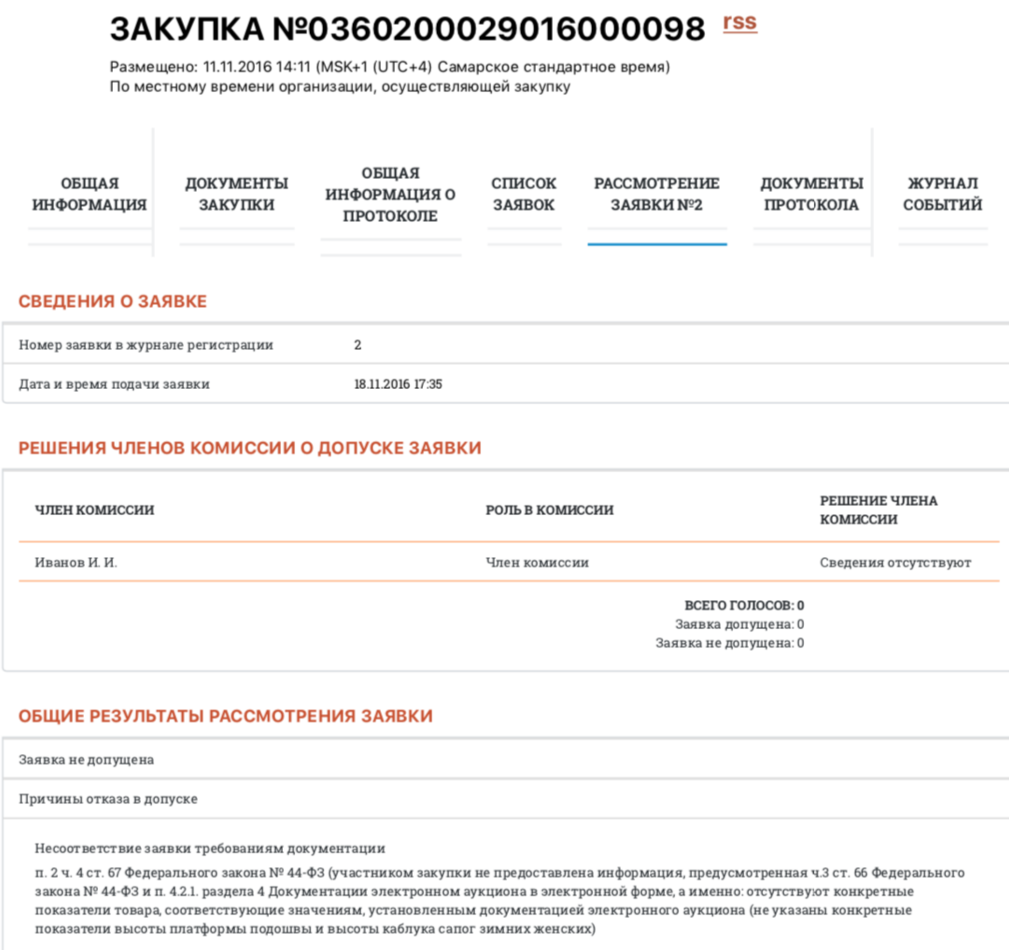

ЗАКУПКА №0360200029016000098 rss

Размещено: 11.11.2016 14:11 (MSK+1 (UTC+4) Самарское стандартное время)
По местному времени организации, осуществляющей закупку

ОБЩАЯ ИНФОРМАЦИЯ | ДОКУМЕНТЫ ЗАКУПКИ | ОБЩАЯ ИНФОРМАЦИЯ О ПРОТОКОЛЕ | СПИСОК ЗАЯВОК | РАССМОТРЕНИЕ ЗАЯВКИ №2 | ДОКУМЕНТЫ ПРОТОКОЛА | ЖУРНАЛ СОБЫТИЙ

СВЕДЕНИЯ О ЗАЯВКЕ

Номер заявки в журнале регистрации	2
Дата и время подачи заявки	18.11.2016 17:35

РЕШЕНИЯ ЧЛЕНОВ КОМИССИИ О ДОПУСКЕ ЗАЯВКИ

ЧЛЕН КОМИССИИ	РОЛЬ В КОМИССИИ	РЕШЕНИЕ ЧЛЕНА КОМИССИИ
Иванов И. И.	Член комиссии	Сведения отсутствуют

ВСЕГО ГОЛОСОВ: 0
Заявка допущена: 0
Заявка не допущена: 0

ОБЩИЕ РЕЗУЛЬТАТЫ РАССМОТРЕНИЯ ЗАЯВКИ

Заявка не допущена

Причины отказа в допуске

Несоответствие заявки требованиям документации

п. 2 ч. 4 ст. 67 Федерального закона № 44-ФЗ (участником закупки не предоставлена информация, предусмотренная ч.3 ст. 66 Федерального закона № 44-ФЗ и п. 4.2.1. раздела 4 Документации электронном аукциона в электронной форме, а именно: отсутствуют конкретные показатели товара, соответствующие значениям, установленным документацией электронного аукциона (не указаны конкретные показатели высоты платформы подошвы и высоты каблука сапог зимних женских)

Notes: This screenshot is taken from the official protocol for Request #0360200029016000098, an electronic auction for winter shoes conducted by an orphanage in November 2016 in Saratov, Russia. Applicant supplier #2 was rejected by the five-member commission on the grounds that the supplier's application did not adequately described the goods offered. More specifically, the application did not contain information about the height of the shoe sole nor the heel of the boot. Bureaucrats applying these requirements so tightly limit the number of suppliers that can participate in the auction.

TABLE H.1: PRODUCTS COVERED BY PREFERENCE LAWS, BY YEAR

2011	2012	2013	2014
Live animals	Live animals	Live pigs	Meat and meat products
Textiles	Fresh, chilled, and frozen pork	Fresh, chilled, and frozen pork	Fish and fish products
Clothing and fur products	Sugar	Meat, sausage and other meat products	Salt
Leather and leather goods	Textiles	Cheese, cream and milk	Rice, starches and flour
Chemical products and pharmaceuticals	Clothing and fur products	Rice	Grains, fruits and vegetables (various)
Ratio and television equipment	Leather and leather goods	Textiles	Bread, desserts, and chocolate
Medical and measurement equipment	Chemical products and pharmaceuticals	Clothing and fur products	Pharmaceuticals
Cars, trailers and semitrailers	Combine harvesters	Leather and leather goods	Medical and measurement equipment
Transport vehicles (excluding cars)	Self-propelled vehicles	Pharmaceuticals	Ceramic products
	Machinery parts	Agricultural machinery	Iron, steel and ferroalloys (incl. pipes)
	Agricultural machinery	Ratio and television equipment	Steam boilers
	Ratio and television equipment	Medical and measurement equipment	Agricultural machinery
	Medical and measurement equipment	Cars, trailers and semitrailers	Metals and mining equipment
	Cars, trailers and semitrailers	Transport vehicles (excluding cars)	
	Transport vehicles (excluding cars)	Sporting equipment (various)	

TABLE H.2: TOTAL PROCUREMENT IN RUSSIA BY TYPE OF MECHANISM USED

Type	**2011**	%	**2012**	%	**2013**	%	**2014**	%	**2015**	%	**2016**	%	**2011-2016**	%
Electronic Auctions	76.60	46.5	107.65	54.55	106.78	57.98	72.62	51.80	45.13	51.12	45.95	56.39	454.73	53.12
Single Supplier	39.08	23.7	42.95	21.76	39.30	21.34	24.60	17.54	19.61	22.22	19.54	23.98	185.08	21.62
Request for Quotations	6.07	3.7	5.66	2.87	5.32	2.89	1.67	1.19	0.91	1.03	0.77	0.94	20.39	2.38
Open Tender	30.70	18.6	40.86	20.70	32.58	17.69	34.08	24.31	15.82	17.92	10.47	12.85	164.50	19.22
Other Methods	12.17	7.4	0.22	0.11	0.17	0.09	7.23	5.16	6.81	7.72	4.75	5.83	31.36	3.66
Total Procurement	164.62		197.33		184.15		140.19		88.28		81.49		856.06	
Russian Non-Resource GDP	1,720.89		1,873.42		1,989.28		1,786.30		1,231.35		1134.47		9,735.72	
Procurement / Non-Resource GDP (%)	9.6		10.5		9.3		7.8		7.2		7.2		8.8	
Exchange Rate (RUB/USD)	29.37		30.96		31.97		39.20		62.01		66.34		43.31	

Notes: This table presents summary statistics about how much procurement was completed under federal laws 94FZ and 44FZ each year according to the mechanism used. All sums are measured in billions of US dollars at current prices using the average ruble-dollar exchange rates shown. Data on Russian procurement comes from the central nationwide Register for public procurement in Russia (`http://zakupki.gov.ru/epz/main/public/home.html`). Data on Russian GDP comes from International Financial Statistics (IFS) at the International Monetary Fund (`http://data.imf.org/`), which we adjust using the percentage of GDP coming from natural resources rents as calculated by the World Bank (`http://data.worldbank.org/indicator/NY.GDP.TOTL.RT.ZS?locations=RU&name_desc=true`).

Appendix References

ABOWD, JOHN M., CREECY, ROBERT H., & KRAMARZ, FRANCIS. 2002. *Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data*. Census Bureau Technical Paper TP-2002-06.

BERTRAND, MARIANNE, & SCHOAR, ANTOINETTE. 2003. Managing with style: The effect of managers on firm policies. *Quarterly Journal of Economics*.

CARD, DAVID, HEINING, JÖRG, & KLINE, PATRICK. 2013. Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics*, **128**, 967–1015.

CHETTY, RAJ, FRIEDMAN, JOHN N., & ROCKOFF, JONAH E. 2014. Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood. *American Economic Review*, **104**, 2533–2679.

FRIEDMAN, JEROME, HASTIE, TREVOR, & TIBSHIRANI, ROBERT. 2013. *The Elements of Statistical Learning*. Berlin: Springer Series in Statistics.

LACETERA, NICOLA, LARSEN, BRADLEY, POPE, DEVING G., & SYDNOR, JUSTIN. 2016. Bid Takers or Market Makers? The Effect of Auctioneers on Auction Outcomes. *American Economic Journal: Microeconomics*, **8**, 195–229.

MAS, ALEXANDRE, & MORETTI, ENRICO. 2009. Peers at Work. *American Economic Review*.

MILGROM, PAUL. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.

PANG, SINNO JIALIN, & YANG, QIANG. 2010. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, **22**, 1345–1359.

RIFKIN, RYAN, & KLAUTAU, ALDEBARO. 2004. In Defense of One-Vs-All Classification. *Journal of Machine Learning Research*, **5**, 101–141.

SCULLEY, D. 2010. Web-Scale K-Means Clustering. *Proceedings of the 19th International Conference on World Wide Web*.

SILVER, DAVID. 2016. *Haste or Waste? Peer Pressure and the Distribution of Marginal Returns to Health Care*. Mimeo: UC Berkeley.

TORREY, LISA, & SHAVLIK, JUDE. 2009. Transfer Learning. *In: Handbook of Research on Machine Learning Applications*. IGI Global.