

Individuals and Organizations as Sources of State Effectiveness*

Michael Carlos Best[†] Jonas Hjort[‡] David Szakonyi[§]

October 2019

Abstract

How important are bureaucrats for the productivity of the state? And to what extent do the tradeoffs between different policies depend on the implementing bureaucrats' effectiveness? Using data on 16 million public procurement purchases in Russia during 2011–2016, we show that over 40 percent of the variation in quality-adjusted prices paid—our measure of performance—is due to the individual bureaucrats and organizations that manage procurement processes. Such differences in effectiveness matter for policy design. To illustrate, we show that a common procurement policy—bid preferences for domestic suppliers—dramatically improves performance, but only when implemented by *ineffective* bureaucrats.

JEL codes: O1, H1

*We are grateful to Ben Olken for his very helpful discussion of the paper at the AEA meetings. We also thank 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, Ori Shelef, Andrei Shleifer, David Silver, Eric Verhoogen, Guo Xu, and Danny Yagan for valuable comments; Andrei Yakovlev for institutional guidance; Andrey Druzaev, 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, 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, University College London, 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”. Any remaining errors are ours alone.

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

[‡]hjort@columbia.edu. Columbia University, 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 “mindless” bureaucracy—a collection of Weberian “machines” translating policy into output, thereby ensuring uniform provision of public services (Weber, 1921). In reality, the skills, organizational capacity, and priorities of bureaucrats differ. But by how much? And to what extent do bureaucrats help explain why some public entities are so much more effective than others at implementing the same policies?

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, in part because bureaucracies produce a wide array of outputs which cannot be measured in public-sector wide, administrative data. However, one task—the procurement of off-the-shelf goods—is performed throughout the state enterprise, and has a well-defined and quantifiable 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 over 40 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. Differences in effectiveness of such magnitude have far-reaching implications for policy design. To illustrate, we study the introduction of a bid preference regime common throughout the world. Under Russia’s bid preferences, contract-winners offering goods manufactured abroad are paid only 85 percent of their bid. Consistent with our model’s predictions, we find that preferences can 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).

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

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 enterprise. Third, Russia’s massive and diverse bureaucracy spans a wide range of state effectiveness. Fourth, the labor market of 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 of public procurement, 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 sellers wishing to bid on government contracts. As a result, less effective bureaucracies 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—for which we do not need to rely on a machine learning classifier. To estimate the causal impacts of individual bureaucrats and organizations on procurement performance, 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 identify the impact of individual bureaucrats and organizations on prices paid under weak assumptions on the nature of 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. The event studies also provide clear evidence supporting a causal interpretation of these effects.⁴

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

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

⁴Importantly, 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 bureaucrats do not sort across organizations based on unmodelled *match* effects, and that drift in effectiveness and switches are uncorrelated. The event studies provide compelling support for

To aggregate the impacts of individual bureaucrats and organizations on prices paid into an estimate of the share of the total variation that is 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 bureaucrats switching between organizations do not link *all* organizations and how the combined productivity effect of bureaucrats and organizations can nevertheless be identified.

We find that the individuals and organizations of the bureaucracy together account for more than 40 percent of the variation in quality-adjusted prices paid, of which individuals and organizations account for roughly equal shares. These results imply that moving the worst-performing quartile of procurers to 75th percentile-effectiveness would reduce procurement expenditures by around 11 percent, or USD 13 billion each year—roughly one fifth of the total amount spent on health care by the Russian government at federal, regional, and municipal level combined.

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 explore correlates of their estimated effectiveness (see also [Lacetera et al. , 2016](#)). Consistent with our model, we find that effective procurers set lower reservation prices, attract more applicants, and allow a higher share of applicants to participate in their auctions. While some other measures of bureaucrat behavior also predict bureaucrat effectiveness, a wide range—including regional measures of corruption—do not.

The second part of the paper focuses on the implications of heterogeneity in policy implementer effectiveness for the design of policy. We focus on bid preferences—a common form of industrial policy implemented through public procurement—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 partic-

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). In the public sector there are additional institutional reasons to expect these assumptions to hold (see Section 2).

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

ipation 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

To test our model's heterogeneous treatment effect predictions, we interact the bid preference regime with our estimates of the effectiveness of the bureaucrats in charge of implementation. We find that the small negative average effect on prices paid masks considerable heterogeneity. Our estimates imply savings of 17.5 percent when the policy is implemented by the least effective quartile of bureaucrats, but only 0.7 percent when implemented by the most effective quartile of bureaucrats, and that prices increase for the *most* effective bureaucrats (as has been shown for similar policies implemented in the U.S.).⁷ We also find that most procurer behaviors and processes that predict how the policy affects prices paid when implemented by a given bureaucrat and organization also predict these procurers' effectiveness in a constant policy regime. This suggests that policy changes can markedly affect state productivity even absent significant changes in policy implementer behavior.

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

⁶The average treatment effect of Russia's "buy local" program suggests that industrial policies in public procurement may on the whole be more successful in countries with *low* average bureaucratic effectiveness, such as Russia. The average treatment effect we estimate contrasts, in particular, with the effect of similar policies found in higher state effectiveness contexts (see e.g. Marion, 2007; Krasnokutskaya & Seim, 2011). This foreshadows our findings on how the impact of the policy varies with the effectiveness of the policy implementers within Russia.

⁷In 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.

ineffectiveness.

We contribute to two main strands of literature on state effectiveness. The first focuses on individuals and the incentives they face as sources of public sector productivity (see, among many others, [Dal Bo et al., 2013](#); [Duflo et al., 2013, 2018](#); [Bertrand et al., forthcoming](#); [Khan et al., 2016, 2018](#); [Rasul & Rogger, 2018](#)).⁸ 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 that avoids the limitations that arise from comparing workers and/or organizations (e.g. firms) (i) pursuing multiple objectives or engaging in different activities and/or (ii) based on wages and profits.⁹

The second strand of work on state effectiveness we contribute to 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 policy *implementation* is delegated to bureaucracies is often overlooked. Bureaucracies are likely to differ in effectiveness across contexts. We provide tools for the measurement of the effectiveness of a bureaucracy and show that effectiveness affects the relative costs and benefits of different policies (see also [Dehejia et al., forthcoming](#)).¹⁰ We are not aware of previous 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).

The rest of the paper is organized as follows. Section 2 presents background on the Rus-

⁸[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.

⁹The 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](#)).

¹⁰The treatment effect heterogeneity we find resonates with the findings 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](#)).

sian public procurement system and the data we use. Our conceptual framework is in Section 3, and in Section 4, we estimate the effectiveness of individual bureaucrats and organizations and their contribution to public sector output. In Section 5 we estimate the impact of the preference policy and its interaction with bureaucratic effectiveness. Section 6 concludes.

2 Background and Data on Public Procurement in Russia

2.1 A decentralized system with centralized rules

In 1991, Russia created an extremely decentralized system to perform procurement, a key government function comprising 10 percent of Russia's non-resource GDP.¹¹ 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 method of procurement, used for 53.5 percent of Russian procurement during our data period—in order to study bureaucrats and organizations performing exactly the same task. Auctions are conducted through one of five designated, 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 making the purchase.

Figure OA.1 traces out the steps involved in a purchase together with the number of purchases in 2011–2016 that followed each path from announcement to contract. Each purchase starts with an auction announcement, drawn up by a procurement officer. The announcement contains technical details on the item(s) to be purchased (from clients), a maximum price for the lot, the required security deposit (between 0.5 and 5 percent of the maximum price), other participation requirements, and the date of the auction. Suppliers can then prepare a formal application, consisting of two parts. 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 from the platform before the auction. The purchasing bureaucrat directs the commission's review to deny applications from suppliers that are not accredited, cannot

¹¹The Soviet Union operated a centralized bureaucracy. Since 1991, the Russian bureaucracy has become very decentralized (see e.g. [Enikolopov & Zhuravskaya, 2007](#)).

pay the security deposit, or whose proposal does not comply with the requested item specifications.¹² If only one supplier is approved to participate, the auction is declared “not held” and a contract is drawn up with that supplier at the maximum price. This is relatively common, occurring in 1.4 million cases (22 percent of purchases). 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 participant number and remain anonymous. All participants log in to the platform and participate in a descending, open-outcry auction. When a participant enters a bid lower than the current winning bid, the bid amount, time, and participant number are displayed to all auction participants. The auction continues until ten minutes pass without a lower qualifying bid.

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 materials to make sure the suppliers’ accreditations, licenses, names, registration and tax ID numbers are correct. Among those who are approved, the contract is signed with the bidder who submitted the lowest bid.

2.2 The role of bureaucrats and organizations in procurement

The labor market for Russian procurement officers very much resembles that of private sector jobs. In particular, individuals interested in working in public procurement 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 entity that we will refer to as an *organization*. The organization 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—we refer to these individuals as *bureaucrats*. Together, the organization and bureaucrat (the *procurers*) are tasked with acquiring the good

¹²The 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.

¹³Examples 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.

the organization requires 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 appropriate 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—from working for one organization (or external procurement agency) to another. The other is that external bureaucrats may conduct purchases with multiple organizations, and a given organization may work with multiple external bureaucrats. On average, bureaucrats in our data are observed working with 5.2 different organizations, and organizations with 4.8 different bureaucrats. This high degree of churn is a powerful source of variation for this paper's empirical exercise.¹⁵

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 organization and bureaucrat then together designate the commission to oversee the auction process. The bureaucrat manages all consultations with specialists, collects the information needed to design the tender, and works with the committee to conduct the first stage review, the auction itself, 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 in purchases conducted before 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 bid-

¹⁴Each regional authority sets rules dictating the type of bureaucrat used for each type of purchase, as defined by the maximum price of the contract and the nature of the item. External procurement agencies can be organized by a given authority (for example an education or health ministry), at the federal, regional, or municipal level. Part of the motivation for creating such agencies was to allow different organizations purchasing similar goods to join forces and achieve lower per-unit 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 below.

¹⁵Our 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](#)).

der 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.

Each year from 2011 to 2014 a list of good categories for which a preference for domestic goods was to apply was drawn up.¹⁶ 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). Preferred goods spanned many categories, including automobiles, clocks, various food products, medical equipment, pharmaceuticals, textiles and furs (see Table OA.3). Procurers filing requests for goods on the list were required to publicly inform potential suppliers that the preference would be applied.

Our analysis of the role of the bureaucratic apparatus in driving variation in procurement performance in Section 4 restricts attention to the standard policy regime without preferences. In Section 5 we analyze the impact of the preference regime.

2.4 Building a dataset of comparable procurement purchases

Since 2011, a centralized procurement website (<http://zakupki.gov.ru/>) has provided information to the public and suppliers 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, covering 6.5 million auction announcements for the purchase of 21 million items. Purchases of services and works contacts are highly idiosyncratic, so we remove them from our sample, resulting in a sample of 15 million purchases of relatively homogeneous goods. Table 1 summarizes our procurement data.

To make comparisons of procurement performance across buyers, we must hold constant the precise good being purchased. A great deal of previous research in economics has faced this challenge, but typically achieve within-category homogeneity at the cost of losing generality.¹⁷ 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

¹⁶Preferences 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.

¹⁷Broadly, 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).

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 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 OA.1.¹⁸

To complement this approach, we collect additional data on purchases of pharmaceuticals, a 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.¹⁹ The pharmaceuticals subsample is summarized in column (4) of Table 1.

Finally, 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.

2.5 Corruption

Both public procurement and Russia are associated with widespread corruption (Transparency International, 2016; Szakonyi, 2018). Corruption in procurement can affect the quality of the goods being purchased. However, since our performance measure—the price paid conditional on the good being purchased—carefully controls for the precise good being purchased, it will not be affected by the presence of this type of corruption.²⁰

The quality-adjusted price paid is an attractive measure of performance in the potential

¹⁸Online Appendix OA.1 also analyzes the sensitivity of our main findings to the choices made when developing our text analysis methodology. As Figure OA.2 and Tables OA.9 and OA.10 show, the findings are remarkably robust.

¹⁹We 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.

²⁰Note also that if auction winners were not to sign the contract—a rare occurrence accounting for under one percent of purchases (see Figure OA.1)—any relevant consequences will be captured by our effectiveness measures since the outcome we focus on when estimating procurer effectiveness is the price ultimately *paid* for the item. Non-delivery of goods is also very rare: our contract execution dataset is unusual in that it includes information on whether the organization paying for the items signed for delivery, and less than one percent of the auctions in our sample suffered from “bad execution”.

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, the quality-adjusted prices a government pays for its inputs is 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.²¹ However, the above arguments hold irrespective of whether high prices are due to corruption or to “intrinsic” ineffectiveness, and so in the theoretical framework and the empirical analysis below, we will remain largely agnostic about their relative contributions to qualityadjusted price differences.²²

3 A Simple 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 4. 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 5.2.

3.1 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: $\log(\bar{\theta}) = \mathbf{X}'\beta + \alpha_\theta + \psi_\theta$. \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

²¹Such 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.

²²In Sub-section 4.5 we provide some evidence that corruption is likely not the primary driver of variation in bureaucratic effectiveness in Russia. This is consistent with Bandiera *et al.* (2009), who find that 83 percent of waste in Italian public procurement purchases is due to low bureaucratic ability rather than corruption.

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 length of time allowed to prepare bids, the clarity of the tender documents, bribes to be 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.²³ 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 ($\delta_F < \delta_L$) but face higher participation costs: $c_i = \frac{\bar{\theta}}{1+\delta_i} - \frac{\bar{\theta}}{1+\delta_L} \sqrt{1-\alpha_c-\psi_c}$.²⁴ 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 $\bar{\theta}$.²⁵ 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](#)).

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 [OA.2](#). 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*

²³Note 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 whether to enter as in [Gentry & Li \(2014\)](#). This significantly complicates the analysis, but the qualitative conclusions are the same. Such a model is available from the authors upon request.

²⁴This 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 .

²⁵A more realistic assumption would 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.

log prices are

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

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 $\alpha_\theta, \psi_\theta$ 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 the costs imposed by bureaucrats ($\tilde{\alpha}$) and organizations ($\tilde{\psi}$) managing the procurement process, and forms the basis of our empirical approach.

3.2 Policy change with heterogeneous state effectiveness: bid preferences

We now study the impact of introducing bid preferences favoring the local bidder L .²⁶ 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. 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 it very unlikely F can win the auction. Formally, we focus on the case when $1 + \frac{\delta_L}{\delta_F + \delta_L} (1 - \gamma^{1+\delta_F}) < \gamma^{-\delta_F} < 1 - \log(\gamma^{\delta_L})$.²⁷ In this case, introducing bid preferences has heterogeneous effects depending on the effectiveness of the bureaucracy that we summarize in the following proposition:

²⁶In our empirical application the bid preferences favor locally manufactured goods not local bidders, but in the model we will treat the identity of the firm as a shorthand for the origin of the products being offered.

²⁷When γ is below this interval, F drops out for sure and prices go up. When γ is above this range, participation always decreases and prices always increase.

Proposition 2. When $1 + \frac{\delta_L}{\delta_F + \delta_L} (1 - \gamma^{1+\delta_F}) < \gamma^{-\delta_F} < 1 - \log(\gamma^{\delta_L})$, 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

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

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 5.2

4 How Important is a Good Bureaucracy?

In this section we estimate the extent to which differences in 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.

4.1 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 price paid in a purchase. We use an event study analysis tracking prices paid by organizations that switch bureaucrats. This happens frequently in Russia. As detailed in Table [OA.5](#), 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. We classify the two bureaucrats involved in the event into effectiveness quartiles 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), akin to [Card et al. \(2013\)](#). In the figure, the horizontal axis displays event time, i.e. purchase weeks. The vertical axis displays the average quality-adjusted prices paid in a given week.²⁸

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. This suggests that drift in effectiveness and switches are uncorrelated. Third, we do not see a systematic dip or spike in performance just before a bureaucrat switch, indicating that switches are

²⁸We 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 explained in more detail in the next sub-section). Table [OA.5](#) 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 in a given switch, 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.

not driven by temporary improvements or deteriorations in performance. Fourth, the price changes associated with switching bureaucrats appear symmetric: organizations switching from a bureaucrat in the best quartile of average prices to a bureaucrat in the worst quartile experience a price *increase* of similar magnitude to those switching in the other direction. In Online Appendix OA.4 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 that arise from organizations switching bureaucrats and vice versa in Russian public procurement can be used to estimate specific procurers' causal impact on procurement performance, and that this impact is large.²⁹

4.2 Variance decomposition method

We now aggregate the causal effects of specific bureaucrats and organizations documented in Sub-section 4.1 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 specific bureaucrats or organizations on prices paid. We use these predictions to examine the mechanisms through which procurers affect prices in Sub-section 4.5 and how bureaucratic effectiveness impacts the way policy rules map into public sector output in Section 5.

We model the price paid for an item i procured by an organization j and a bureaucrat $b(i,j)$ as a function of a vector of item attributes \mathbf{X}_i , a price premium that is 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 3 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

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

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.³⁰

²⁹We also construct analogous event study figures for organizations and bureaucrats switching from purchasing one type of *good* to another. The results are presented in Figure OA.4. Each event study shows the same general patterns as in Figure 1.

³⁰By lot size we mean the maximum allowable price for all the items to be purchased in a given auction. We divide the maximum allowable price into bins so as to allow our estimates of procurer effectiveness to capture the impact on prices of the procurers' choice of the exact maximum price posted. The interactions help address, for example, concerns that systematic variation in the average prices of different types of

Identifying the bureaucrat and organization premia is made possible by the switches we documented in Sub-section 4.1. As Abowd *et al.* (2002) show, individual and organization effects are only identified *within* sets of organizations connected by individuals moving between them.³¹ However, such switches do not connect all bureaucrats and organizations that conduct procurement in Russia. Our data contain 984 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 bureaucrat-organization pair that only ever occurs 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. 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 more likely to be federal or regional, and less likely to be in internal affairs or agriculture, but their purchases are of similar size and quantity to those in the full sample, reassuring us that the sample we use for analysis is fairly representative.³²

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:

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

In Online Appendix OA.3, we show that while the $\tilde{\alpha}$ s and $\tilde{\psi}$ s in equation (2) are not identified, the α s, ψ s and γ s in equation (3) are. These are related to the underlying bureau-

goods across space, in combination with differences across procurers in the items purchased, confound our estimates of bureaucrat and organization effectiveness. Russian regions are highly heterogeneous (Enikolopov & Zhuravskaya, 2007; Acemoglu *et al.*, 2011; Yakovlev & Zhuravskaya, 2014). Hereafter we refer to the goods categories constructed using the method described in Sub-section 2.4 as “goods”.

³¹More 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 dimensions—the bureaucrat and the organization effects.

³²In Table OA.7 we show that our results are robust to using only the largest set of connected organizations. Table OA.6 compares the Analysis Sample to its largest connected set.

cocrat and organization effects as follows: $\alpha_b = \tilde{\alpha}_b - \bar{\alpha}_{s(b)}$, $\psi_j = \tilde{\psi}_j - \bar{\psi}_{s(j)}$, and $\gamma_{s(b,j)} = \bar{\alpha}_{s(b,j)} + \bar{\psi}_{s(b,j)}$, where $\bar{\alpha}_{s(b)}$ is the mean bureaucrat effect in the connected set containing bureaucrat b , and similarly $\bar{\psi}_{s(j)}$ is the mean organization effect in organization j 's connected set.³³

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

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

all of which can be identified. Since $\text{Var}\left(\alpha_{b(i,j)}\right)$ and $\text{Var}(\psi_j)$ are variances within connected sets, they are lower bounds on the true variances of bureaucrat and organization effects.³⁴ However, we can combine our estimates to capture the variance in prices that is attributable to the bureaucrats and the organizations *jointly* using the law of total variance: $\text{Var}(\tilde{\alpha}_b + \tilde{\psi}_j) = \text{Var}(\alpha_b + \psi_j) + \text{Var}\left(\gamma_{s(b,j)}\right)$.³⁵

We can obtain unbiased estimates of bureaucrat and organization effects using OLS under the assumption that the residuals ε_i 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 4.1, we do not see any evidence of such pre-trends.

Second, equation (3) assumes that prices are log-linear in the bureaucrat and organization effects—an assumption about the degree of complementarity between the bureaucrat and the organization working on a purchase. If the model is misspecified, then the omitted complementary terms are a component of the residuals in (3).³⁶ 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 they require. Under such sorting, estimates from (3) would recover a mixture of the true effect and the average complementarity of bureaucrat-organization matches.

However, this sorting would imply that organizations switching from bureaucrats

³³Faced with this issue, previous work on private sector workers and firms 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. \(2016\)](#) who study the largest male and female connected sets in Portuguese data.

³⁴Formally, $\text{Var}(\tilde{\alpha}_b) \equiv \mathbb{E}[\text{Var}(\tilde{\alpha}_b|s(b))] + \text{Var}(\mathbb{E}[\tilde{\alpha}_b|s(b)]) = \text{Var}(\alpha_b) + \text{Var}(\mathbb{E}[\tilde{\alpha}_b|s(b)]) \geq \text{Var}(\alpha_b)$. Similarly, $\text{Var}(\tilde{\psi}_j) = \text{Var}(\psi_j) + \text{Var}(\mathbb{E}[\tilde{\psi}_j|s(j)]) \geq \text{Var}(\psi_j)$.

³⁵ $\text{Var}(\tilde{\alpha}_b + \tilde{\psi}_j) \equiv \mathbb{E}[\text{Var}(\tilde{\alpha}_b + \tilde{\psi}_j|s(b,j))] + \text{Var}(\mathbb{E}[\tilde{\alpha}_b + \tilde{\psi}_j|s(b,j)]) = \text{Var}(\alpha_b + \psi_j) + \text{Var}\left(\gamma_{s(b,j)}\right)$

³⁶Note that our identifying assumption does not rule out high effectiveness bureaucrats and organizations matching with each other.

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 1.³⁷ The striking symmetry of the event study evidence indicates that omitted complementarities are unlikely to bias our estimates. Online Appendix 4.4 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 *a priori* be confident 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 of bureaucrats and organizations 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.

We address these sampling error issues in three ways. First, we bootstrap to estimate standard errors for our variance decomposition.³⁸ 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: $\widehat{\text{Var}}^{SS}(\alpha_b) = \text{Cov}(\hat{\alpha}_b^1, \hat{\alpha}_b^2)$, $\widehat{\text{Var}}^{SS}(\psi_j) = \text{Cov}(\hat{\psi}_j^1, \hat{\psi}_j^2)$, $\widehat{\text{Var}}^{SS}(\gamma_s) = \text{Cov}(\hat{\gamma}_s^1, \hat{\gamma}_s^2)$, and

³⁷If anything, Figure 1 showed slightly *smaller* price decreases when organizations switch to lower average-price bureaucrats than when organizations switch to higher average-price bureaucrats.

³⁸We construct partial residuals $\epsilon_i = p_i - \mathbf{X}_i\beta$ 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.

$$\widehat{\text{Var}}^{SS}(\alpha_b + \psi_j) = \text{Cov}(\hat{\alpha}_b^1 + \hat{\psi}_j^1, \hat{\alpha}_b^2 + \hat{\psi}_j^2).$$

Third, we adopt two shrinkage approaches to create predictions of each bureaucrat's and each organization's 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](#)).³⁹ 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.⁴⁰ We label this method "covariance shrinkage". It yields our preferred estimates of the price variance decomposition in equation (4).⁴¹

4.3 Results

Table 2 shows the results of implementing 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 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). Rows 1–3 show standard deviations across bureaucrats, organizations, and connected sets, while rows 4–8 show the decomposition across items purchased.

³⁹Formally, we find $\lambda_b = \arg\min_{\lambda} \mathbb{E}[\alpha_b - \lambda\hat{\alpha}_b] = \sigma_\alpha^2 / (\sigma_\alpha^2 + \sigma_{\mu_b}^2)$, 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 $\hat{\psi}_j^{Sh} = \hat{\lambda}_j\hat{\psi}_j$.

⁴⁰Formally, 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 $\hat{\theta}$ and the matrix of weights by Λ , the objective is $\min_{\Lambda} \mathbb{E}[(\theta - \Lambda\hat{\theta})'(\theta - \Lambda\hat{\theta})]$, which has solution $\Lambda^* = \mathbb{E}[\theta\hat{\theta}'](\mathbb{E}[\hat{\theta}\hat{\theta}'])^{-1}$. Replacing the expectations with their sample analogues yields the shrinkage matrix $\hat{\Lambda}^* = \text{diag}(\hat{\sigma}_\alpha^2, \hat{\sigma}_\psi^2)(\text{diag}(\hat{\sigma}_\alpha^2, \hat{\sigma}_\psi^2) + \Sigma)^{-1}$, where Σ is the covariance matrix of the bootstrap estimates and $\text{diag}(\hat{\sigma}_\alpha^2, \hat{\sigma}_\psi^2)$ is the diagonal matrix with $\hat{\sigma}_\alpha^2$ in entries corresponding to entries for bureaucrats in θ and $\hat{\sigma}_\psi^2$ in entries corresponding to organizations.

⁴¹We thus use "covariance shrunk" estimates in our analysis of the determinants of bureaucratic capacity in Sub-section 4.5 and the analysis of the effects of procurement policy changes in Section 5. 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.

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 1.283. Compared to this, the bureaucrat fixed effects have a standard deviation of 0.747 and the organization fixed effects' standard deviation is 0.827. 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.⁴² However, our covariance shrinkage approach yields a more plausible estimate of the correlation of 0.33.⁴³ As a result, the covariance shrunk estimates of share of the variation in performance explained by bureaucrats and organizations—24 and 26 percent respectively—represent our preferred estimates of the importance of bureaucrats and organizations for state effectiveness in procurement.

Third, the combined importance of bureaucrats and organizations for policy performance is large. As shown in Sub-section 4.2, the estimates of the variation in bureaucrat and organization effects are to be interpreted as estimates of *within*-connected set variation, but these can be combined with the variation in the connected set intercepts to yield an estimate of the variation in the combined impact of bureaucrats and organizations on effectiveness Russia-wide. These estimates are shown in row 8 of Table 2. The estimates are remarkably consistent across the four methods, ranging from 0.63 for the raw fixed effects estimates down to our preferred estimate of 0.512, or 40 percent of the standard deviation of log unit prices, for the covariance-shrunk estimates. 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

⁴²The 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 4.2) 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)).

⁴³Of course, such assortative matching does not violate the no-sorting-on-match-effects assumption discussed in sub-sections 4.1 and 4.2.

and/or organizations with low effectiveness are improved, for example through changes in recruiting, training of existing bureaucrats, or improved organizational management. Increasing the effectiveness of the lowest quartile of bureaucrats to the 75th percentile would save the Russian government 3.3 percent of annual procurement expenses. Moving all bureaucrats *and* organizations below 25th percentile-effectiveness to 75th percentile-effectiveness would save the government 10.7 percent of procurement expenditures.⁴⁴ Annual procurement expenses are USD 86 billion, so this implies savings of USD 13 billion each year, or 0.9 percent of non-resource GDP (see Table OA.4)—roughly one fifth, for example, of the total amount spent on health care in 2013 and 2014.⁴⁵

4.4 Robustness: Like-for-like good comparisons

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 at minimizing the price paid for each specific good it procures. However, if our goods classification based on the contract texts is inaccurate, our estimates will conflate the true effects on prices with differences across bureaucrats and organizations in the products that they are buying. To probe this concern, we perform three robustness checks.

First, we show that our findings are remarkably 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 2.4 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, 46 percent is attributable to the combination of bureaucrats and organizations. 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, Column (6) of table 4 shows that the results from our variance decomposition exercise are also essentially unaffected if we restrict the sample to items the text-based classification method is confidently able to assign a 10-digit Harmonized-System product

⁴⁴Figure OA.5 shows how these counterfactuals affect the distributions of effectiveness.

⁴⁵Online Appendix OA.6 compares these magnitudes to other studies of individuals’ and organizations’ effects on output in other settings.

code to.⁴⁶ Third, our results are robust to focusing on more homogeneous subsets of goods in our full sample. We split the sample into quintiles of good homogeneity as defined by the commonly-used measure of the scope for quality differentiation developed by Sutton (1998). We then reestimate (3) on successive subsamples. Table 4 show the results. Column (5) includes all observations for which the Sutton (1998) measure is available.⁴⁷ 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 decreases with good homogeneity. However, the estimated share of the variance explained by bureaucrats and organizations remains very similar across the columns. In Table OA.8 we repeat this exercise using an alternative measure of scope for quality differentiation developed by Khandelwal (2010) and find the same result.⁴⁸

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

4.5 What do effective bureaucracies do differently?

In this sub-section we show evidence on what distinguishes effective bureaucracies. We leverage the richness of our procurement data, which contain detailed information on the evolution of each of the 6.5 million procurement processes in the sample, from the initial request document, through the auction itself, to the final contract signed with the supplier. We complement these data with information about participating firms from the Bureau Van Dijk's *Ruslana* database, and with data on corruption and other measures of institutions across regions.⁴⁹ We then use the resulting dataset to investigate which features of the procurement process, the firms participating in it, and the procurers themselves co-vary with the estimated effectiveness of the implementing bureaucrats and organizations. The resulting dataset contains 160 potential explanatory variables, listed in Table OA.12.

To avoid overfitting and for the sake of parsimony, we use a LASSO procedure to select 50 predictor variables and 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

⁴⁶The algorithm developed in Step 2 of the procedure outlined in Sub-section 2.4 and Online Appendix OA.1 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.

⁴⁷We 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.

⁴⁸Another possibility is that organizations endogenously respond to the effectiveness of the bureaucrats available to them by purchasing more/fewer, or different types of, goods. This would lead us to underestimate the true variance in procurer effectiveness and its consequences.

⁴⁹The latter come from Schulze *et al.* (2016) and the ICSID Russian Regions database.

the organization effects).⁵⁰ Figures 2 and 3 show the results. The left panels show regression coefficients from a series of bivariate regressions of the bureaucrat effect (in Figure 2) and the organization effect (in Figure 3) on each of the selected observables. The right panels show the coefficients from the multivariate regression of the procurer effects on all of the selected variables. 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 on prices. Of course, the relationships displayed in Figures 2 and 3 need not be causal, in part because we do not observe everything different procurers do differently.

Seven key findings emerge. First, effective bureaucrats both attract more applicants and allow more of the applicants to participate in their auctions. To provide an example, Figure OA.8 shows that the bureaucrat overseeing a request for winter boots by a Saratov orphanage 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, effective bureaucrats set lower reservation prices. A common procedure is to apply officially standardized algorithms to market research on the average price paid for a given good or service. Effective bureaucrats are able to set lower reservation prices by soliciting accurate commercial information from trusted suppliers and established market players.⁵¹ These two findings figures 2 and 3 resonate with our theoretical framework in Section 3, which predicts that some procurers pay higher prices than others because they impose high costs of fulfilling government contracts and high participation costs, consequently attracting fewer suppliers to auctions.

Third, more experienced bureaucrats—for example those who have run more auctions in the past—are more effective, consistent with them having a larger network of contacts with suppliers to draw on. Fourth, “in-house” bureaucrats who are employees of the organization acquiring the item are less effective. Fifth, effective *organizations* also attract more applicants and set lower reservation prices in their auctions. Additionally, certain types of organizations—for example municipal ones and those located in less remote areas—are more effective than other types of organizations. Sixth, figures 2 and 3 also reveal important differences in the types of bidders effective buyers are able to attract. Effective buyers

⁵⁰To 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 OA.9 and OA.10 show that the results are not sensitive to placing more weight on the Ridge regression.

⁵¹Research has shown the use of flawed market information to be one of the main ways ineffective bureaucrats drive up the price of Russian procurement (see Sapozhkov, Oleg. “Krivye Putyi Goszakazchikov.” *Kommersant*, April 12, 2019).

tend to buy from high-profit firms or directly from manufacturers, and avoid “fly-by-night” entrepreneurs who charge higher prices. In 2019, the State Duma deputies acknowledged the severity of this problem.⁵²

Finally, there are some notable variables among the 110 that are *not* selected by the LASSO. In particular, the wide range of regional measures of corruption have very weak predictive power. In Figure 2 we see that two such measures—the Regional Number of Corruption Cases and the Regional Number of Corruption Convictions—do predict respectively higher and lower bureaucratic effectiveness⁵³, but the magnitude of the estimated coefficients on these variables is very small. 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 favorable types of firms.

5 Policy Design with a Heterogeneous Bureaucracy

Section 4 documents the large variation in procurement performance under a constant policy regime attributable to heterogeneity in bureaucratic agents’ effectiveness. This naturally raises the question of how introducing a different policy regime would affect the performance of the average bureaucrat, and of more versus less effective bureaucrats. This question is important for all enterprises since productivity can be enhanced either through workers and organizational units themselves—by modifying hiring, training, and incentive practices—or by optimizing the tasks and policies the workforce is directed to implement. However, the magnitude of the potential benefits of designing policy to match the effectiveness of the implementing agents is *especially* important for states since altering human resource practices can be infeasible or costly in the public sector.

In this section, we study the impact of a particular policy change in Russia. We show that the introduction of bid preferences favoring locally manufactured goods led to firms supplying such goods winning more procurement auctions, without significant impacts on prices or overall participation. However, these average treatment effects mask dramatic heterogeneity across “good” versus “bad” procurers, suggesting that there is significant

⁵²The Duma proposed new legislation to require more information from suppliers about their past procurement experience before receiving contract advances (see Filonenko, Valerii. “Firmy-Odnodnevki Otstranyat ot Goszakupki” *Parliamentskaya Gazeta*, April 11, 2019.) “Fly-by-night” entrepreneurs had consistently failed to fulfill contracts on-time and to high standards, forcing organizations to revise contracted prices upward.

⁵³Controlling for corruption convictions, a higher number of cases initiated could indicate better monitoring of bureaucrats by superiors and law enforcement authorities.

scope for tailoring policy design to the effectiveness of the implementing bureaucracy in Russian procurement.

5.1 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 on prices and participation 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. As described in detail in Sub-section 2.3, Russia's policy imposed a bid penalty of 15 percent on foreign-manufactured goods. 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.⁵⁴ We exploit this variation in a generalized difference-in-differences design, estimating

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta \text{Preferred}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt} \quad (5)$$

where y_{igt} is the outcome in purchase i of good g in month t . Preferred_{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 preferred goods has been published. \mathbf{X}_{igt} are the same controls we use in Section 4, 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 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. We also 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:

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

where terms are defined as above, and ListMonth_t is the month closest to month t in which a preference list is published.

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 Preferred_{gt} as a dummy equal to one if good g is on that year's list. Since pharmaceuticals are always on

⁵⁴Preferred goods spanned many categories, including automobiles, clocks, various food products, medical equipment, pharmaceuticals, and textile and furs (see Table OA.3 for the full list).

the list, for pharmaceuticals we instead define Preferred_{gt} as equal to one if the drug is manufactured both in Russia and abroad (since several drugs consumed in Russia are manufactured either only abroad or only domestically).

Panel A of Figure 4 shows the event study coefficients δ_s from estimation of (6) for prices in the Analysis Sample. They are all close to zero and statistically indistinguishable from zero in the months leading up to the publication of the preference list, lending 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 OA.7 shows the evolution of the share of purchases for preferred 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.

Table 5 shows the results of estimating (5). Column (5) shows that the policy achieves its main goal: the good purchased is 14 percent more likely to be domestically manufactured.⁵⁵ However, columns (1)–(4) show that this does not come at the cost of higher prices paid or lower participation. In both samples, the preferences policy decreases the log price achieved by about two percentage points, despite decreasing the average number of bidders per auction slightly, though these estimates are statistically insignificant in the pharmaceuticals sample and only marginally significant in the full sample. The policy's discouragement of foreign manufacturers is offset by a combination of encouragement of local manufacturers and the mechanical decrease in prices paid when the winning bidder supplies foreign manufactured goods.

These findings contrast with the results from studies of similar preference policies in the U.S., which suggest that prices increase when the government introduces bid preferences (see e.g. Marion, 2007; Krasnokutskaya & Seim, 2011; Athey *et al.*, 2013). Our heterogeneity analysis in the next section points towards a possible explanation: we find effects similar to those in the U.S. when preferences are implemented by procurers of high effectiveness, but opposite effects when procurers are ineffective. Since the overall impact in Russia averages over a population of procurers with very different effectiveness than in the U.S. (where procurers are likely to be more effective on average), these findings can be reconciled.

⁵⁵We only observe goods' country of origin for pharmaceuticals, so we cannot assess the impact of the policy on the likelihood a domestically produced good wins in the full sample. In Column (5) we restrict the sample to purchases in which an auction takes place in order to be consistent with Column (5) 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).

5.2 Bureaucratic performance heterogeneity under different policy regimes

We are interested in whether the policy change also impacted how the bureaucrats and organizations who implement policy affect procurement performance. Proposition 2 in Section 3 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. The proposition implies that bid preferences will lead to a compression of the variation in performance driven by bureaucrats and organizations. To test for this 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 4.

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

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \theta_b \hat{\alpha}_b + \theta_j \hat{\psi}_j + \delta \text{Pref}_{gt} \text{Active}_t + \gamma_b \text{Pref}_{gt} \hat{\alpha}_b + \gamma_j \text{Pref}_{gt} \hat{\psi}_j + \eta_b \text{Active}_t \hat{\alpha}_b + \eta_j \text{Active}_t \hat{\psi}_j + \pi_b \text{Pref}_{gt} \text{Active}_t \hat{\alpha}_b + \pi_j \text{Pref}_{gt} \text{Active}_t \hat{\psi}_j + \varepsilon_{igt} \quad (7)$$

The parameters of interest are π_b , the heterogeneity of the treatment effect by bureaucrat effectiveness, and π_j , the heterogeneity of the treatment effect by organization effectiveness. We also estimate less parametric versions of (7) by including separate triple-interaction terms for each decile of bureaucrat effectiveness $\hat{\alpha}_b$ and organization effectiveness $\hat{\psi}_j$, and extend the event study (6) to estimate effects separately by quartile of bureaucrat- and organization-effectiveness.

Table 6 and Figure 5 show the results.⁵⁶ We see that the small negative average price effect found in Table 5 masks substantial heterogeneity in the impact of bid preferences across bureaucracies. Consistent with the predictions of Proposition 2, prices drop significantly more for bureaucrats who pay higher prices (i.e., who have a higher $\hat{\alpha}_b$) when there are no bid preferences, both in the full sample and the pharmaceuticals sample. Similarly, prices drop more for organizations who pay higher prices (have a higher $\hat{\psi}_j$). However, this latter effect is more muted in the full sample, and not statistically significant in the pharmaceuticals sample.

The estimates in Table 6 illustrate that the degree to which bureaucracies deviate from the Weberian ideal of mechanistic, uniform performance depends not just on heterogeneity in the participants in the bureaucracy (Weber, 1921), but also on the task that the bureaucracy is asked to perform. When a given group of Russian bureaucrats are asked to implement a policy regime featuring bid preferences, the standard deviation of the impact of bureaucrats and organizations on prices falls by 12 percent.

Figure 5 confirms the findings in Table 6 graphically. Panel B shows a clear pattern of

⁵⁶In Table 6 we use the covariance shrunk estimates of the bureaucrat and organization effects. Table OA.11 uses the raw fixed effect estimates and shows very similar results.

larger price drops for bureaucrats of lower effectiveness. The effects are decreasing throughout the observed range, rather than being concentrated among especially effective or ineffective bureaucrats. Meanwhile, Panel D does not show such a clear pattern for organizations.

The event studies in panels A and C of Figure 5 also help rule out potential confounds like mean reversion or differences in seasonality across different types of bureaucrats and organizations. The figure shows no discernible trends in prices before the introduction of bid preferences and then a marked divergence after the introduction of preferences for high versus low effectiveness bureaucrats, but not for organizations. These patterns provide compelling evidence that the estimates in Table 6 capture the causal differential of interest.

Proposition 2 also predicts differential changes in the number of bidders participating in procurement processes, mirroring the effects on prices. Columns (2) and (4) of Table 6 show that this is indeed what we see. The average number of participants decreases in auctions administered by effective procurers when bid preferences apply, but increases in auctions administered by ineffective procurers.⁵⁷ Finally, Proposition 2 predicts differential impacts of the policy on the likelihood that an auction is won by a bidder offering locally manufactured goods. Column (5) of Table 6 shows that we do indeed see strong heterogeneity in the impact on goods' origin: 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.

Overall, these results suggests 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 3. For organizational effectiveness, the policy design implications of our findings in Table 6 are less pronounced. Our estimates suggest that the bid preference policy saved the government 17.5 percent when it was implemented by the least effective quartile of bureaucrats, but only 0.7 percent when implemented by the most effective quartile of bureaucrats. Similarly, the pharmaceuticals estimates suggest the probability that an auction was won by a local manufacturer increased by 15.9 percent when the policy was administered by the least effective quartile of bureaucrats, but only 10.3 percent when the policy was administered by the most effective quartile of bureaucrats. The results for the most effective bureaucrats in Russia are comparable to results from the U.S. (Marion, 2007; Krasnokutskaya & Seim, 2011; Athey *et al.*, 2013), but the overall effect in Russia also includes a large number of bureaucrats with lower capacity.

⁵⁷In the full sample the difference in the change in the number of bidders for effective and ineffective organizations is not statistically significant.

5.3 Drivers of performance heterogeneity under different policy regimes

We saw in the previous sub-section that the preferences policy alters the relationship between bureaucrats' "types" and their ultimate procurement performance. Understanding if this is because the mapping from bureaucrat type to how procurement processes are carried out changes, or because the mapping from how procurement processes are carried out to the ultimate prices paid changes, is important. Doing so can help us begin to unpack why the relationship between bureaucratic heterogeneity and performance heterogeneity changes under different policy regimes.

To investigate, we turn again to our rich data on procurement processes and take an approach similar to the one we used in Sub-section 4.5 to study the determinants of bureaucrats' and organizations' performance in the baseline policy regime. We estimate a triple difference regression akin to (7):

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \theta_b\hat{\alpha}_b + \theta_j\hat{\psi}_j + \delta \text{Pref}_{gt}\text{Active}_t + \gamma_b \text{Pref}_{gt}\hat{\alpha}_b + \gamma_j \text{Pref}_{gt}\hat{\psi}_j \\ + \eta_b \text{Active}_t\hat{\alpha}_b + \eta_j \text{Pref}_{gt}\hat{\psi}_j + \text{Pref}_{gt}\text{Active}_t \mathbf{Y}_{igt}\pi + \varepsilon_{igt} \quad (8)$$

The triple difference terms in (7) are replaced with interactions between $\text{Pref}_{gt}\text{Active}_t$ and a vector of observables \mathbf{Y}_{igt} . Since our data contains a large number of observables (listed in Table OA.12), the vector \mathbf{Y}_{igt} is chosen by the same regularization procedure used in Sub-section 4.5.⁵⁸ Under the assumption that $\hat{\alpha}_{igt}$ and $\hat{\psi}_{igt}$ capture all the ways that the observables \mathbf{Y}_{igt} affect prices and participation in the policy regime without preferences, the coefficients π capture the effect of the preference regime on the relationship between the observables \mathbf{Y}_{igt} and prices and participation.⁵⁹ The high R^2 of the regressions in Section 4 from which the $\hat{\alpha}$ s and $\hat{\psi}$ s are derived suggests this assumption is reasonable.⁶⁰

If the observables with large π coefficients in (8) are also predictors of bureaucrat type in Figure 2, then the covariates that are important for the effect of the policy are also the ones on which high and low effectiveness bureaucrats differ in a constant policy regime. If so, we expect significant differences in the effect of the preference regime between effective and ineffective bureaucrats since they have different levels of these covariates, even if the

⁵⁸We first run a LASSO procedure with the full set of observables in our data to select the elements of \mathbf{Y}_{igt} . For the selected variables, we run regression (8). As in Sub-section 4.5, we also use an elastic net procedure so that the regularization takes greater account of the correlation between the observables. Figures OA.12 (for prices) and OA.13 (for the number of bidders) show that the results are very robust to how much weight we place on the ridge criterion in the elastic net.

⁵⁹Formally, we require that $\mathbb{E}[y_{igt}|\mathbf{X}_{igt}, \mu_g, \lambda_t, \hat{\alpha}_b, \hat{\psi}_j, \text{Pref}_{gt}\text{Active}_t=0, \mathbf{Y}_{igt}] = \mathbb{E}[y_{igt}|\mathbf{X}_{igt}, \mu_g, \lambda_t, \hat{\alpha}_b, \hat{\psi}_j, \text{Pref}_{gt}\text{Active}_t=0]$, or that \mathbf{Y}_{igt} does not contain information on baseline prices conditional on $\hat{\alpha}_b$ and $\hat{\psi}_{igt}$.

⁶⁰We could relax this by adding in the requisite linear and interaction terms with \mathbf{Y}_{igt} , but at the cost of making the model computationally very hard to estimate.

relevant covariates do not *change* with the policy. On the other hand, if the covariates that matter for the preference policy's impact are very different from those that predict bureaucrats' effectiveness when there are no preferences, it must be the case that the preference regime significantly altered the behavior of bureaucrats in a way that led to changes in the observables that predict the policy's impact.

Figure 6 shows the estimates of the heterogeneity of the effect of the preferences policy on prices.⁶¹ The figure is constructed analogously to figures 2 and 3. We see that attracting a large number of applicants becomes even more important, as does attracting smaller, local firms that sell a more focused range of products, under the policy. Additionally, setting high reservation prices becomes even more damaging to performance; bundling items together becomes more important (possibly because larger, more diverse lots are able to draw in competitive foreign bidders despite the preferences); and actually holding the auction becomes more important. Finally, bureaucrats' overall experience matters less for the impact of the policy, but their experience with the product being purchased becomes more important, under bid preferences.

Turning to how the preferences policy altered the relationship between bureaucrats' types and their procurement performance, two points are worth noting. First, most of the predictors in Figure 6 relating to the type of the bureaucrat, the end-user organization, and the execution of the purchase requests and the auction are also predictors of bureaucrat effectiveness in Figure 2. These are features of the procurement process for which differences in bureaucrats' effectiveness will drive differences in performance under the preference regime, even without bureaucrats changing anything about how they conduct procurement. These predictors are important: they represent half of the observables selected by the regularization procedure. This suggests that policy choices—optimizing the tasks and policies the bureaucracy is directed to implement, taking into account bureaucratic effectiveness—will have important consequences for the productivity of the state, even without changes in the behavior of the policy-implementing agents.

Second, the variables that predict heterogeneity of the effect of bid preferences in Figure 6 *but not* bureaucratic effectiveness in Figure 2 overwhelmingly relate to the type of firms that bid for, and win, procurement contracts. In particular, where in the distribution of potential bidders the bidders and winners lie are important determinants of the effect of bid preferences, but not of bureaucratic effectiveness. This highlights the importance of bidders' participation decisions and how they interact with policy implementers for understanding the impact of procurement policies on performance. In sum, the results described in this sub-

⁶¹Figure OA.11 shows the analogous results for the number of bidders. The findings are very consistent with those for prices, so for brevity we focus here on prices.

section 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 5.2 suggest.

6 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 40 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, permitting 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.

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. 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. Naturally, doing so will involve tradeoffs between these two approaches. We see studying these tradeoffs as a promising direction for future research.

References

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

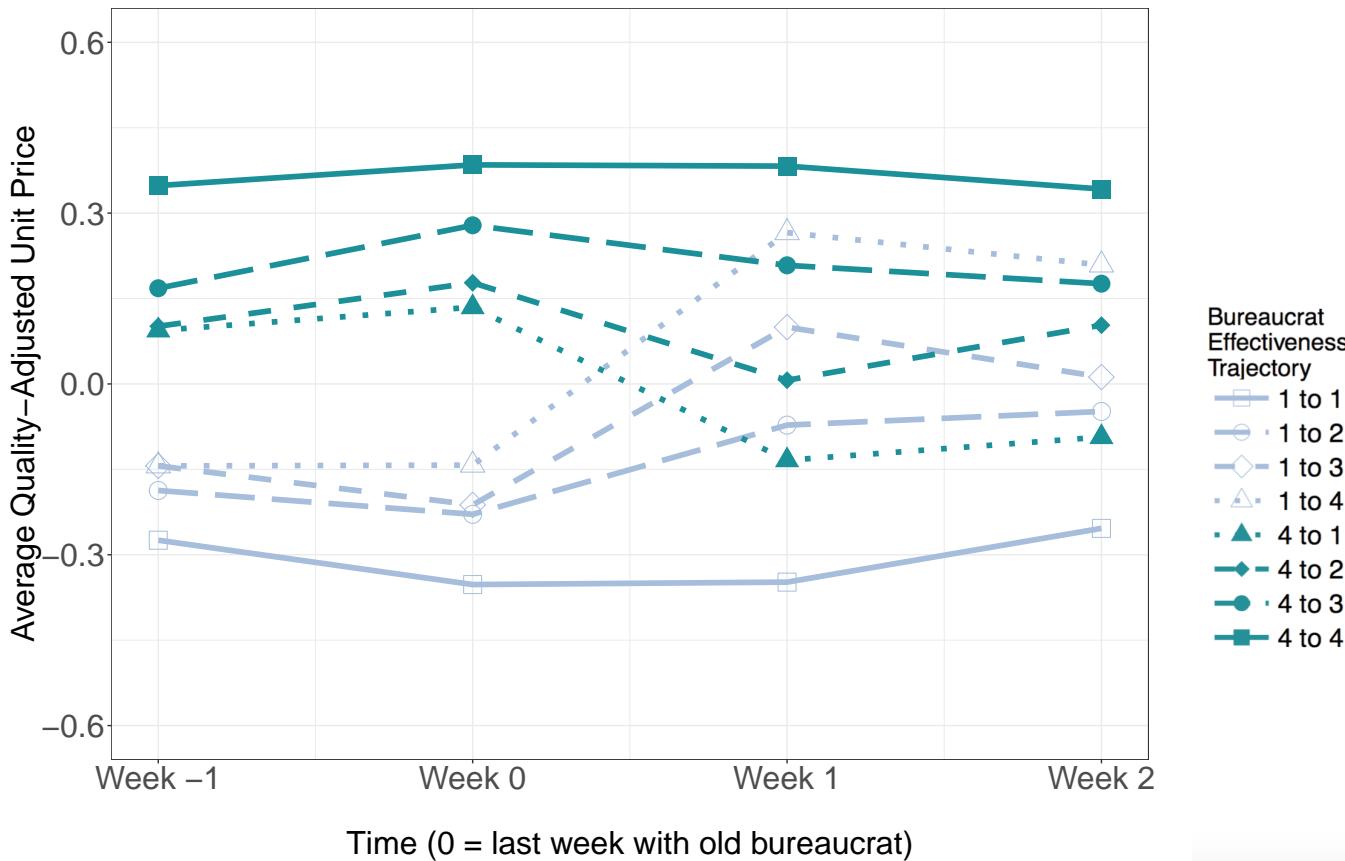
- ABOWD, JOHN M., CREECY, ROBERT H., & KRAMARZ, FRANCIS. 2002. *Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data*. Census Bureau Technical Paper TP-2002-06.
- ACEMOGLU, DARON, HASSAN, TAREK A., & ROBINSON, JAMES A. 2011. Social structure and development: a legacy of the holocaust in Russia. *Quarterly Journal of Economics*.
- 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.
- 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. forthcoming. The Glittering Prizes: Career Incentives and Bureaucrat Performance. *Review of Economic Studies*.
- 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, KLEVÉN, HENRIK, SPINNEWIJN, JO-HANNES, & 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.
- 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.
- 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, NEGTERS, 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.
- 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, SJAHIRIR, 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.

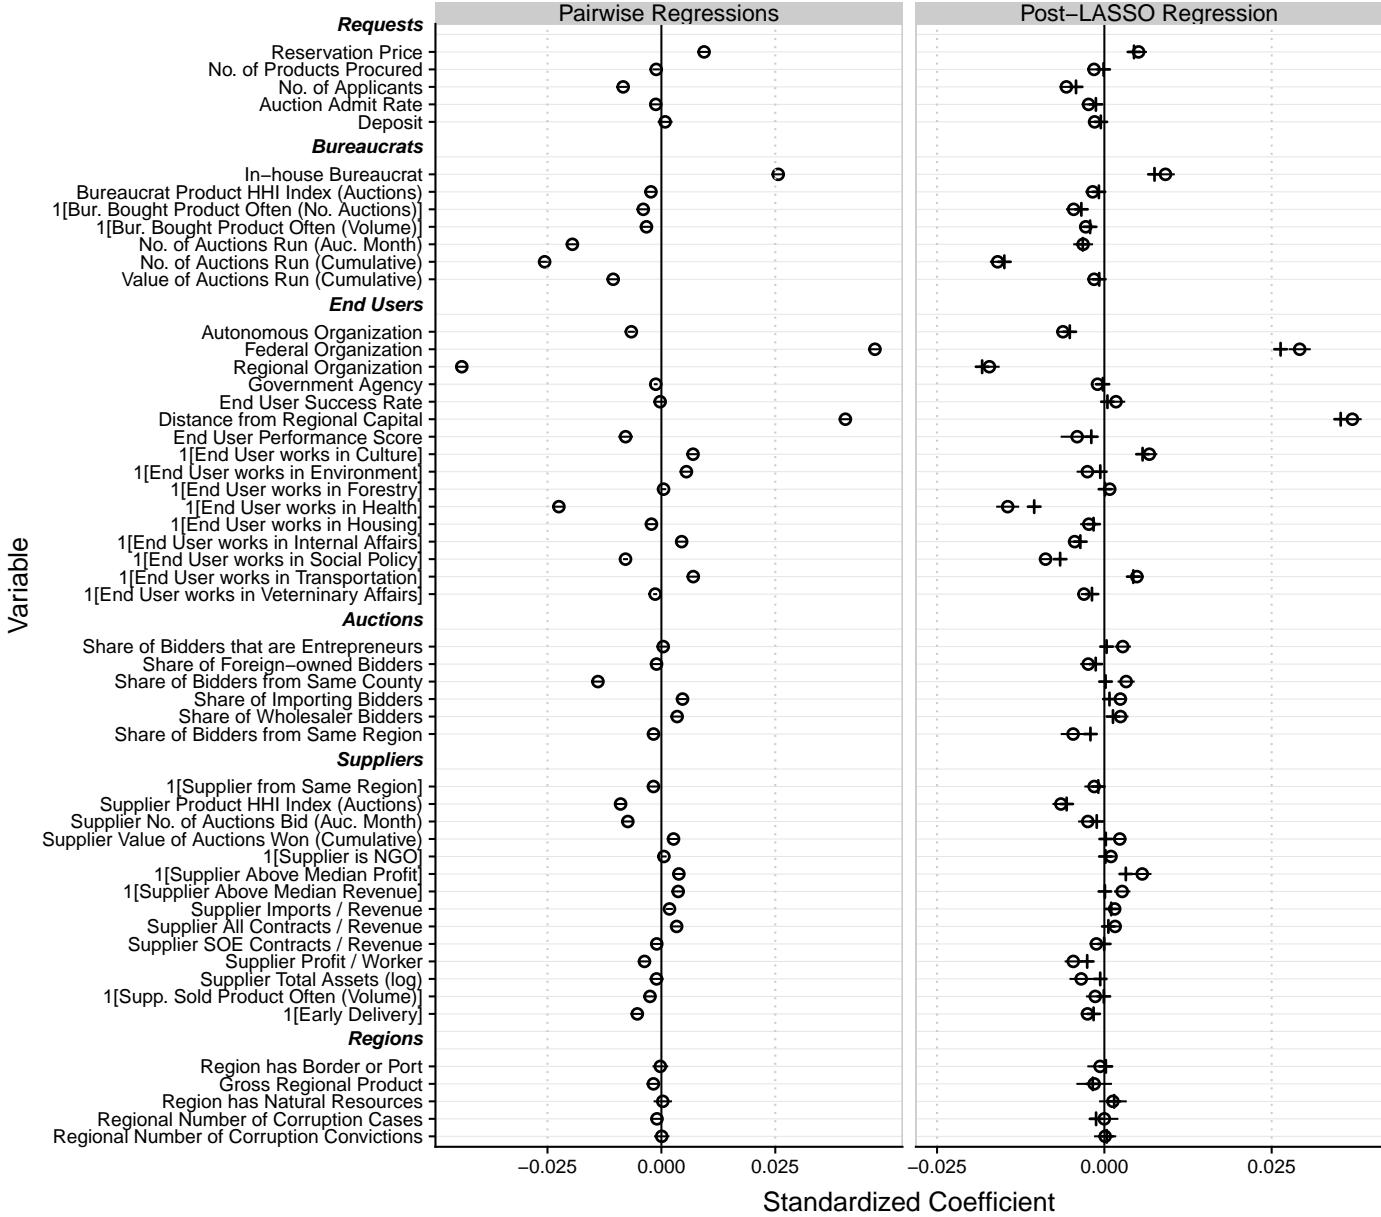
- SYVERSON, CHAD. 2004. Market Structure and Productivity: A Concrete Example. *Journal of Political Economy*.
- SZAKONYI, DAVID. 2018. Businesspeople in Elected Office: Identifying Private Benefits from Firm-Level Returns. *American Political Science Review*, **112**(2), 322–338.
- TRANSPARENCY INTERNATIONAL. 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, 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



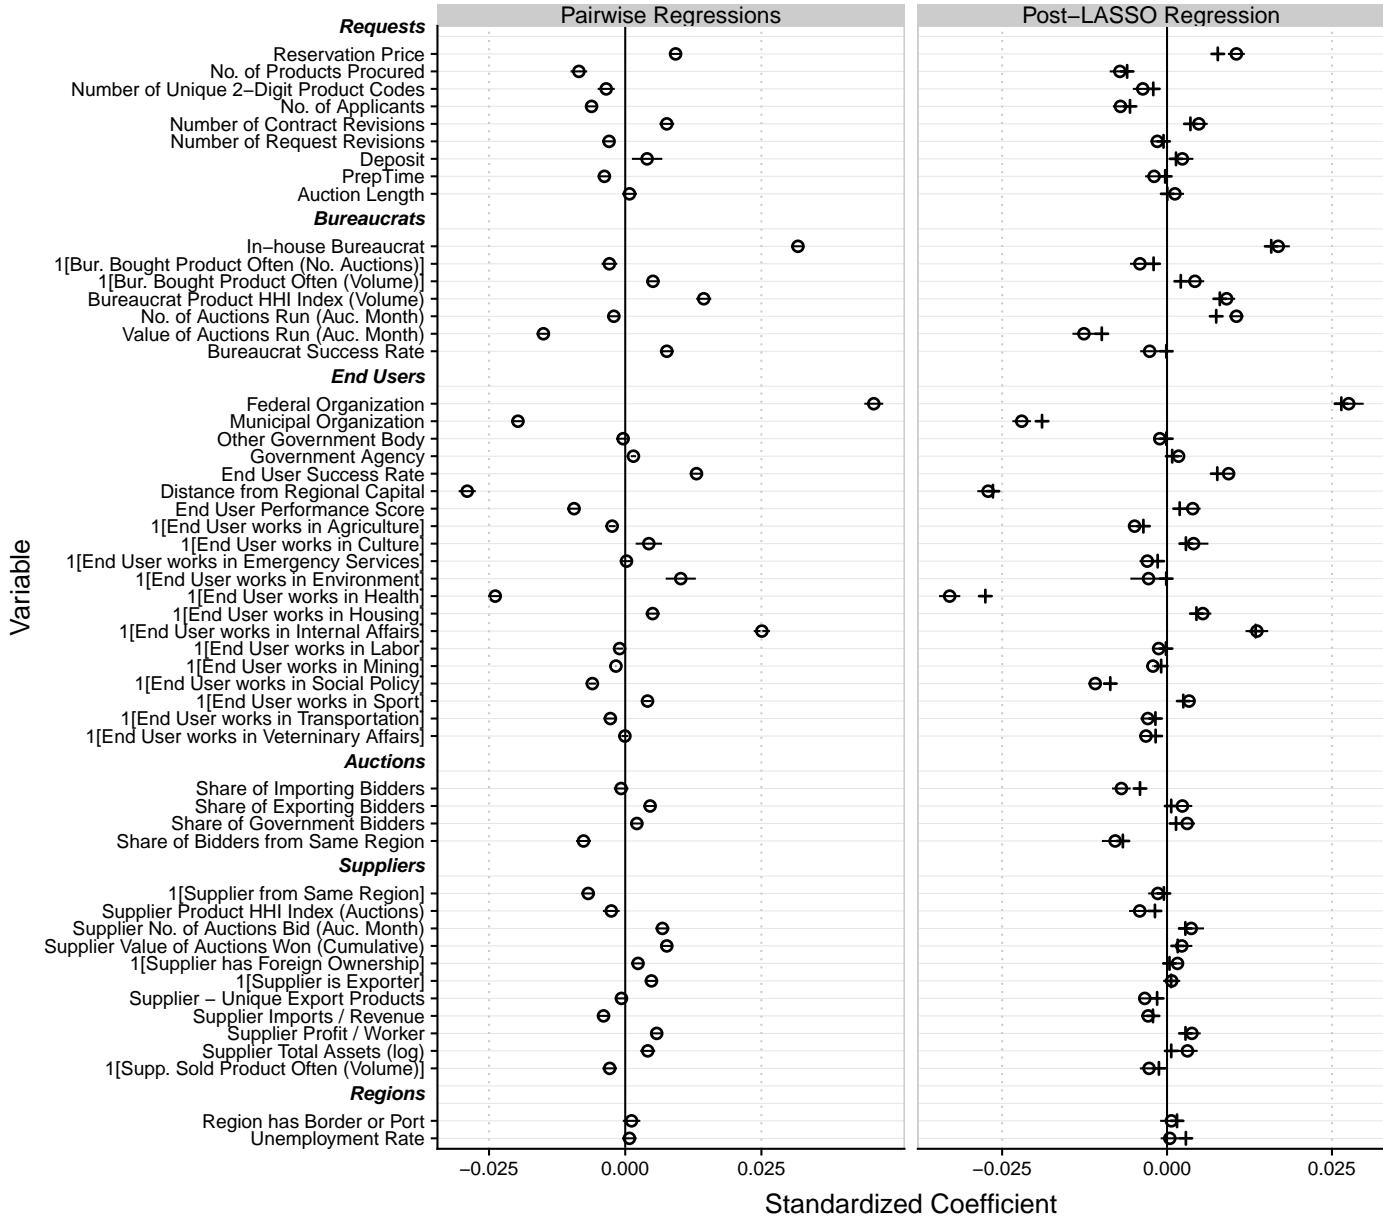
The figure shows time trends in prices around the time that organizations switch which bureaucrat makes purchases on their behalf. The horizontal axis measures fortnight on which bureaucrat-organization pairs work together, with time 0 being the last fortnight on which the organization works with the old bureaucrat just before switch, and time 1 being the first fortnight 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 on two separate fortnights 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: CORRELATES OF BUREAUCRAT EFFECTIVENESS



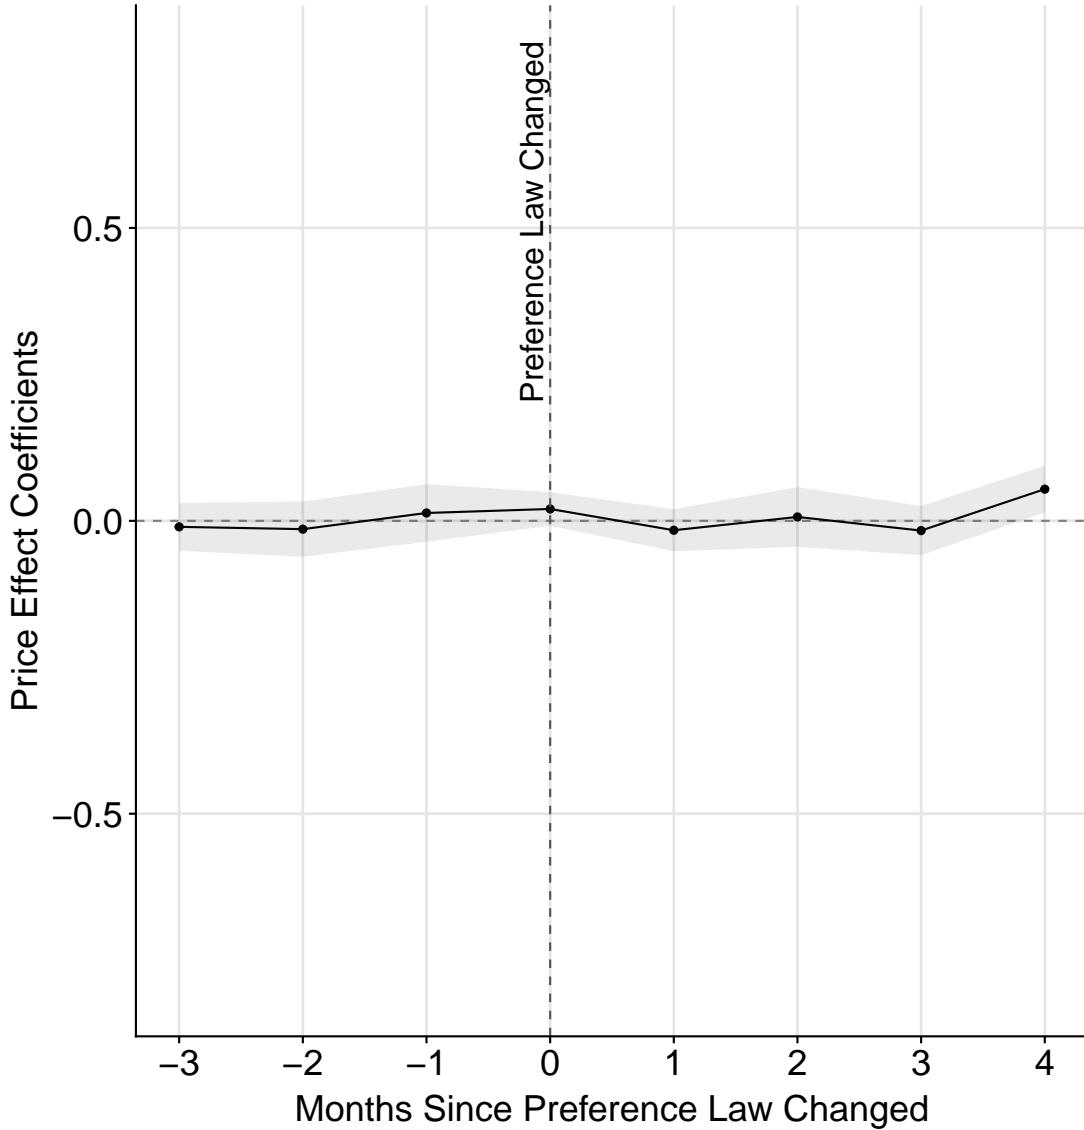
The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (3): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 50 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 3: CORRELATES OF ORGANIZATION EFFECTIVENESS



The figure shows the results of regressions of estimated organization effects $\hat{\alpha}_b$ from estimation of equation (3): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. We use a LASSO procedure to select 50 predictor variables and regress each purchase's covariance-shrunk organization effect on these variables, the purchase's bureaucrat effect, and the controls in (3). The left panels show regression coefficients from a series of bivariate regressions of the organization 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.

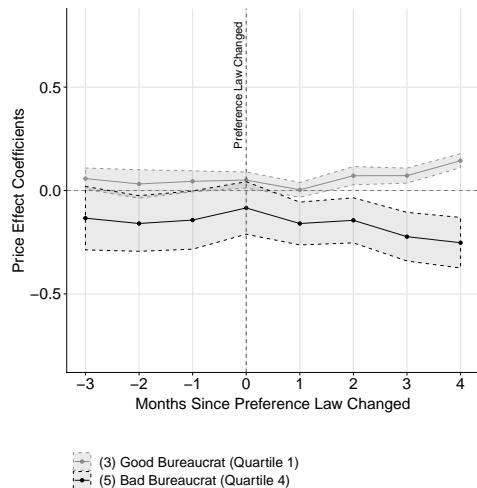
FIGURE 4: GRAPHICAL EVIDENCE ON EFFECT OF BID PREFERENCES



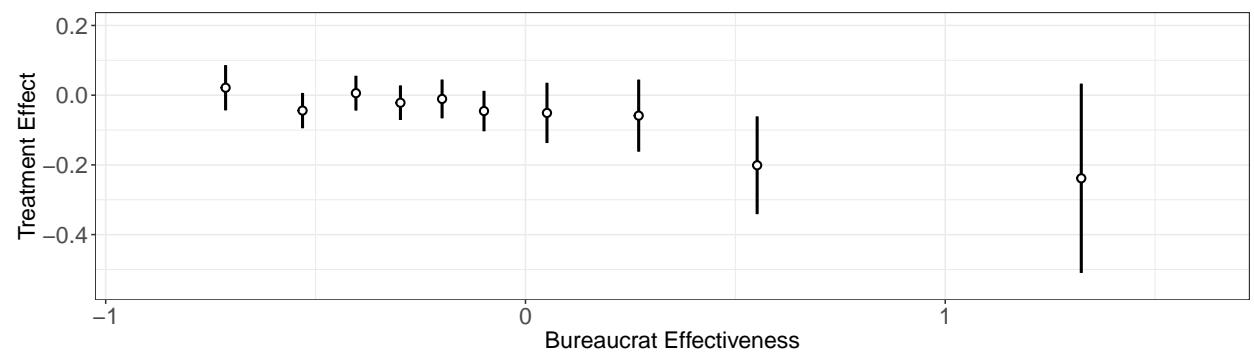
Notes: The figure shows a graphical analysis of the preferences policy over the period of study. The x-axis is measured in the number of months preceding or following the activation of the annual preferences laws in 2011, 2012, 2013, and 2014. The dotted vertical lines indicates when the policy was became active. The y-axis in each plot shows the month-specific coefficients from estimation of equation (6): $p_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \sum_{s=-3}^4 \delta_s \text{Preferred}_{gt} \times \mathbf{1}\{t - \text{ListMonth}_t = s\} + \varepsilon_{igt}$, where p_{igt} is the price paid for purchase i of good g in month t . Preferred_{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 4, 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.

FIGURE 5: HETEROGENEITY OF EFFECT OF BID PREFERENCES BY BUREAUCRAT AND ORGANIZATION EFFECTIVENESS

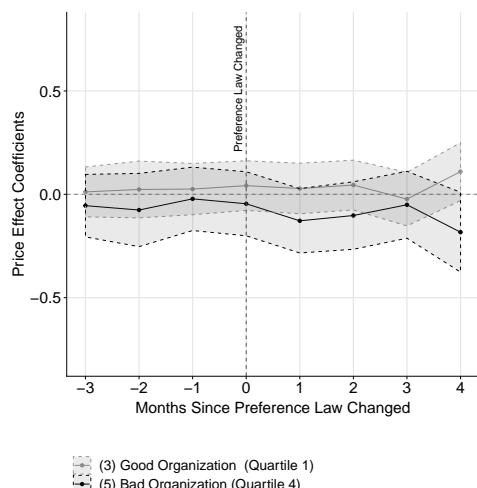
PANEL A: EVENT STUDY BY BUREAUCRAT EFFECTIVENESS



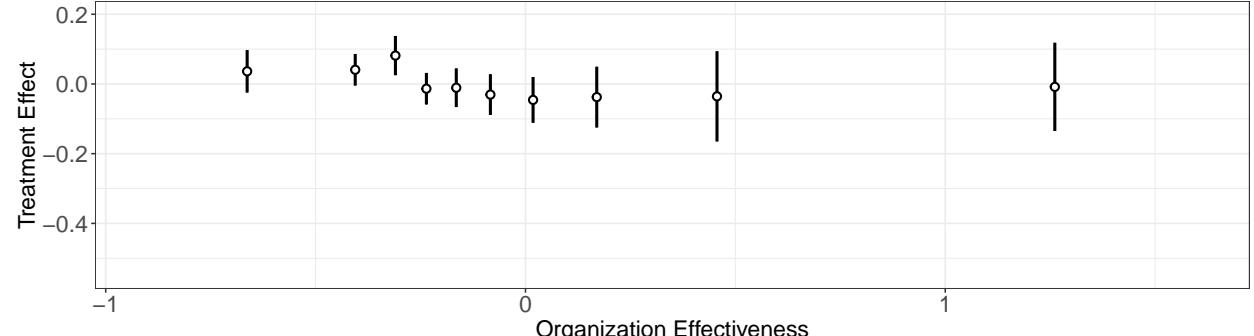
PANEL B: DIFFERENCE IN DIFFERENCES BY BUREAUCRAT EFFECTIVENESS



PANEL C: EVENT STUDY BY ORGANIZATION EFFECTIVENESS

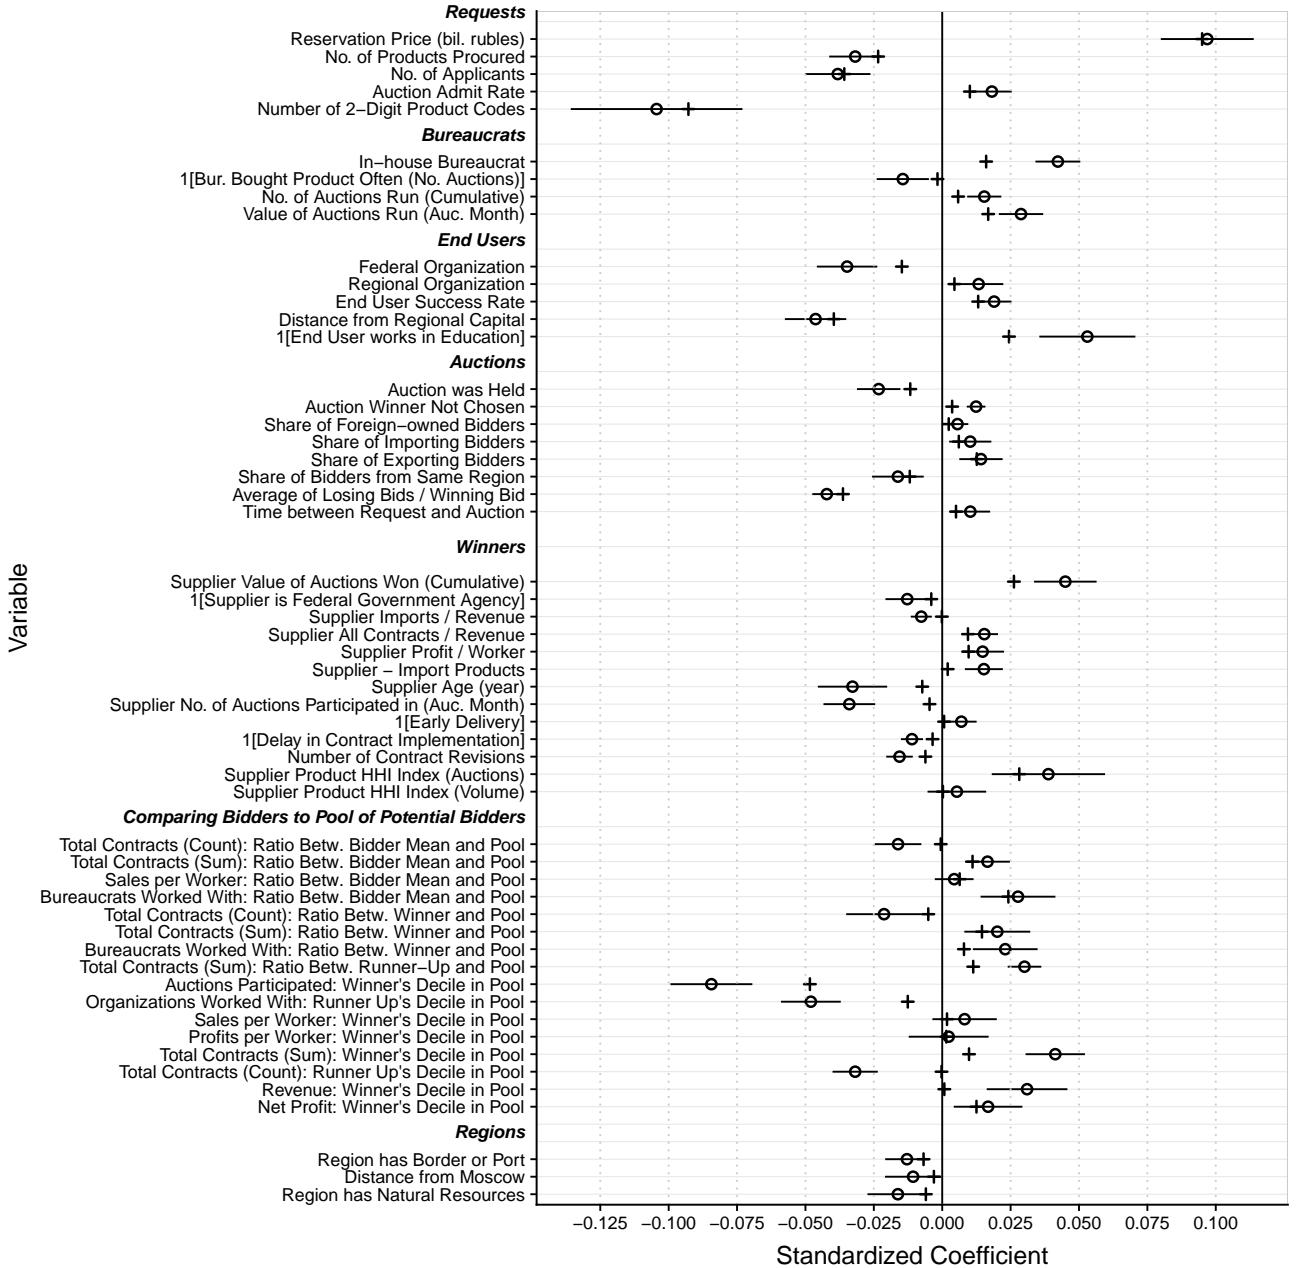


PANEL D: DIFFERENCE IN DIFFERENCES BY ORGANIZATION EFFECTIVENESS



The figure shows how the impacts of the introduction of bid preferences varies by the effectiveness of the implementing bureaucracy. Panels A and C extend the event study (6) shown in figure 4 to estimate separate effects by quartile of bureaucrat (panel A) and organization (panel C) effectiveness. Panels B and D extend the triple difference model (7) shown in table 6 to estimate separate effects by decile of bureaucrat (panel B) and organization (panel D) effectiveness. The horizontal axis plots the average effectiveness within the relevant decile.

FIGURE 6: 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 (8):

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \theta_b \hat{\alpha}_b + \theta_j \hat{\psi}_j + \delta \text{Pref}_{gt} \text{Active}_t + \gamma_b \text{Pref}_{gt} \hat{\alpha}_b + \gamma_j \text{Pref}_{gt} \hat{\psi}_j \\ + \eta_b \text{Active}_t \hat{\alpha}_b + \eta_j \text{Pref}_{gt} \hat{\psi}_j + \text{Pref}_{gt} \text{Active}_t \mathbf{Y}_{igt} \pi + \varepsilon_{igt}$$

where the elements of the vector of observables \mathbf{Y}_{igt} are picked by LASSO using the largest regularization penalty that returns 50 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 50 observables.

TABLE 1: SUMMARY STATISTICS

	All Products			Pharmaceuticals Subsample		
	(1) All No Bid Preferences	(2) Analysis Sample	(3) Analysis Sample incl. Bid Preferences	(4) All No Bid Preferences	(5) Analysis Sample	(6) Analysis Sample incl. Bid Preferences
(1) # of Bureaucrats	115,859	54,771	54,771	5,561	2,505	2,505
(2) # of Organizations	88,326	59,574	59,574	3,662	1,884	1,884
(3) # of Connected Sets	26,239	984	984	0	129	129
(4) # of Bureaucrats with >1 Org.	14,093	12,538	12,888	965	929	1,101
(5) # of Organizations with >1 Bur.	54,580	42,438	42,995	2,076	1,454	1,604
(6) # of Federal Organizations	12,890	2,022	2,022	496	26	26
(7) # of Regional Organizations	25,164	19,014	19,014	2,786	1,613	1,613
(8) # of Municipal Organizations	50,272	38,538	38,538	380	245	245
(9) # of Health Organizations	10,167	7,896	7,896	3,172	1,719	1,719
(10) # of Education Organizations	42,062	33,223	33,223	109	63	63
(11) # of Internal Affairs Organizations	3,126	867	867	105	3	3
(12) # of Agr/Environ Organizations	1,032	339	339	26	1	1
(13) # of Other Organizations	31,939	17,249	17,249	250	98	98
(14) # of Goods	16,376	15,727	16,223	4,220	3,863	4,354
(15) # of Regions	86	86	86	85	79	79
(16) # of Auction Requests	1,733,449	1,249,770	1,930,936	62,755	42,929	115,318
(17) Mean # of Applicants	3.6	3.6	3.46	2.98	3.03	3.02
(18) Mean # of Bidders	2.83	2.14	2.78	1.95	1.98	2.01
(19) Mean Reservation Price	23,460	25,004	25,762	40,708	44,425	38,684
(20) Quantity Mean	1,134	1,117	1,164	1,201	1,718	972
Median	20	25	26	40	45	50
SD	80,831	91,806	174,315	136,260	172,099	108,426
(21) Total Price Mean (bil. USD)	93.2	84.4	84.3	128	91.2	100
Median	4.67	4.35	4.77	6.23	6.7	7.06
SD	577	535	513	5,745	493	524
(22) Unit Price Mean (bil. USD)	72.1	66.5	58.4	20.2	25.4	28.8
Median	0.21	0.17	0.182	0.175	0.18	0.18
SD	21,248	23,341	19,038	226	265	281
(23) # of Observations	15,096,663	11,516,088	16,575,168	290,483	182,060	461,989
(24) Total Procurement Volume (bil. USD)	516	399	635	14.5	9.38	19.9

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 unpreferred auctions. Analysis Sample denotes all unpreferred auctions in connected sets that fulfill three restrictions: singleton bureaucrat-organization, bureaucrat-good, and organization-

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)	1.385	(0.033)	1.483	(0.0328)	0.860	0.626
(2) s.d. of Organization Effects (across orgs)	1.209	(0.0277)	1.317	(0.0214)	0.743	0.519
(3) s.d. of Connected Set Effects (across CS)	0.843	(0.0341)	0.511	(0.0337)	0.318	0.318
(4) s.d. of Bureaucrat Effects (across items)	0.747	(0.0396)	0.775	(0.0202)	0.589	0.311
(5) s.d. of Organization Effects (across items)	0.827	(0.0445)	0.867	(0.0288)	0.644	0.336
(6) s.d. of Connected Set Effects (across items)	0.402	(0.0563)	0.401	(0.0274)	0.136	0.136
(7) Bur-Org Effect Correlation (across items)	-0.665	(0.0166)	-0.432	(0.0315)	-0.602	0.331
(8) s.d. of Total Bur + Org Effects (across items)	0.630	(0.0425)	0.652	(0.0234)	0.525	0.512
(9) s.d. of log unit price	2.197		2.197		2.197	2.197
(10) s.d. of log unit price good, month	1.283		1.283		1.283	1.283
(11) Adjusted R-squared	0.964		0.964		0.964	0.964
(12) Number of Bureaucrats	54,771		54,771		54,771	54,771
(13) Number of Organizations	59,574		59,574		59,574	59,574
(14) Number of Bureaucrat-Organization Pairs	284,710		284,710		284,710	284,710
(15) Number of Connected Sets	984		984		984	984
(16) Number of Observations	11,516,088		11,516,088		11,516,088	11,516,088

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 1–3 show the s.d. of the bureaucrat, organization and connected set effects. Rows 4–8 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): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. 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 $\min_{\Lambda} \mathbb{E}[(\theta - \Lambda\hat{\theta})'(\theta - \Lambda\hat{\theta})]$ where $\hat{\theta}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 4.2.

TABLE 3: ROBUSTNESS TO RESTRICTING TO PHARMACEUTICALS SUBSAMPLE WITH BARCODE INFORMATION

	Fixed Effects (1)	Split Sample (s.e.) (2)	Shrinkage (s.e.) (4)	Covariance Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	0.209	(0.0074)	0.216	(0.00728)	0.103
(2) s.d. of Organization Effects (across orgs)	0.165	(0.00627)	0.173	(0.00634)	0.078
(3) s.d. of Connected Set Effects (across CS)	0.228	(0.0116)	0.152	(0.0123)	0.191
(4) s.d. of Bureaucrat Effects (across items)	0.124	(0.0122)	0.132	(0.0105)	0.0885
(5) s.d. of Organization Effects (across items)	0.129	(0.0131)	0.137	(0.0108)	0.0815
(6) s.d. of Connected Set Effects (across items)	0.152	(0.00841)	0.129	(0.00794)	0.137
(7) Bur-Org Effect Correlation (across items)	-0.333	(0.0783)	-0.164	(0.0439)	-0.264
(8) s.d. of Total Bur + Org Effects (across items)	0.190	(0.00779)	0.179	(0.00679)	0.163
(9) s.d. of log unit price	1.915		1.915		1.915
(10) s.d. of log unit price good, month	0.319		0.319		0.319
(11) Adjusted R-squared	0.997		0.997		0.997
(12) Number of Bureaucrats	2,501		2,501		2,501
(13) Number of Organizations	1,880		1,880		1,880
(14) Number of Bureaucrat-Organization Pairs	8,112		8,112		8,112
(15) Number of Connected Sets	129		129		129
(16) Number of Observations	181,493		181,493		181,493

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. All methods are described fully in Section 4.2.

TABLE 4: ROBUSTNESS TO USING SUBSAMPLES OF INCREASINGLY HETEROGENEOUS GOODS

	Quintile 1	Quintiles 1–2	Quintiles 1–3	Quintiles 1–4	Quintiles 1–5	10-Digit Codes
(1) s.d. of Bur + Org Effects Within CS (across items)	0.466	0.487	0.581	0.660	0.674	0.573
(2) s.d. of Total Bur + Org Effects (across items)	0.488	0.561	0.614	0.644	0.735	0.629
(3) s.d. of log P	1.092	1.592	1.801	1.960	2.084	2.043
(4) s.d. of log P good, month	0.745	0.921	1.053	1.151	1.198	1.114
(5) s.d. of Bur+Org Within Efs / s.d. of log P good, month	0.626	0.529	0.551	0.573	0.563	0.515
(6) s.d. of Bur+Org Total Efs / s.d. of log P good, month	0.655	0.608	0.583	0.559	0.613	0.565
(7) Sample Size	713,366	1,244,778	1,860,274	2,461,987	2,960,786	4,227,349

The table shows the components of the variance due to bureaucrats, organizations, and connected sets estimated by implementing the variance decomposition in equation (4) (see notes to Table 2 for details). Column (6) 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. Column (5) uses the sub-set of the sample in Column (6) that we can match to the scope-for-quality-differentiation ladder developed by Sutton (1998). Column (4) removes the quintile with the highest scope-for-quality-differentiation according to the Sutton (1998) ladder, Column (3) the highest two quintiles, and so on.

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

	All Products		Pharmaceuticals		
	Log Price (1)	No. Bidders (2)	Log Price (3)	No. Bidders (4)	Domestic Winner (5)
log Standardized Quantity	-0.308 (0.016)	0.043 (0.003)	-0.030 (0.003)	0.012 (0.004)	0.003 (0.001)
Preferred * Policy Active	-0.019 (0.012)	-0.059 (0.035)	-0.026 (0.017)	-0.025 (0.027)	0.049 (0.011)
Constituent Terms	Yes	Yes	Yes	Yes	Yes
Month, Good FEs	Yes	Yes	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes	Yes	Yes
Outcome Mean	5.57	2.11	6.22	1.89	0.36
Observations	16,575,168	16,575,168	461,989	461,989	293,538
R ²	0.658	0.287	0.948	0.286	0.735

This table estimates the Intent to Treat (ITT) from equation (5): $y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Preferred}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt}$. The sample used is summarized in columns (3) and (6) of Table 1. In the All Products sample an item has $\text{Preferred}_{gt} = 1$ if the type of good appears on the list of goods covered by the preferences policy for that year. In the Pharmaceuticals sample, $\text{Preferred}_{gt} = 1$ if the drug purchased is made both in Russia and abroad. $\text{PolicyActive}_t = 1$ 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)	No. Bidders (2)	Log Price (3)	No. Bidders (4)	Domestic Winner (5)
log Standardized Quantity	-0.310 (0.016)	0.042 (0.003)	-0.027 (0.003)	0.008 (0.003)	0.003 (0.001)
Bureaucrat FE * Preferred * Policy Active	-0.120 (0.025)	0.070 (0.034)	-0.351 (0.102)	1.004 (0.225)	0.162 (0.066)
Organization FE * Preferred * Policy Active	-0.083 (0.030)	0.014 (0.028)	0.029 (0.116)	0.476 (0.249)	-0.119 (0.057)
Constituent Terms	Yes	Yes	Yes	Yes	Yes
Month, Good FEs	Yes	Yes	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes	Yes	Yes
Outcome Mean	5.57	2.11	6.22	1.89	0.36
Observations	16,575,168	16,575,168	461,989	461,989	293,538
R ²	0.664	0.293	0.950	0.302	0.736

This table estimates the triple-difference from equation (7): $y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Pref}_{gt}\text{Active}_t + \gamma_b\text{Pref}_{gt}\hat{\alpha}_b + \gamma_j\text{Pref}_{gt}\hat{\psi}_j + \eta_b\text{Active}_t\hat{\alpha}_b + \eta_j\text{Active}_t\hat{\psi}_j + \pi_b\text{Pref}_{gt}\text{Active}_t\hat{\alpha}_b + \pi_j\text{Pref}_{gt}\text{PolicyActive}_t\hat{\psi}_j + \varepsilon_{igt}$. The sample used is summarized in columns (3) and (6) of Table 1. In the All Products sample an item has $\text{Pref} = 1$ if the type of good appears on the list of goods covered by preferences that year. In the Pharmaceuticals sample, $\text{Pref} = 1$ if the drug purchased is made both in Russia and abroad. $\text{Active} = 1$ 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 4. Standard errors are clustered by month and good.

OA Online Appendix (For Web Publication Only)

OA.1 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.

OA.1.1 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 2.4. Each product description is parsed in the following way, using the Russian libraries for Python’s Natural Language Toolkit⁶²

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’.

⁶²Documentation on the Natural Language Toolkit (NLTK) can be found at <http://www.nltk.org/>

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].⁶³ Similarly, the product description “sodium bicarbonate - solution for infusion 5%,200ml” would result in the following tokens: [sodium, bicarbonate, solution, infusion, 5%, 200ml].⁶⁴

OA.1.2 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 $\mathbf{x}_i \in \mathcal{X}_C$.⁶⁵ 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 $\hat{g}_C(\mathbf{x}) : \mathcal{X}_C \rightarrow \mathcal{H}_C$ 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:

$$\hat{g}_C(\mathbf{x}_i) = \operatorname{argmax}_{g \in \mathcal{G}_C} p_g(\mathbf{x}_i) \quad (\text{OA.1})$$

Each binary classifier is a logistic regression solving

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

⁶³The 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].

⁶⁴The 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].

⁶⁵The 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].

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 (OA.1). We implement this procedure using the Vowpal Wabbit library for Python.⁶⁶ This simple procedure is remarkably effective; when trained on a randomly selected half of the customs data and then implemented on the remaining data for validation, the classifications are correct 95% of the time. Given this high success rate without regularization, 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.⁶⁷ That is, a classifier that solves

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

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 4.3 presented in figure 1 and table 2 is repeated using these alternative classifiers in figure OA.2 panels C and D and in table OA.9. As they show, the results are remarkably 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 OA.1 (translated into English) and OA.2 (in the original Russian). The three

⁶⁶See <http://hunch.net/~vw/>.

⁶⁷A 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 OA.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

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 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 OA.1.3 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 (OA.1). Second, the similarity between the sentence and the average sentence with the sentence's assigned classification in the *customs* data used

TABLE OA.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, возраст, 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

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} .⁶⁸ For each good, we then calculate the mean sentence vector in the customs data as $\mathbf{v}_g^C = \sum_{\mathbf{v}_{gi}, \mathbf{x}_i \in \mathbf{X}^C} \mathbf{v}_{gi} / |\mathbf{t}_g|$. Then, to identify outlier sentences in the procurement data, we calculate each sentence's normalized cosine similarity with the good's mean vector,

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

where $s(\mathbf{v}_{gi}, \mathbf{v}_g) \equiv \cos(\mathbf{v}_{gi}, \mathbf{v}_g) = \frac{\mathbf{v}_{gi} \cdot \mathbf{v}_g}{\|\mathbf{v}_{gi}\| \|\mathbf{v}_g\|} = \frac{\sum_{k=1}^{K_g} t_{gik} t_{gk}}{\sqrt{\sum_{k=1}^{K_g} t_{gik}^2} \sqrt{\sum_{k=1}^{K_g} t_{gk}^2}}$ is the cosine similarity of the sentence vector \mathbf{v}_{gi} with its good mean \mathbf{v}_g ,⁶⁹ 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 OA.2 panels A and B and Table OA.10 show the robustness of our results to using the 45th or

⁶⁸Note 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

⁶⁹Note that the cosine similarity ranges from 0 to 1, with 0 being orthogonal vectors and 1 indicating vectors pointing in the same direction.

55th percentile as thresholds.

OA.1.3 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

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

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 its 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

$$\eta(i) = \frac{b(i) - a(i)}{\max\{b(i), a(i)\}} \quad (\text{OA.6})$$

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 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.

OA.2 Proofs of Propositions

OA.2.1 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

$$q_j \mathbb{E}[\pi_i] + (1 - q_j)(\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i]) = c_i \iff q_j = \frac{\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i] - c_i}{\bar{\theta} - \mathbb{E}[\bar{\theta}/\theta_i] - \mathbb{E}[\pi_i]}, \quad (\text{OA.7})$$

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

For the firms to be indifferent between entering and not entering, equation (OA.7) 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 : $G_i(\theta_i) = 1 - \theta_i^{-\delta_i}$.

$$H_i(b) \equiv \mathbb{P}(b_i \leq b) = \mathbb{P}\left(\theta_i \geq \frac{\bar{\theta}}{b}\right) = \left(\frac{b}{\bar{\theta}}\right)^{\delta_i} \quad (\text{OA.8})$$

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

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

$$\begin{aligned}\mathbb{E}[\pi_i|b_i] &= \mathbb{E}_{b_j}[b_j - b_i | b_j > b_i] \mathbb{P}(b_j > b_i) = \int_{b_i}^{\bar{\theta}} (b_j - b_i) dH_j(b_j) \\ &= \frac{\delta_j}{1 + \delta_j} \bar{\theta} - b_i + b_i \left(\frac{b_i}{\bar{\theta}} \right)^{\delta_j} \frac{1}{1 + \delta_j},\end{aligned}\tag{OA.9}$$

where the final equality follows by inserting (OA.8) and integrating. Now we can derive unconditional expected profits by the law of iterated expectations:

$$\mathbb{E}[\pi_i] = \mathbb{E}_{b_i}[\mathbb{E}[\pi_i|b_i]] = \int_0^{\bar{\theta}} \mathbb{E}[\pi_i|b_i] dH_i(b_i) = \left(\frac{1}{1 + \delta_i} - \frac{1}{1 + \delta_F + \delta_L} \right) \bar{\theta}.\tag{OA.10}$$

Inserting these and the definition of the entry costs c_i into (OA.7) and rearranging yields the statement in the proposition

$$q_i = \sqrt{\kappa(1 - \alpha_c - \psi_c)},\tag{OA.11}$$

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.

$$\mathbb{P}(p \leq x) = \mathbb{P}(\max\{b_F, b_L\} \leq x) = H_F(x)H_L(x) = \left(\frac{x}{\bar{\theta}} \right)^{\delta_F + \delta_L}\tag{OA.12}$$

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

$$\begin{aligned}\mathbb{P}(\log(p) \leq x) &= \mathbb{P}(p \leq e^x) = \left(\frac{e^x}{\bar{\theta}} \right)^{\delta_F + \delta_L} \\ \mathbb{E}[\log(p) | \text{both enter}] &= \int_{-\infty}^{\log(\bar{\theta})} x \frac{\delta_F + \delta_L}{M^{\delta_F + \delta_L}} e^{(\delta_F + \delta_L)x} dx = \log(\bar{\theta}) - \frac{1}{\delta_F + \delta_L}\end{aligned}\tag{OA.13}$$

The expected log price is then simply $\mathbb{E}[\log(p)] = q_F q_L \mathbb{E}[\log(p) | \text{both enter}] + (1 - q_F q_L) \log(\bar{\theta})$. Inserting (OA.13) 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

OA.2.2 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

$$\begin{aligned}\mathbb{E}[\pi_F|b_F, \gamma] &= \mathbb{E}[\gamma(b_L - b_F)|b_L > b_F]\mathbb{P}(b_L > b_F) \\ \mathbb{E}[\pi_L|b_L, \gamma] &= \mathbb{E}[b_F - b_L|\bar{\theta} \geq b_F > b_L]\mathbb{P}(\bar{\theta} \geq b_F > b_L) + \mathbb{P}(\theta_F < 1/\gamma)(\bar{\theta} - b_L).\end{aligned}\tag{OA.14}$$

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 $\alpha_c + \psi_c \leq \underline{c}$. 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 (OA.14) over all bids,

$$\mathbb{E}[\pi_F|\gamma < 1] = \int_0^M \mathbb{E}[\pi_F|b_F, \gamma < 1] dH_F(b_F|\gamma < 1) = \gamma^{1+\delta_F} M \left(\frac{1}{1+\delta_F} - \frac{1}{1+\delta_F + \delta_L} \right)\tag{OA.15}$$

Setting (OA.15) 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

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

In turn, the distribution of log prices is given by

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

making the expected log price in the auction

$$\begin{aligned}\mathbb{E}[\log(p) | \text{both enter}] &= \int_{-\infty}^{\log(\gamma\bar{\theta})} \frac{\delta_L \gamma^{-\delta_L + \delta_F \gamma^{\delta_F}}}{\bar{\theta}^{\delta_F + \delta_L}} x e^{(\delta_F + \delta_L)x} dx + \int_{\log(\gamma\bar{\theta})}^{\log(\bar{\theta})} \frac{\delta_F \gamma^{\delta_F}}{\bar{\theta}^{\delta_F + \delta_L}} x e^{(\delta_F + \delta_L)x} dx \\ &\quad + [1 - H_F(\bar{\theta})] \log(\bar{\theta}) \\ &= \log(\bar{\theta}) - \frac{\gamma^{\delta_F} (1 - \log(\gamma^{\delta_L}))}{\delta_F + \delta_L}. \tag{OA.16}\end{aligned}$$

Comparing (OA.16) to the expected price without preferences (OA.13), prices rise as long as $\gamma^{\delta_F} [1 - \log(\gamma^{\delta_L})] > 1$.

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

$$\mathbb{P}(L \text{ wins}) = \mathbb{P}(b_L < b_F) = \int_0^{\bar{\theta}} H_L(b_F | \gamma = 1) dH_F(b_F | \gamma = 1) = 1 - \frac{\delta_L}{\delta_F + \delta_L}, \tag{OA.17}$$

while when there are preferences this increases to

$$\mathbb{P}(L \text{ wins}) = \mathbb{P}(b_L < b_F | \gamma < 1) = \int_0^{\bar{\theta}} H_L(b_F | \gamma < 1) dH_F(b_F | \gamma < 1) = 1 - \gamma^{\delta_F} \frac{\delta_L}{\delta_F + \delta_L}. \tag{OA.18}$$

Buyers with $\underline{c} < \alpha_c + \psi_c \leq \bar{c}$. 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 $\mathbb{E}[\pi_L|\gamma]$:

$$\mathbb{E}[\pi_L|\gamma < 1] = \int_0^{\bar{\theta}} \mathbb{E}[\pi_L|b_L, \gamma < 1] dH_L(b_L|\gamma < 1) = \bar{\theta} \left(\frac{1}{1+\delta_L} - \frac{\gamma^{\delta_F}}{1+\delta_F+\delta_L} \right). \quad (\text{OA.19})$$

Setting (OA.19) 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

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

Inserting the entry probabilities and the price equation (OA.16) and rearranging, the expected price when there are preferences is lower whenever

$$\begin{aligned} & q_F(\gamma < 1) q_L(\gamma < 1) (\log(\bar{\theta}) - \mathbb{E}[\log(p)|\text{both enter}, \gamma < 1]) \\ & - q_F(\gamma = 1) q_L(\gamma = 1) (\log(\bar{\theta}) - \mathbb{E}[\log(p)|\text{both enter}, \gamma = 1]) \geq 0 \\ & \iff -\log(\gamma^{\delta_L}) - \frac{\delta_L}{1+\delta_F} (1 - \gamma^{1+\delta_F}) \geq 0 \end{aligned} \quad (\text{OA.20})$$

Noting that (OA.20) holds with equality when $\gamma = 1$ and that the left hand side of (OA.20) has slope $-\delta_L(\gamma^{-1} - \gamma^{\delta_F}) < 0 \forall \gamma < 1$ shows that (OA.20) holds for all $\gamma < 1$. Participation in

the auction is $\mathbb{E}[N] = q_F + q_L$. When there are no preferences

$$\mathbb{E}[N|\gamma=1] = q_F(\gamma=1) + q_L(\gamma=1) = 2 \frac{1+\delta_F+\delta_L}{1+\delta_L} \sqrt{1-\alpha_c-\psi_c}, \quad (\text{OA.21})$$

while with preferences participation is

$$\begin{aligned} \mathbb{E}[N|\gamma<1] &= q_F(\gamma<1) + q_L(\gamma<1) \\ &= \left(\frac{1}{\gamma^{\delta_F}} + \frac{1}{\gamma^{1+\delta_F} + (1-\gamma^{1+\delta_F}) \frac{1+\delta_F+\delta_L}{1+\delta_F}} \right) \frac{1+\delta_F+\delta_L}{1+\delta_L} \sqrt{1-\alpha_c-\psi_c}. \end{aligned} \quad (\text{OA.22})$$

Comparing (OA.21) to (OA.22) shows that participation increases whenever

$$\frac{1}{\gamma^{\delta_F}} + \frac{1}{1 + \frac{\delta_L}{\delta_F+\delta_L} (1-\gamma^{1+\delta_F})} > 2 \quad (\text{OA.23})$$

Equation (OA.23) 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

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

Combining these two components shows that the statement is correct as long as $\gamma^{\delta_F} \left[1 + \frac{\delta_L}{\delta_F+\delta_L} (1-\gamma^{1+\delta_F}) \right] < 1$. \square

OA.3 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

$$\mathbf{p} = \mathbf{X}\beta + \mathbf{B}\alpha + \mathbf{F}\psi \quad (\text{OA.24})$$

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; α 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 ψ is the $N_o \times 1$ vector of organization effects.

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

$$\begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \hat{\alpha}_{OLS} \\ \hat{\psi}_{OLS} \end{bmatrix} = \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \mathbf{p} \quad (\text{OA.25})$$

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 α s and ψ s.

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

$$\mathbf{p} = \mathbf{X}\beta + \mathbf{B}\tilde{\alpha} + \mathbf{F}\tilde{\psi} + \mathbf{S}\gamma \quad (\text{OA.26})$$

where \mathbf{S} is the $N \times N_s$ design matrix indicating which connected set each item belongs to; γ 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:

$$\begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} & \begin{bmatrix} \mathbf{B} & \mathbf{F} & \mathbf{S} \end{bmatrix} \\ \begin{bmatrix} \mathbf{S}_b & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_o & \mathbf{0} \end{bmatrix} \end{bmatrix} \begin{bmatrix} \hat{\alpha} \\ \hat{\psi} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \mathbf{p} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (\text{OA.27})$$

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 (OA.24), then $\hat{\alpha}$, $\hat{\psi}$, and $\hat{\gamma}$, the estimators of $\tilde{\alpha}$, $\tilde{\psi}$ and γ in the augmented model (OA.26) that solve the augmented normal equations (OA.27) (i) are uniquely identified, and (ii) are related to the true bureaucrat and organization effects α and ψ by*

$$\begin{bmatrix} \hat{\alpha} \\ \hat{\psi} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \alpha - \mathbf{S}_b' \bar{\alpha} \\ \psi - \mathbf{S}_o' \bar{\psi} \\ \bar{\alpha} + \bar{\psi} \end{bmatrix} \quad (\text{OA.28})$$

where $\bar{\alpha}$ is the $N_s \times 1$ vector of connected-set bureaucrat effect means, and $\bar{\psi}$ 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 (OA.26) 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 $[\mathbf{BFS}]$ 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 (OA.27) has rank $N_b + N_o + N_s$ and hence the solution to (OA.27) is unique.

To see the second part, it suffices to show that (OA.28) solves (OA.27). First, substitute the estimators out of (OA.27) using (OA.28) and substitute in the true model using (OA.24)

to rewrite (OA.27) as

$$\begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} [\mathbf{B}(\alpha - \mathbf{S}_b' \bar{\alpha}) + \mathbf{F}(\psi - \mathbf{S}_o' \bar{\psi}) + \mathbf{S}(\bar{\alpha} + \bar{\psi})] \\ \mathbf{S}_b(\alpha - \mathbf{S}_b' \bar{\alpha}) \\ \mathbf{S}_o(\psi - \mathbf{S}_o' \bar{\psi}) \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} [\mathbf{B}\alpha + \mathbf{F}\psi] \\ 0 \\ 0 \end{bmatrix}$$

From here, noting again that $\mathbf{B}\mathbf{S}_b' = \mathbf{F}\mathbf{S}_o' = \mathbf{S}$; that $\mathbf{S}_b\alpha$ 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. \square

OA.4 Robustness of Event Study Design

In Appendix Figure OA.3, we change a series of choices made in constructing the event studies discussed in Sub-section 4.1: the length of time bureaucrats and organizations are required to work together to be part of an event (the top four panels require, rather than two active weeks working together as in Figure 1, two active days, two active fortnights, two active months, and three active weeks, respectively); how coarsely we define effectiveness categories (Panel E categorizes bureaucrats by tercile rather than quartile); and the sample based on which the cut-offs between the different categories are defined (Panel F construct quartiles based on the entire sample period rather than each quarter separately). In each of the panels, the patterns observed when an organization switches bureaucrats are very similar to those in Figure 1.

OA.5 Probing the log-linearity assumption

Misspecification 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 4.1 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, in Appendix OA.5, we examine patterns in the size of residuals across the bivariate distribution of the estimated bureaucrat and organization effects. If match effects omitted from (3) are important, we should see residuals that are systematically larger for

large values of the estimated bureaucrat and/or organization effects. We see no evidence of this in the top panel of Appendix Figure OA.6. In the bottom panel we repeat this exercise, but now we plot the residuals from a model that is analogous to (3) but specified in levels rather than logs. The systematic patterns seen in the bottom panel provide clear evidence that such a model is misspecified.

Third, we re-estimate equation (3) with fixed effects for each bureaucrat-organization *pair* added in Appendix OA.5. The improvement in the model's fit from adding pair effects is minuscule, indicating that a log-linear model is a good approximation to the true, underlying production function (see also Card *et al.*, 2013).

The model we have estimated assumes that the price achieved is approximately log-linear in the bureaucrat and organization effects. 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. Appendix Figure OA.6 shows a heat map of residuals for the analysis sample. The map reveals no clear patterns in the residuals. Appendix Figure shows an analogous heat map of residuals from running (3) in levels rather than logs. The figure suggests 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.

As a further test of our log-linear model of prices, 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 1.139 to 1.112 and increases the R^2 from 0.964 to 0.965, 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.

OA.6 Comparison to existing estimates of individuals' and organizations' effects on output

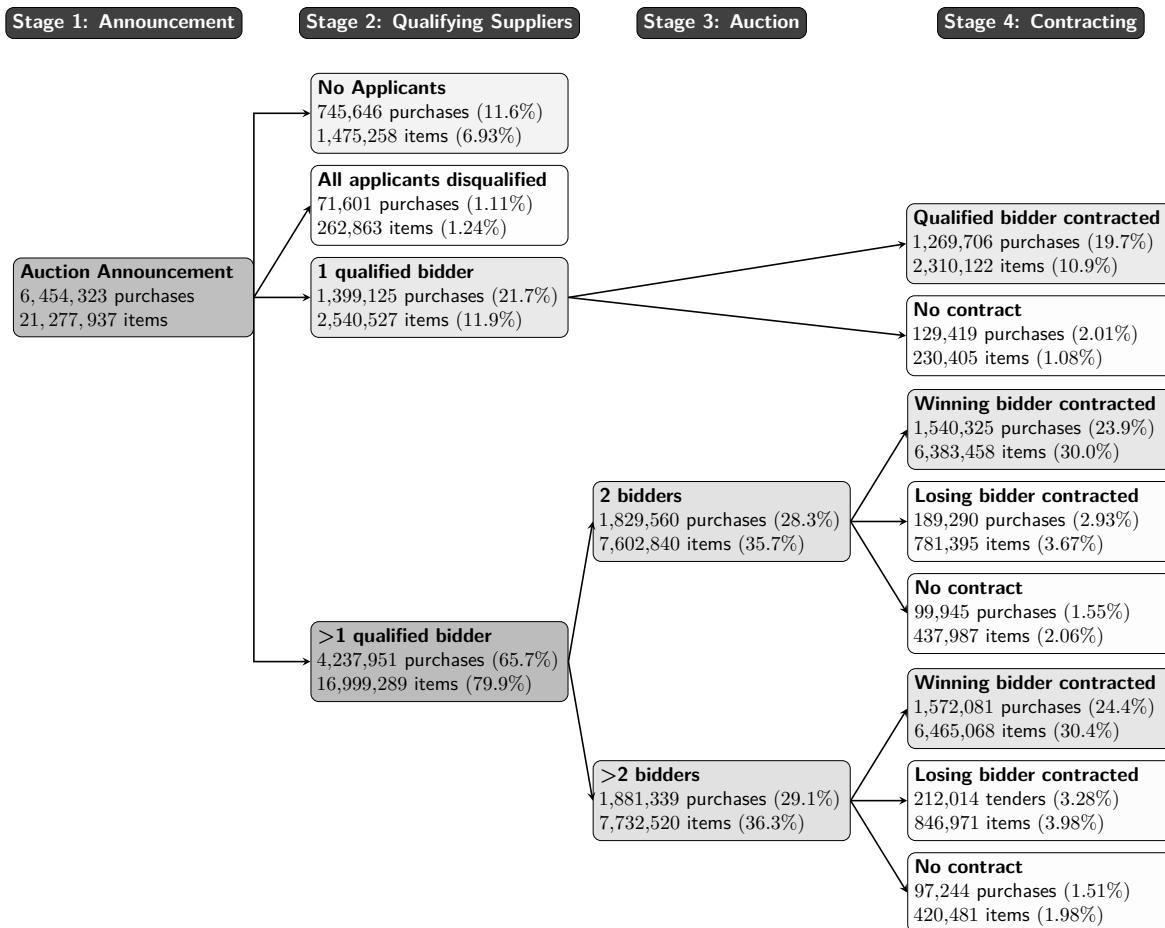
How do our results compare to existing estimates of the extent to which individuals and organizations affect output in other settings? While we are not aware of comparable estimates of the causal effects of workers and organizations on output in a low or middle-income country government context, several 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 29.0 and 44.0 percent respectively.⁷⁰ 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 55.1 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 \(2003\)](#) 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 72.1 percent relative to the top quartile due solely to the bureaucrat effects.

⁷⁰We perform these calculations separately in each connected set and report the average, weighting by the number of items.

OA.7 Additional Figures and Tables

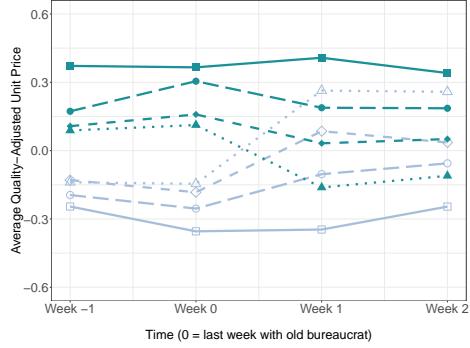
FIGURE OA.1: PROCUREMENT PROCESS FLOW-CHART



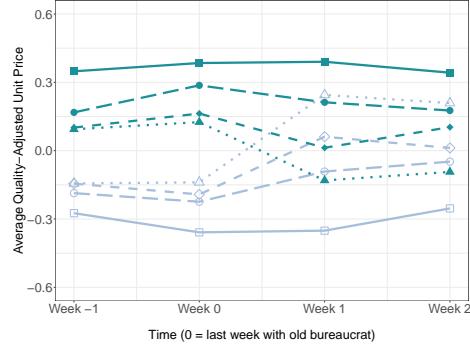
This figure lays out the stages of the process public procurement purchases of off-the-shelf goods through electronic auctions follow in Russia. Numbers are based on all purchases made under laws 94 and 44 in 2011-2016. The stages are described in detail in Sub-section 2.1.

FIGURE OA.2: ROBUSTNESS OF EVENT STUDIES TO ALTERNATIVE TEXT CLASSIFIERS, CLASSIFICATION ACCURACY THRESHOLDS, AND OUTLIER TRIMMING

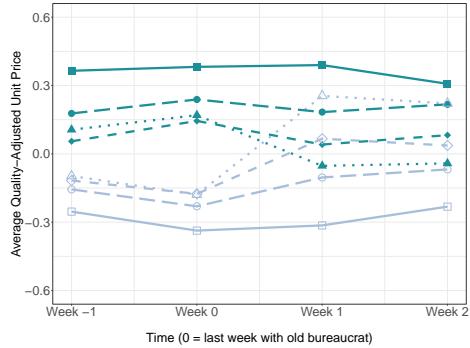
PANEL A: CLASSIFIER ACCURACY THRESHOLD 45TH PERCENTILE



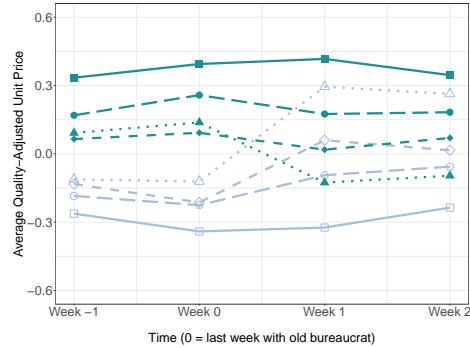
PANEL B: CLASSIFIER ACCURACY THRESHOLD 55TH PERCENTILE



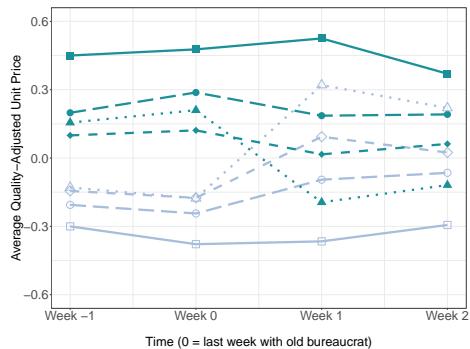
PANEL C: SUPPORT VECTOR MACHINE CLASSIFIER



PANEL D: HIERARCHICAL MODEL CLASSIFIER



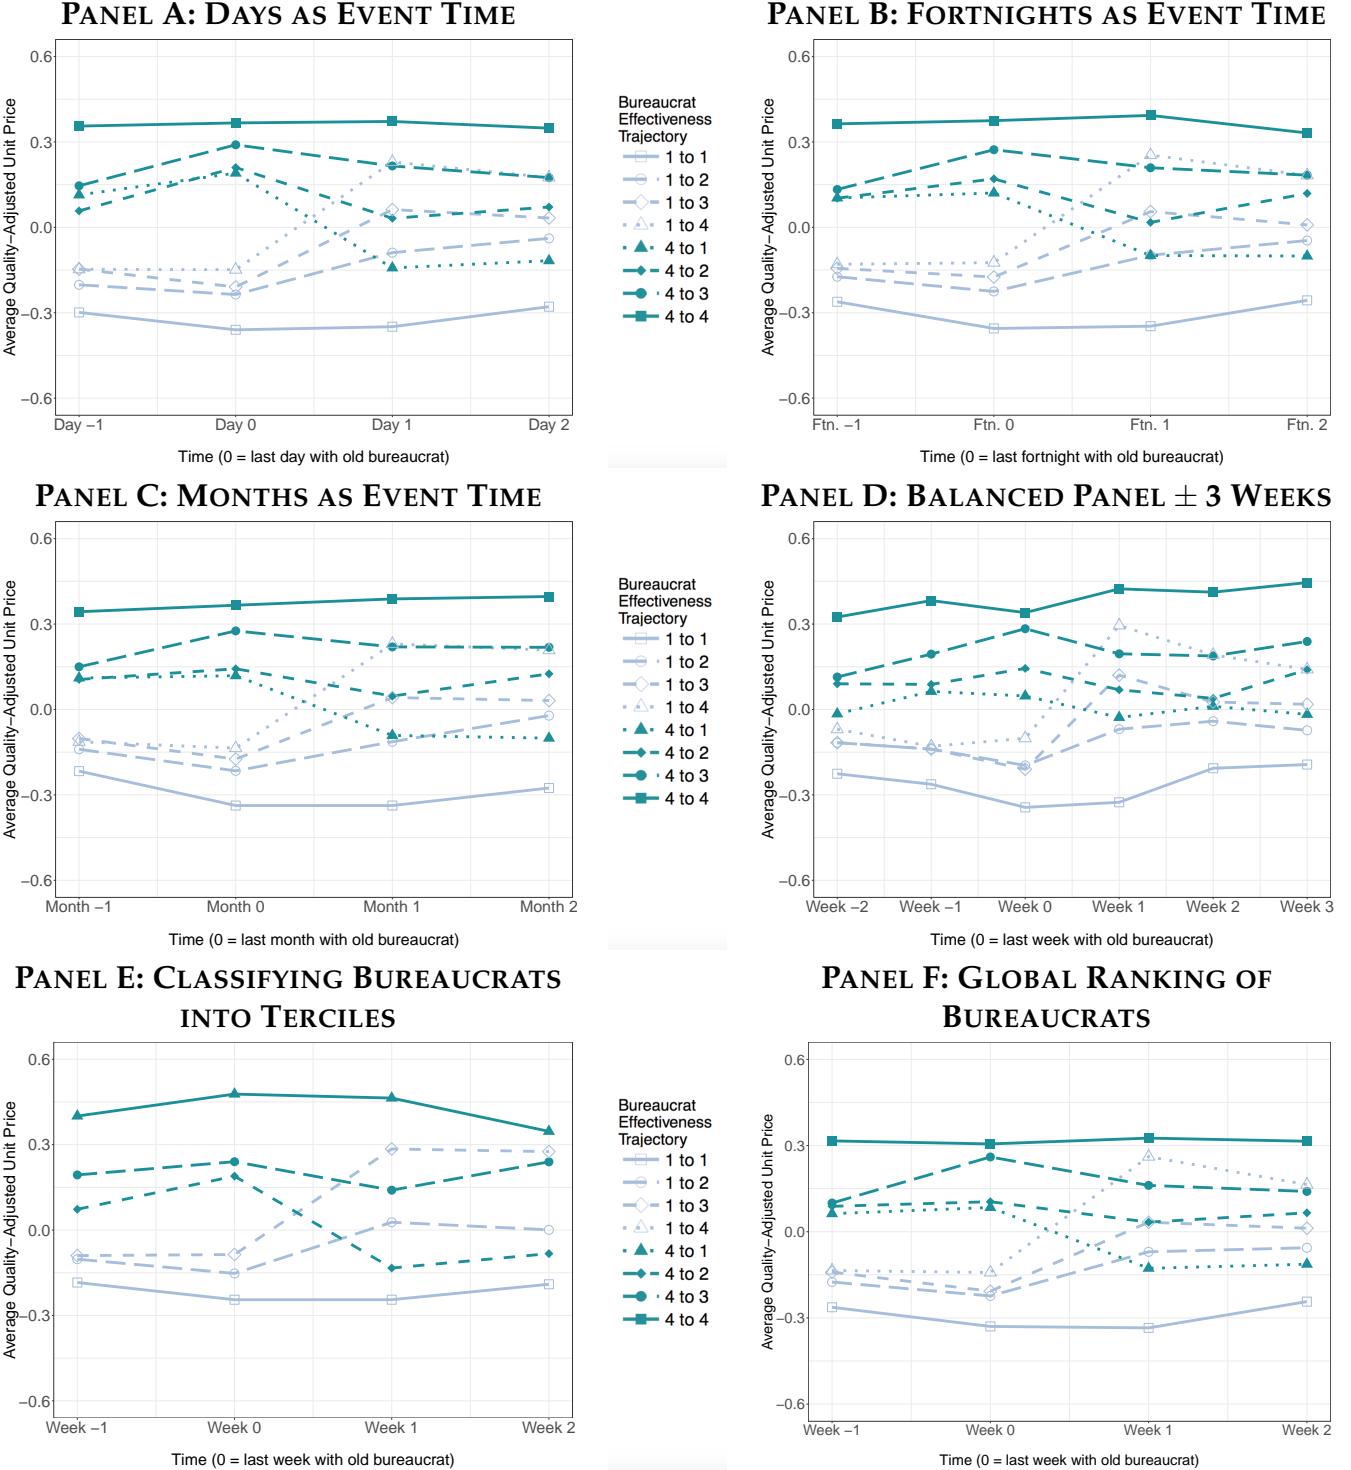
PANEL E: DROPPING TOP AND BOTTOM 2.5% OUTLIERS



Bureaucrat Effectiveness Trajectory
 1 to 1
 1 to 2
 1 to 3
 1 to 4
 4 to 1
 4 to 2
 4 to 3
 4 to 4

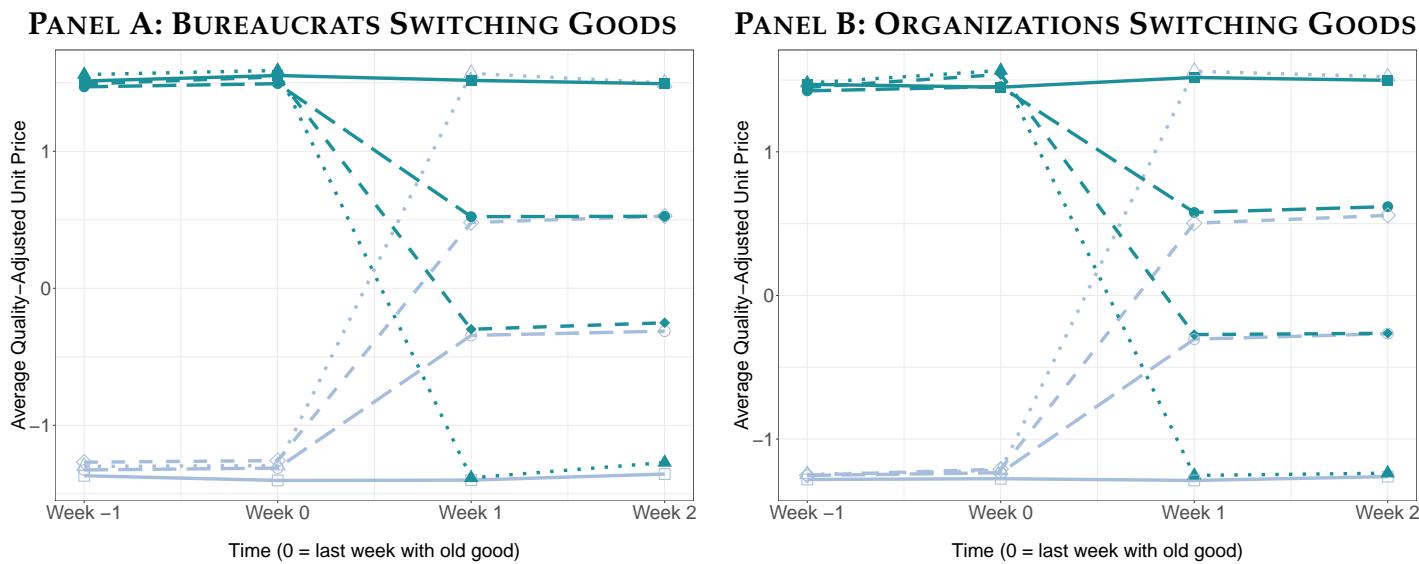
Each panel in the figure is analogous to Figure 1 (see notes to that figure for details of construction), with the following changes. In Panel A, rather than requiring the bureaucrat-organization pair to work together in two separate weeks, we require the pair to work together on two separate days. In 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. In Panel E we categorize bureaucrats by terciles rather than quartiles, and in Panel F we construct quartiles by ranking bureaucrats based on the entire sample period rather than each semester separately.

FIGURE OA.3: ROBUSTNESS OF EVENT STUDIES TO DESIGN CHOICES



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. In Panel E we categorize bureaucrats by terciles rather than quartiles, and in Panel F we construct quartiles by ranking bureaucrats based on the entire sample period rather than each semester separately.

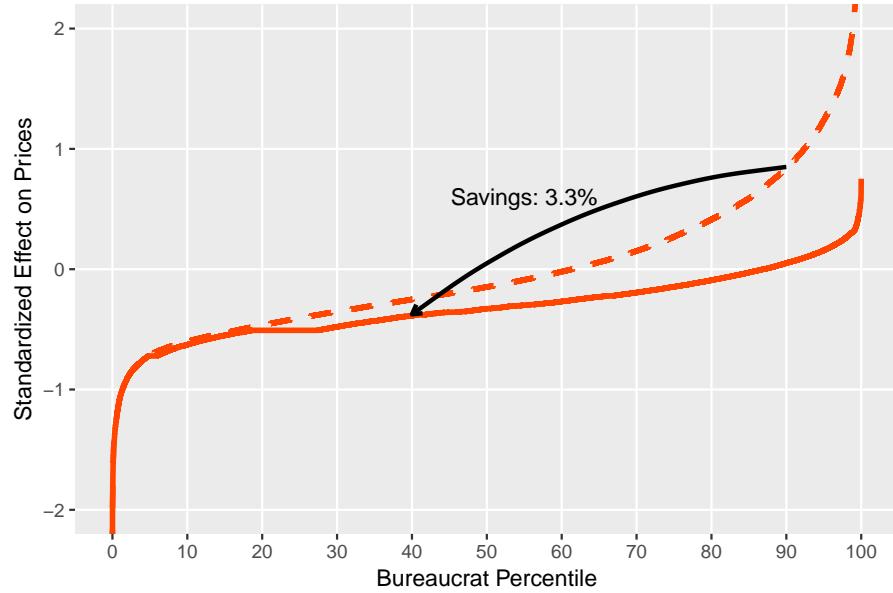
FIGURE OA.4: EVENT STUDY OF PROCUREMENT PRICES AROUND TIMES BUREAUCRATS AND ORGANIZATIONS SWITCH GOODS



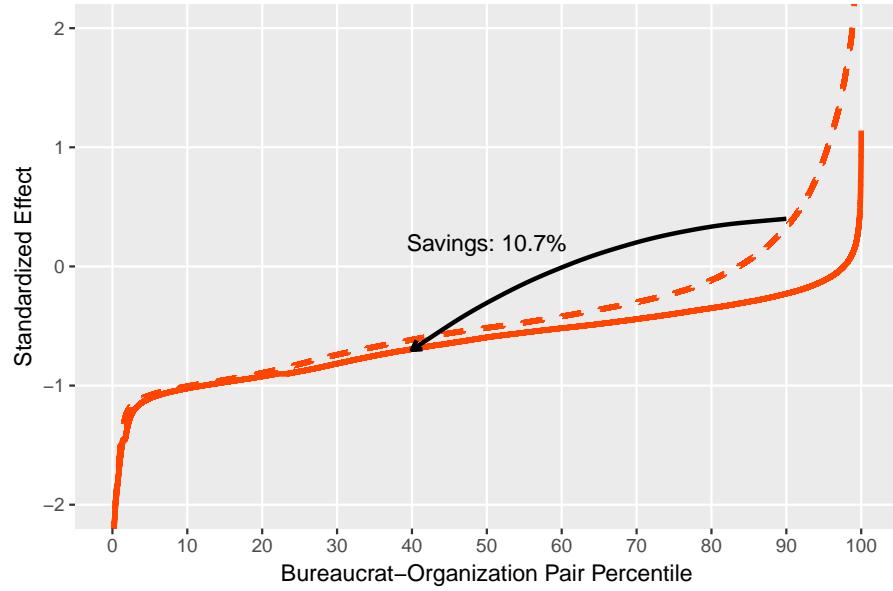
Each panel in the figure is analogous to Figure 1 that studies price changes around the time that organizations switch the bureaucrat making their purchases (see notes to that figure for details of construction). Panel A shows price changes around the time that bureaucrats switch the good they are purchasing.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

FIGURE OA.5: CRUDE COUNTERFACTUALS

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



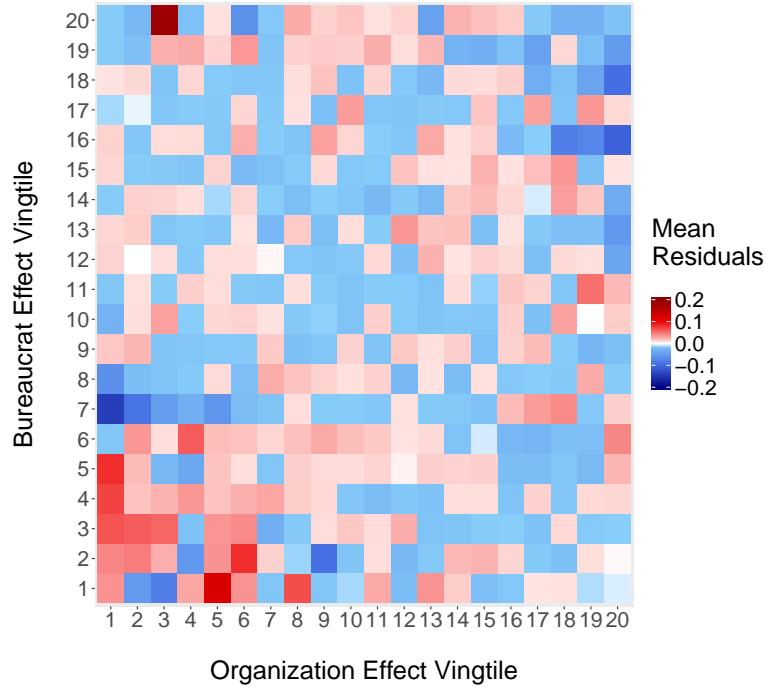
Panel B: Moving Least Effective 25% of Bureaucrats and Organizations to 75th Percentile Effectiveness



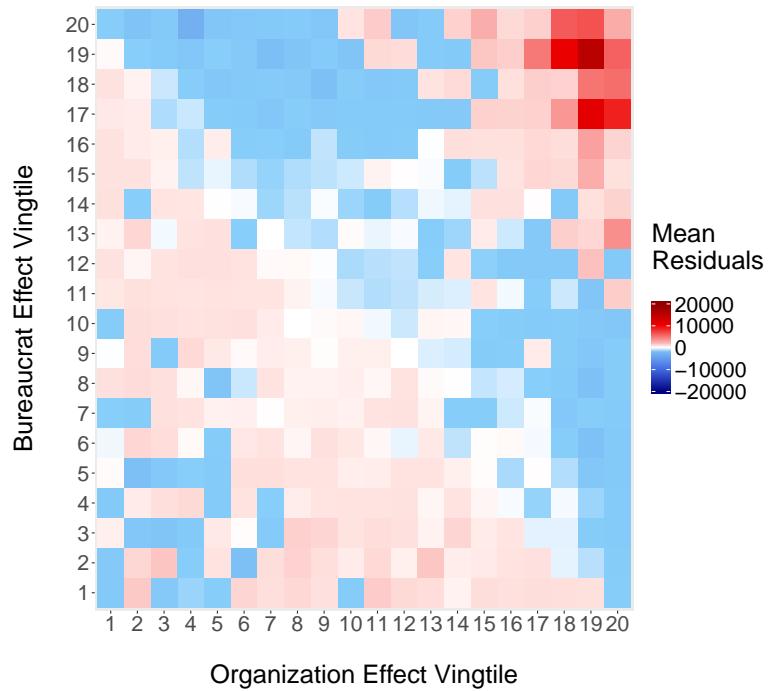
The figure shows the impact of two counterfactual scenarios on the distribution of our estimated price effects. Panel A considers moving all bureaucrats above the 75th percentile of their connected set's distribution of covariance shrunken price effects down to their connected set's 25th percentile. 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.

FIGURE OA.6: CORRELATION OF RESIDUALS WITH ESTIMATED BUREAUCRAT AND ORGANIZATION EFFECTS

PANEL A: PRICES IN LOGS (MAIN SPECIFICATION)

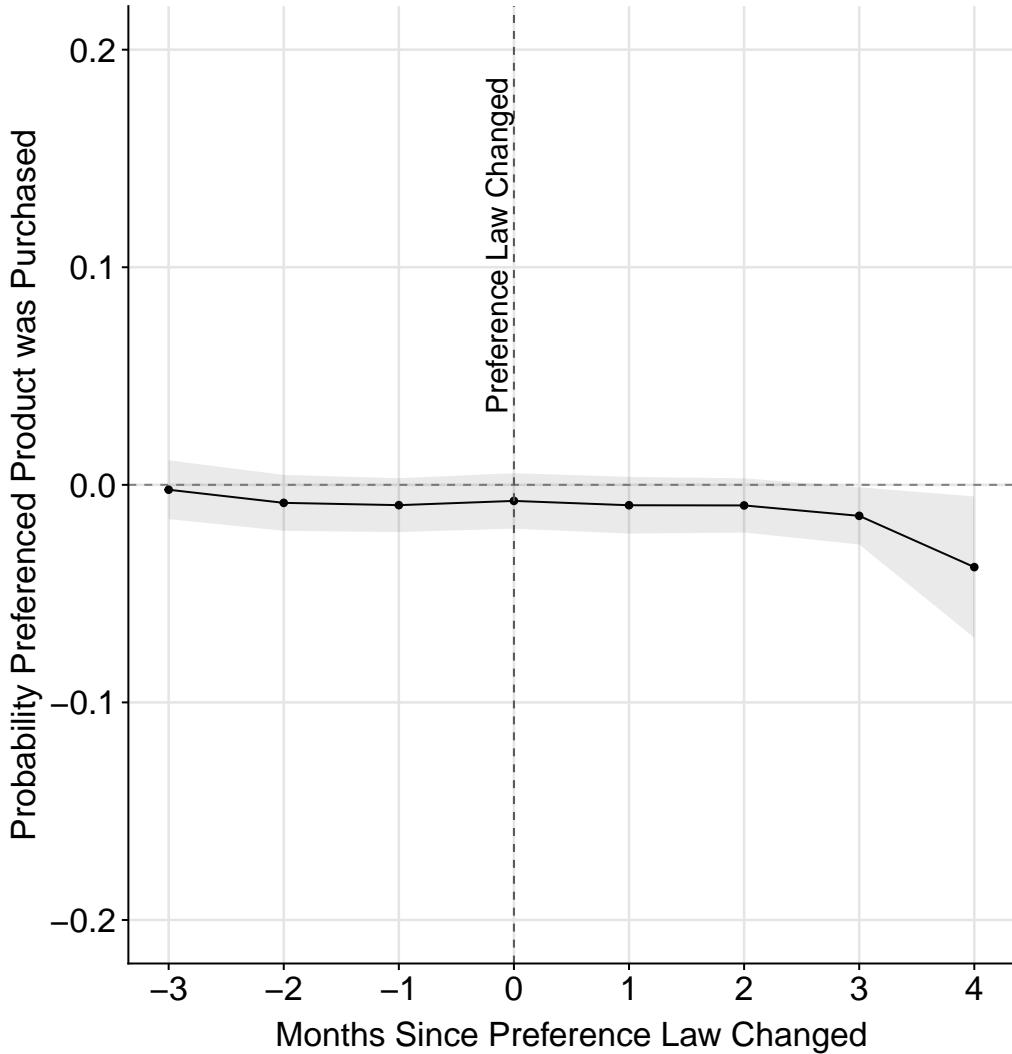


PANEL B: PRICES IN LEVELS (ILLUSTRATING MISSPECIFICATION)



The figure presents heatmaps of averages of the residuals from the estimation of equation (3): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ — 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 OA.7: END USERS DO NOT CHANGE THE TIMING OF THEIR PROCUREMENT IN ANTICIPATION OF PREFERENCE LAWS



Notes: The figure shows a graphical analysis of the timing of procurement around preferences policy over the period of study. 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: $\text{Preferred}_{gt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \mathbf{1}\{t - \text{ListMonth}_t = s\} + \varepsilon_{igt}$, where Preferred_{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 4, but we remove the month fixed effects. ε_{igt} is an error term we allow to be clustered by month and good.

FIGURE OA.8: 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

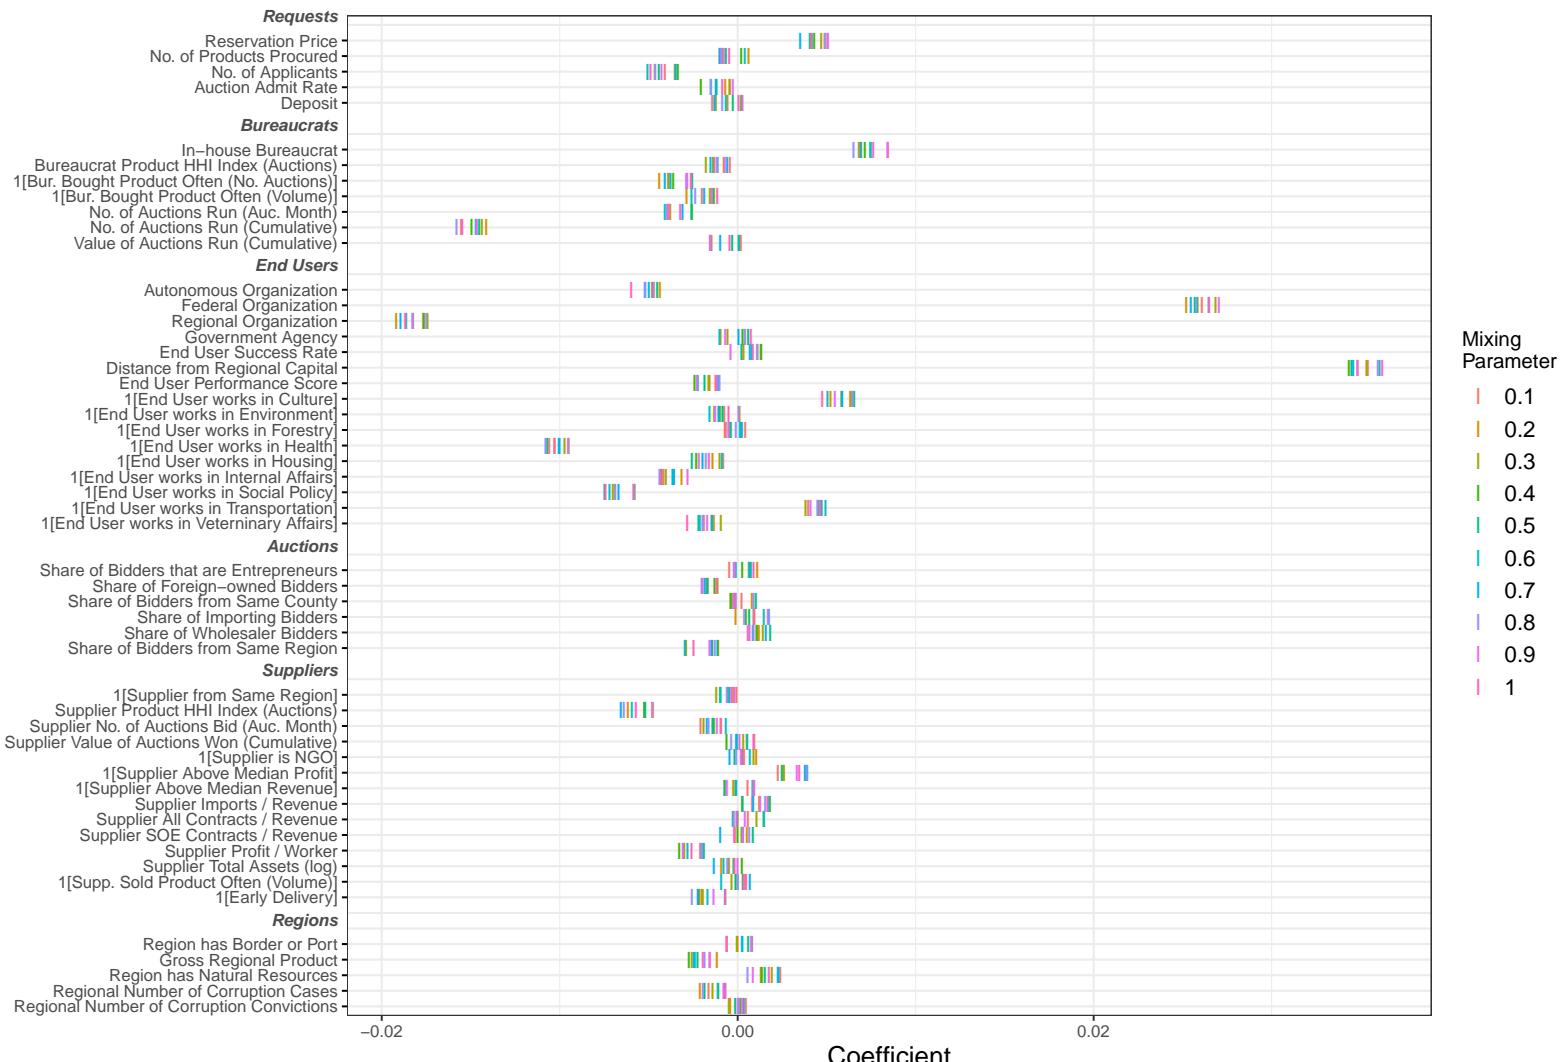
ВСЕГО ГОЛОСОВ: 0
Заявка допущена: 0
Заявка не допущена: 0

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 электронного 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)

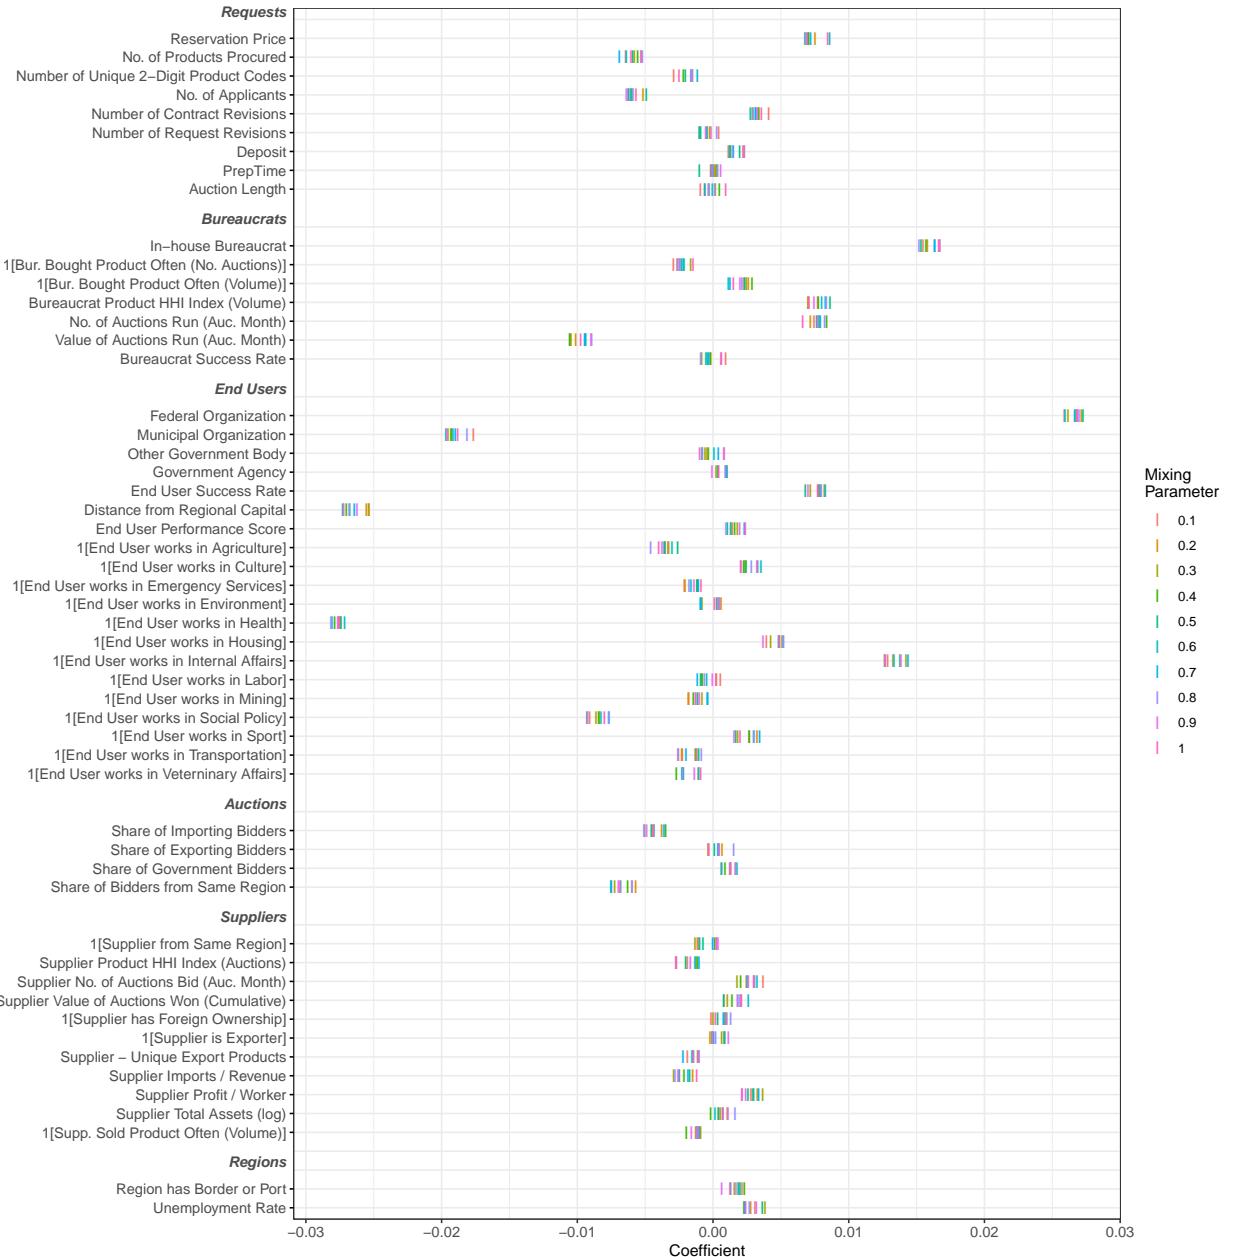
This screenshot is taken from the official protocol for Request #0360200029016000098, an electronic auction for winter shoes conducted by an orphanage in November 2016 in Saratov, Russia. Applicant supplier #2 was rejected by the five-member commission on the grounds that the supplier's application did not adequately describe 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.

FIGURE OA.9: CORRELATES OF BUREAUCRAT EFFECTIVENESS: ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS



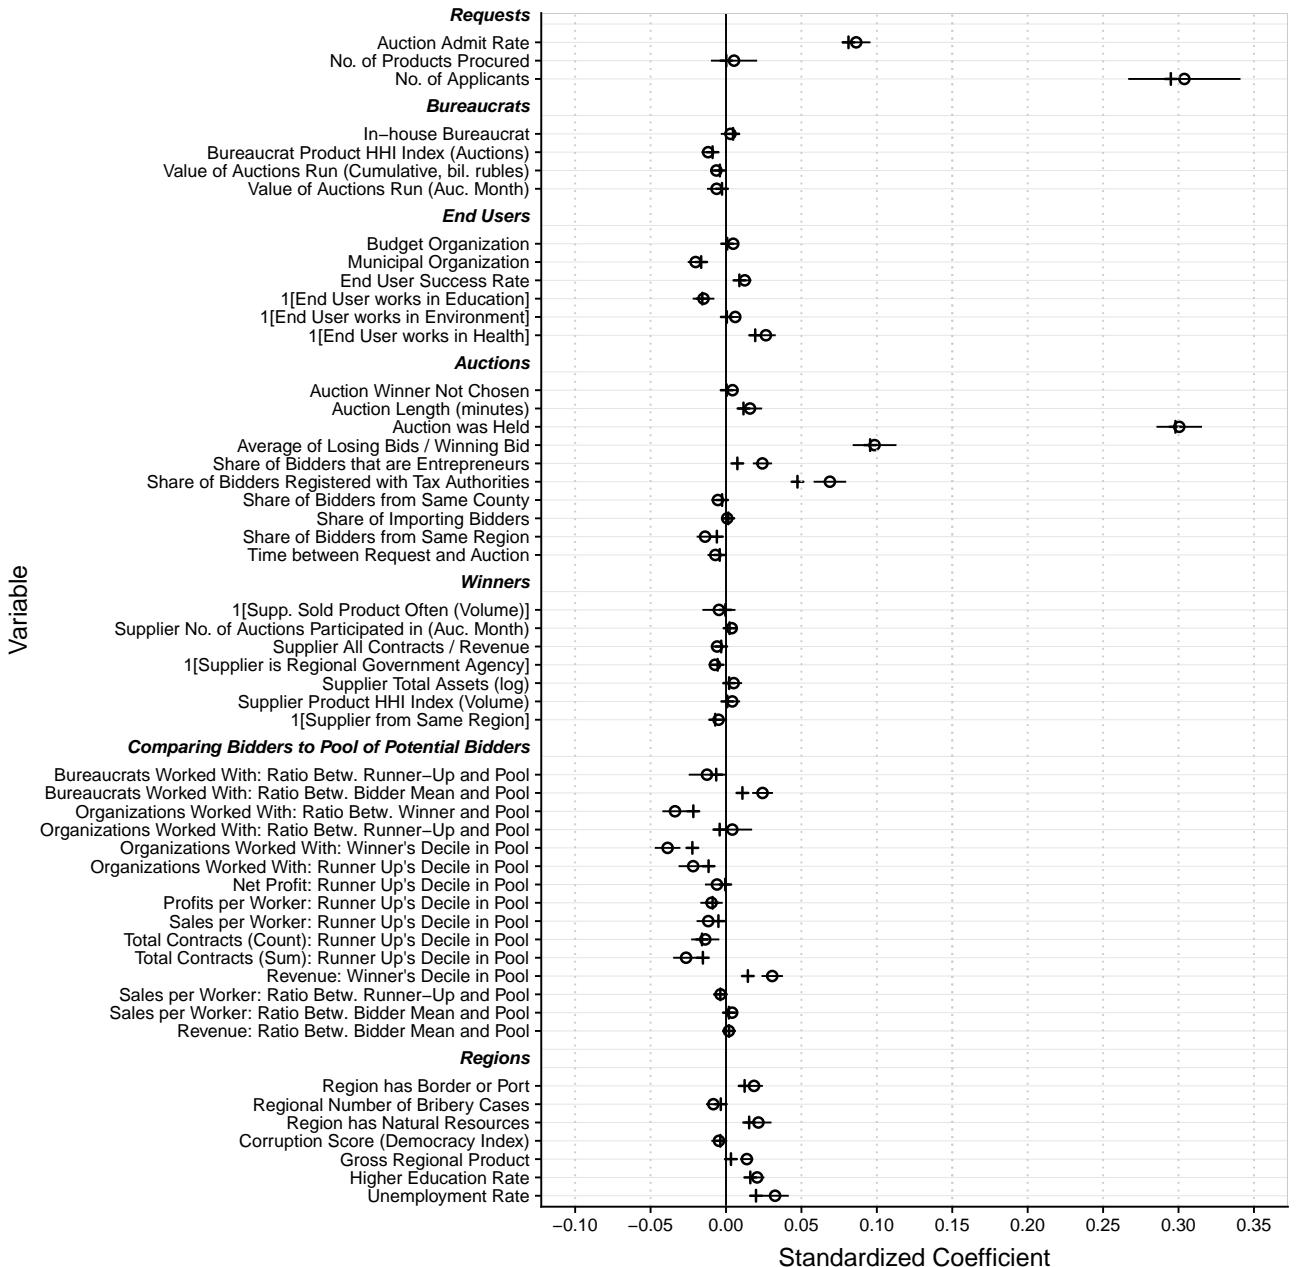
The figure shows the coefficients from the elastic net regularization procedure on the estimated bureaucrat effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. A mixing parameter of 1 represents LASSO, our baseline model. The variables shown are from the base model shown in Table 2 where the values of the regularization penalty lambda λ are chosen to return a 50 variables.

FIGURE OA.10: CORRELATES OF ORGANIZATION EFFECTIVENESS: ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS



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 3 where the values of the regularization penalty lambda λ are chosen to return 50 variables.

FIGURE OA.11: PREDICTORS OF HETEROGENEITY OF EFFECT OF BID PREFERENCES ON PARTICIPATION

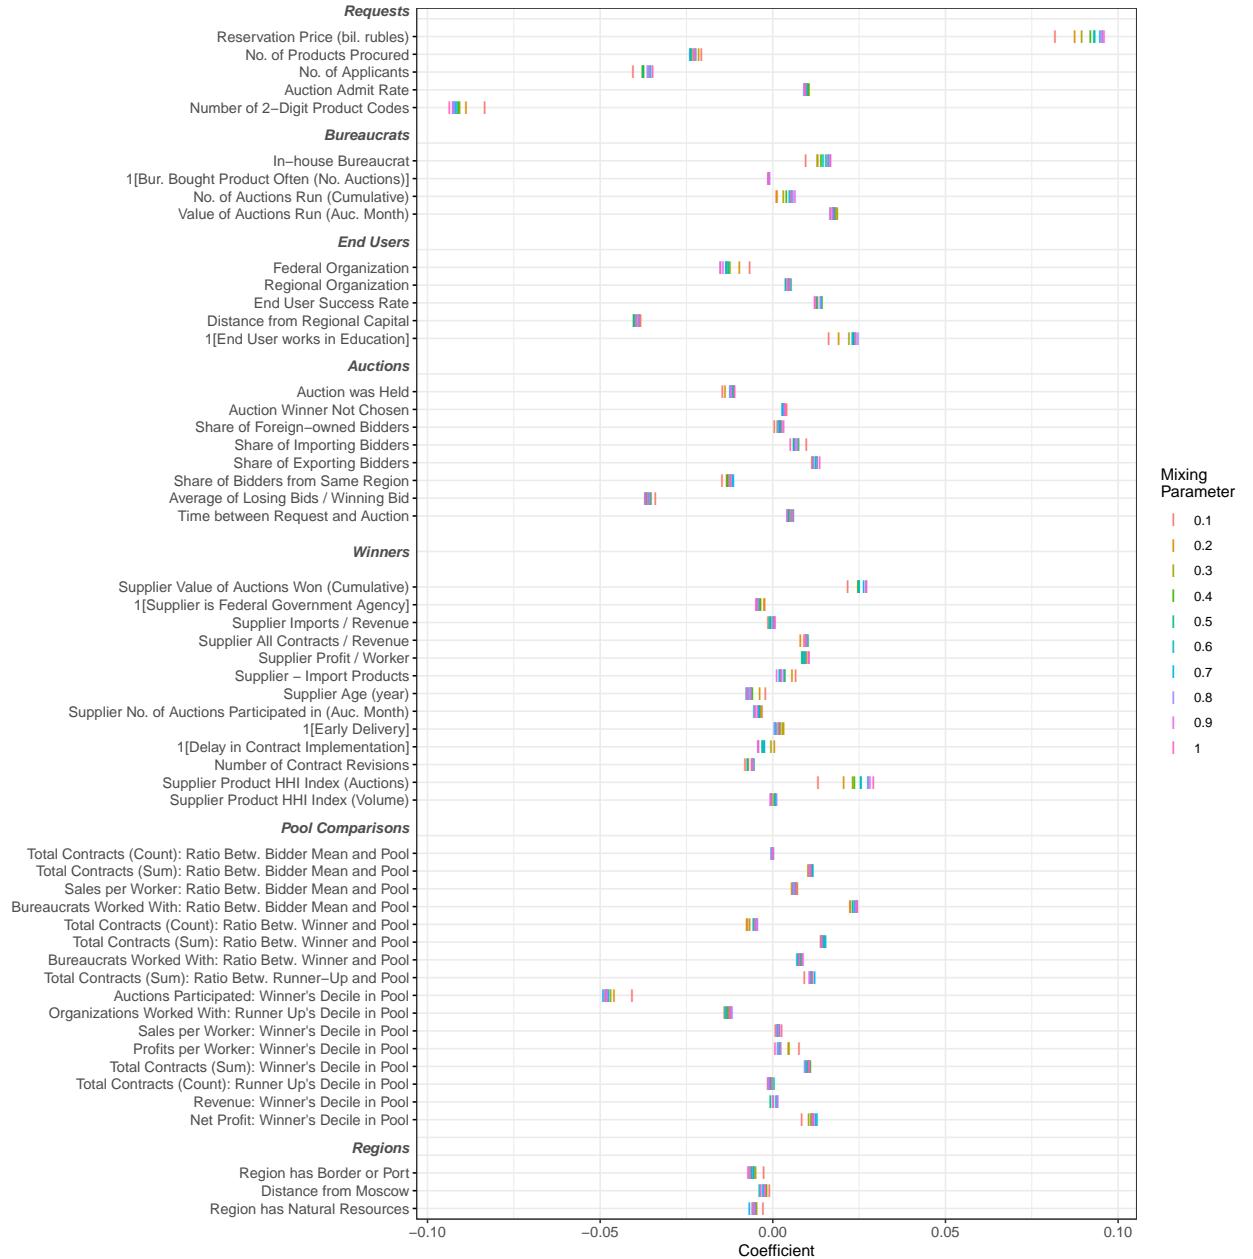


Notes: The figure shows the results of estimating our triple-differences specification for heterogeneity of the effect of bid preferences (8):

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \theta_b \hat{\alpha}_b + \theta_j \hat{\psi}_j + \delta \text{Pref}_{gt} \text{Active}_t + \gamma_b \text{Pref}_{gt} \hat{\alpha}_b + \gamma_j \text{Pref}_{gt} \hat{\psi}_j \\ + \eta_b \text{Active}_t \hat{\alpha}_b + \eta_j \text{Pref}_{gt} \hat{\psi}_j + \text{Pref}_{gt} \text{Active}_t \mathbf{Y}_{igt} \pi + \varepsilon_{igt}$$

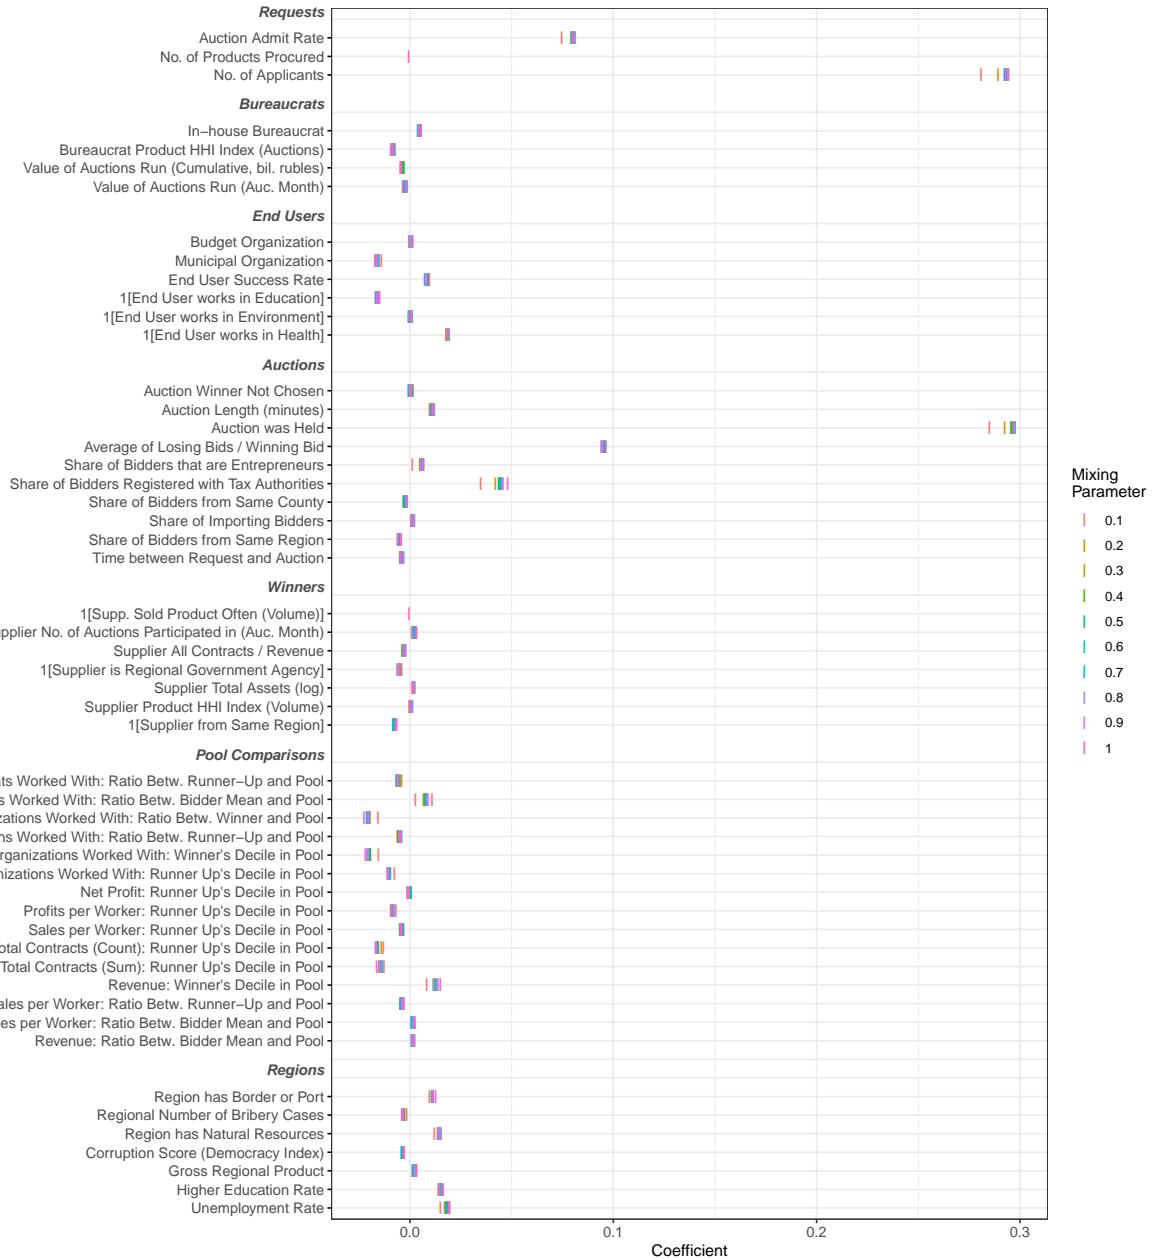
where the elements of the vector of observables \mathbf{Y}_{igt} are picked by LASSO using the largest regularization penalty that returns 50 non-zero coefficients. The coefficients from the LASSO are shown as crosses, while the circles show the coefficients and 95% confidence intervals of a multivariate regression including the 50 observables.

FIGURE OA.12: CORRELATES OF PRICE DiD: ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS



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 Table 6 where the values of the regularization penalty lambda λ is chosen to return 50 variables.

FIGURE OA.13: CORRELATES OF BIDDERS DID: ELASTIC NET REGULARIZATION COEFFICIENTS ACROSS DIFFERENT MIXING PARAMETERS



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 Table OA.11 where the values of the regularization penalty lambda λ are chosen to return 50 variables.

TABLE OA.3: 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

TABLE OA.3: PRODUCTS COVERED BY PREFERENCE LAWS, BY YEAR

2011	2012	2013	2014
	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 OA.4: 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	

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).

TABLE OA.5: EVENT STUDIES SUMMARY STATISTICS

Origin/destination Quartile*	Number of Moves (1)	Number of Observations (2)	Mean Log Residuals of Bureaucrat Movers				Mean Weeks Betw. Cols:		
			Week -1 (3)	Week 0 (4)	Week 1 (5)	Week 2 (6)	(3)-(4) (7)	(4)-(5) (8)	(5)-(6) (9)
1 to 1	5,605	240,974	-0.274	-0.359	-0.351	-0.254	12.518	28.253	12.119
	5,442	224,305	-0.187	-0.224	-0.092	-0.048	12.431	23.619	12.538
	3,393	136,756	-0.144	-0.192	0.061	0.012	13.780	28.949	13.366
	1,736	70,098	-0.144	-0.139	0.245	0.209	14.171	35.299	16.808
	5,604	229,646	-0.043	-0.094	-0.211	-0.168	13.056	27.703	13.575
	9,659	484,044	-0.034	-0.050	-0.010	-0.033	12.200	25.723	11.923
	6,122	277,266	-0.035	-0.043	0.087	0.047	12.789	28.965	14.853
	2,243	87,754	0.066	0.010	0.246	0.186	11.985	35.071	15.127
	3,173	132,081	0.015	0.003	-0.115	-0.159	15.436	26.144	12.123
	5,822	262,335	0.001	0.043	0.017	0.008	13.601	24.306	12.060
	5,608	239,169	0.030	0.089	0.117	0.112	15.337	27.186	14.642
	2,649	112,986	0.183	0.166	0.292	0.267	13.623	31.675	16.031
4 to 1	1,356	55,993	0.094	0.125	-0.131	-0.093	15.422	30.282	12.537
	1,728	73,812	0.101	0.163	0.013	0.103	15.639	29.743	12.104
	2,159	90,069	0.168	0.287	0.212	0.176	15.145	29.638	12.552
	2,490	110,435	0.348	0.385	0.390	0.342	14.998	28.177	14.932
Totals	64,789	2,827,723							

The table shows information on events in which organizations switch bureaucrats. 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 4.1. 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 ≤ 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): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. 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 4.2. 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

TABLE OA.6: SUMMARY STATISTICS - LARGEST CONNECTED SET

	All Products		Pharmaceuticals Subsample	
	(1) Analysis Sample	(2) Largest Connected Set	(3) Analysis Sample	(4) Largest Connected Set
(1) # of Bureaucrats	54,771	19,257	2,505	153
(2) # of Organizations	59,574	19,546	1,884	142
(3) # of Connected Sets	984	1	129	0
(4) # of Bureaucrats with >1 Org.	12,538	4,004	929	24
(5) # of Organizations with >1 Bur.	42,438	13,617	1,454	110
(6) # of Federal Organizations	46,708	513	26	0
(7) # of Regional Organizations	19,014	7,493	1,613	49
(8) # of Municipal Organizations	38,538	11,540	245	93
(9) # of Health Organizations	7,896	2,879	1,719	139
(10) # of Education Organizations	33,223	10,424	63	0
(11) # of Internal Affairs Organizations	867	225	3	1
(12) # of Agr/Environ Organizations	339	99	1	0
(13) # of Other Organizations	17,249	5,919	98	2
(14) # of Goods	16,223	13,952	3,863	1,440
(15) # of Regions	86	48	79	1
(16) # of Auction Requests	1,249,770	469,957	42,929	1,510
(17) Mean # of Applicants	3.6	3.61	3.03	2.79
(18) Mean # of Bidders	2.14	2.2	1.98	1.75
(19) Mean Reservation Price	25,004	35,329	44,425	14,694
(20) Quantity Mean	1,117	1,612	1,718	223
Median	25	29	45	30
SD	91,806	153,179	172,099	1,014
(21) Total Price Mean (bil. USD)	84.4	90.2	91.2	97.2
Median	4.35	4.37	6.7	6.86
SD	535	601	493	532
(22) Unit Price Mean (bil. USD)	66.5	66.7	25.4	26.7
Median	0.17	0.163	0.18	0.279
SD	23,341	14,723	265	211
(23) # of Observations	11,516,088	3,975,113	182,060	7,701
(24) Total Procurement Volume (bil. USD)	399	179	9.38	0.13

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 unpreferred 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 OA.7: SHARE OF VARIANCE OF PROCUREMENT PRICES AND PARTICIPATION EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS: LARGEST CONNECTED SET

	Fixed Effects	(s.e.)	Split Sample	(s.e.)	Shrinkage	Covariance Shrinkage
(1) s.d. of Bureaucrat Effects (across burs)	1.441	(0.0599)	1.502	(0.0383)	0.908	0.764
(2) s.d. of Organization Effects (across orgs)	1.330	(0.104)	1.384	(0.0593)	0.797	0.669
(3) s.d. of Bureaucrat Effects (across pairs)	0.930	(0.0901)	0.941	(0.051)	0.655	0.480
(4) s.d. of Organization Effects (across pairs)	1.017	(0.139)	1.005	(0.0644)	0.707	0.466
(5) Bur-Org Effect Correlation (across pairs)	-0.532	(0.0492)	-0.386	(0.0521)	-0.480	0.301
(6) s.d. of Bur + Org Effects (across pairs)	0.945	(0.0398)	0.972	(0.0255)	0.696	0.763
(7) s.d. of Bureaucrat Effects (across items)	0.841	(0.0398)	0.941	(0.051)	0.650	0.331
(8) s.d. of Organization Effects (across items)	0.921	(0.0398)	1.005	(0.0644)	0.704	0.383
(9) Bur-Org Effect Correlation (across items)	-0.728	(0.0398)	-0.386	(0.0521)	-0.687	0.246
(10) s.d. of Bur + Org Effects (across items)	0.654	(0.0499)	0.635	(0.0361)	0.538	0.564
(11) s.d. of log unit price	2.231		2.231		2.231	2.231
(12) s.d. of log unit price good, month	1.292		1.292		1.292	1.292
(13) Adjusted R-squared	0.963		0.963		0.963	0.963
(14) Number of Bureaucrats	19,257		19,257		19,257	19,257
(15) Number of Organizations	19,546		19,546		19,546	19,546
(16) Number of Bureaucrat-Organization Pairs	101,375		101,375		101,375	101,375
(17) Number of Observations	3,975,113		3,975,113		3,975,113	3,975,113

The table shows the components of the variance due to bureaucrats and organizations, estimated by implementing the variance decomposition in equation (4). The sample used is the All Products-Largest Connected Set Sample summarized in Table 1 and discussed in Sub-Section 4.2. Rows 1–2 show the s.d. of the bureaucrat and organization effects. Rows 3–6 show the components of the variance of prices across bureaucrat-organization pairs, effectively weighting the estimates in rows 1–2 by the number of pairs they appear in. Rows 7–10 show the components of the variance of prices across items, 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): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. 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 as described in Section 4.2. 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 sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates, as described in Section

TABLE OA.8: ROBUSTNESS TO USING SUBSAMPLES OF INCREASINGLY HETEROGENEOUS GOODS (KHANDELWAL (2010) MEASURE)

	Quintile 1	Quintiles 1–2	Quintiles 1–3	Quintiles 1–4	Quintiles 1–5
(1) s.d. of Bur + Org Effects Within CS (across items)	0.931	0.856	0.818	0.806	0.779
(2) s.d. of Total Bur + Org Effects (across items)	0.967	0.860	0.807	0.811	0.850
(3) s.d. of log P	2.120	2.252	2.390	2.348	2.390
(4) s.d. of log P good, month	1.300	1.302	1.355	1.392	1.378
(5) s.d. of Bur+Org Within Efs / s.d. of log P good, month	0.716	0.658	0.604	0.579	0.565
(6) s.d. of Bur+Org Total Efs / s.d. of log P good, month	0.744	0.660	0.595	0.583	0.617
(7) Sample Size	365,653	674,047	1,087,299	1,352,056	1,684,802

37

The table implements the variance decomposition in equation (4) using the estimates from equation (3): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column (5) 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 and that we can match to the scope-for-quality-differentiation ladder developed by Khandelwal (2010). Column (4) removes the quintile with the highest scope-for-quality-differentiation according to the Khandelwal (2010) ladder, Column (3) the highest two quintiles, and so on.

TABLE OA.9: RESULTS ARE ROBUST TO ALTERNATIVE CLASSIFIERS AND TRIMMING FEWER OUTLIERS

	(1)	(2)	(3)	(4)
Machine learning Method	lr	lr	svm	hm
Classification Confidence Threshold	50	50	50	50
Outlier Trimming	2.5	5	5	5
(1) s.d. of Bureaucrat Effects (across burs)	1.731	1.385	1.351	1.360
(2) s.d. of Organization Effects (across orgs)	1.575	1.209	1.225	1.203
(3) s.d. of Connected Set Effects (across CS)	1.149	0.843	0.836	0.855
(4) s.d. of Bureaucrat Effects (across items, merge)	1.104	0.747	0.719	0.748
(5) s.d. of Organization Effects (across items, merge)	1.236	0.827	0.831	0.839
(6) s.d. of Connected Set Effects (across items, merge)	0.487	0.402	0.358	0.420
(7) s.d. of Total Bur + Org Effects (across items, merge)	0.894	0.630	0.594	0.640
(8) s.d. of log unit price	2.434	2.197	2.205	2.196
(9) s.d. of log unit price good, month	1.417	1.283	1.253	1.286
(10) Adjusted R-squared	0.959	0.964	0.965	0.963
(11) Number of Bureaucrats	61,815	54,771	55,187	54,361
(12) Number of Organizations	65,204	59,574	59,685	59,146
(13) Number of Bureaucrat-Organization pairs	309,912	284,710	286,394	283,900
(14) Number of Connected Sets	1,035	984	971	972
(15) Number of Observations	12,287,649	11,516,088	11,539,042	11,527,796

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): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Column (2) replicates the findings in column (1) of table 2. Column (1) removes the top and bottom 2.5% of outlier observations for each good. Column (3) uses the Support Vector Machine classifier described in Section OA.1 instead of logistic regression. Column (4) uses the hierarchical classifier described in Section OA.1 instead of logistic regression.

TABLE OA.10: RESULTS ARE ROBUST TO ALTERNATIVE CLASSIFIER RELIABILITY THRESHOLDS

	(1)	(2)	(3)	(4)	(5)	(6)
Machine learning Method	lr	lr	lr	lr	lr	lr
Classification Confidence Threshold	45	50	55	45	50	55
Outlier Trimming	2.5	2.5	2.5	5	5	5
(1) s.d. of Bureaucrat Effects (across burs)	1.689	1.731	1.661	1.483	1.385	1.406
(2) s.d. of Organization Effects (across orgs)	1.505	1.575	1.477	1.356	1.209	1.251
(3) s.d. of Connected Set Effects (across CS)	1.087	1.149	1.065	0.969	0.843	0.888
(4) s.d. of Bureaucrat Effects (across items, merge)	0.993	1.104	0.948	0.940	0.747	0.777
(5) s.d. of Organization Effects (across items, merge)	1.069	1.236	1.069	1.030	0.827	0.895
(6) s.d. of Connected Set Effects (across items, merge)	0.465	0.487	0.396	0.495	0.402	0.449
(7) s.d. of Total Bur + Org Effects (across items, merge)	0.781	0.894	0.742	0.750	0.630	0.677
(8) s.d. of log unit price	2.445	2.434	2.434	2.214	2.197	2.197
(9) s.d. of log unit price good, month	1.428	1.417	1.417	1.302	1.283	1.283
(10) Adjusted R-squared	0.958	0.959	0.959	0.963	0.964	0.964
(11) Number of Bureaucrats	62,712	61,815	61,815	55,785	54,771	54,771
(12) Number of Organizations	66,063	65,204	65,204	60,018	59,574	59,574
(13) Number of Bureaucrat-Organization pairs	312,281	309,912	309,912	287,173	284,710	284,710
(14) Number of Connected Sets	1,038	1,035	1,035	981	984	984
(15) Number of Observations	12,337,810	12,287,649	12,287,649	11,535,202	11,516,088	11,516,088

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): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Column (5) replicates the findings in column (1) of table 2. As described in Section OA.1, column (4) deemed sentences to be correctly classified if their predicted value $\hat{y}_C(\mathbf{x}_i)$ was above the 45th percentile and their normalized cosine similarity θ_{gi} was above the 45th percentile. Column (6) uses the 55th percentile. Columns (1)–(3) are analogous to columns (4)–(6) but on the sample in which only the top and bottom 2.5% of outlier observations for each good are removed.

TABLE OA.11: BID PREFERENCES ARE MORE EFFECTIVE WHEN IMPLEMENTED BY LESS EFFECTIVE BUREAUCRATS (USING RAW FIXED EFFECTS ESTIMATES OF EFFECTIVENESS)

	All Products		Pharmaceuticals		
	Log Price (1)	No. Bidders (2)	Log Price (3)	No. Bidders (4)	Domestic Winner (5)
log Standardized Quantity	-0.330 (0.014)	0.045 (0.003)	-0.027 (0.003)	0.009 (0.003)	0.003 (0.001)
Bureaucrat FE * Preferred * Policy Active	-0.209 (0.020)	0.064 (0.025)	-0.121 (0.066)	0.366 (0.149)	0.081 (0.029)
Organization FE * Preferred * Policy Active	-0.155 (0.018)	0.058 (0.021)	-0.004 (0.063)	0.276 (0.114)	0.019 (0.020)
Constituent Terms	Yes	Yes	Yes	Yes	Yes
Month, Good FEs	Yes	Yes	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes	Yes	Yes
Outcome Mean	5.57	2.11	6.22	1.89	0.36
Observations	16,575,168	16,575,168	461,989	461,989	293,538
R ²	0.696	0.295	0.950	0.300	0.736

This table estimates the triple-differences equation (7): $p_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Preferred}_{gt} \times \text{PolicyActive}_t + \gamma\text{Preferred}_{gt} \times \hat{\alpha}_b + \zeta\text{Preferred}_{gt} \times \hat{\psi}_j + \eta\text{PolicyActive}_t \times \hat{\alpha}_b + \theta\text{PolicyActive}_t \times \hat{\psi}_j + \pi\text{Preferred}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \nu\text{Preferred}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt}$. The With Bid Preferences samples summarized in columns (3) and (6) of Table 1 are used, i.e. the combination of each Analysis Sample and the treated auctions that procurers therein carried out. Columns (1) and (3) estimate the ITT on the log price paid (P); columns (2) and (4) the ITT on the number of bidders participating in the auction (N); and Column (5) the ITT on an indicator for the winner supplying domestically made goods. In the All Products sample an item has Preferred = 1 if the type of good appears on the list of goods covered by the preferences policy for that year. In the Pharmaceuticals sample, Preferred = 1 if the drug purchased is made—by at least one supplier—both in Russia and abroad. Policy Active = 1 during the part of the relevant year that the preferences policy was in effect. Bureaucrat and Organization FEs are the raw fixed effects estimates from Section 4. Month and good fixed effects are included in all columns, as are interactions between 2-digit HS Product categories, years, region, and lot size. (We use “product” to distinguish the categories used in these interactions from the much more disaggregate goods categories used for the good fixed effects). Standard errors are clustered on month and good.

TABLE OA.12: CORRELATIVES OF BUREAUCRAT AND ORGANIZATION EFFECTIVENESS: VARIABLE DESCRIPTIONS

Auctions	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description
Auction Length (minutes)	-0.00052 (0.00049)	0.00079 (0.00069)	82.97	0	12874	Length of the auction in minutes
Auction Winner Not Chosen	-0.00094 (0.00039)	0.00022 (0.00037)	0.04	0	1	Indicator if the winner of the auction was ultimately not the supplier listed on the contract
Auction was Held	-0.00416 (0.00051)	-0.00248 (0.00057)	0.61	0	1	Indicator if the auction was held (i.e. more than one supplier was admitted to the auction)
Average of Losing Bids / Winning Bid	-0.00354 (0.00029)	-0.00057 (4e-04)	1.04	1	5	Ratio of the average of all losing bids over the final winning bid
Number of Bidders	-0.00509 (0.00044)	-0.00235 (0.00051)	1.66	0	29	Number of bidders that entered bids
Share of Bidders Registered with Tax Authorities	-0.00018 (0.00049)	0.00302 (0.00057)	0.8	0	1	Share of bidders that participated in the auction that were registered with federal tax authorities
Share of Bidders among Firms with High Profit	0.00174 (0.00048)	0.00163 (0.00062)	0.6	0	1	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	0.00097 (5e-04)	0.00238 (0.00056)	0.65	0	1	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	-0.01393 (0.00047)	0.00892 (6e-04)	0.24	0	1	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	-0.00173 (0.00066)	-0.00766 (0.00069)	0.68	0	1	Share of bidders that participated in the auction that were located in the same region as the End User
Share of Bidders that are Entrepreneurs	4e-04 (0.00048)	-0.00296 (0.00047)	0.13	0	1	Share of bidders that participated in the auction that were registered as individual entrepreneurs
Share of Exporting Bidders	0.00033 (0.00051)	0.00457 (0.00052)	0.02	0	1	Share of bidders that participated in the auction that had exporting activities
Share of Foreign-owned Bidders	-0.00106 (0.00058)	0.00233 (0.00052)	0.01	0	1	Share of bidders that participated in the auction that were foreign-owned
Share of Government Bidders	-0.00165 (0.00055)	0.00213 (0.00061)	0.02	0	1	Share of bidders that participated in the auction owned by federal, regional, or municipal governments
Share of Importing Bidders	0.00464	-0.00075	0.12	0	1	Share of bidders that participated in the auction that had importing activities

	(0.00048)	(0.00068)				
Share of Wholesaler Bidders	0.00346 (0.00051)	0.00151 (0.00051)	0.15	0	1	Share of bidders that participated in the auction that operated primarily as wholesale traders
Time between Request and Auction	-0.00272 (0.00065)	-0.00383 (0.00056)	11.14	-51	201	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 (No. Auctions)]	-0.00398 (0.00062)	-0.0029 (0.00074)	0.45	0	1	Indicator if the main product was also the most common product purchased overall by the Bureaucrat (no. auctions)
1[Bur. Bought Product Often (Volume)]	-0.00328 (0.00053)	0.00508 (0.00065)	0.16	0	1	Indicator if the main product was also the most common product purchased overall by the Bureaucrat (volume)
Bureaucrat Product HHI Index (Auctions)	-0.00229 (0.00054)	0.00536 (0.00066)	0.25	0	1	Hirschmann-Herfindahl index measuring the distribution of auctions (count) by each bureaucrat
Bureaucrat Product HHI Index (Volume)	0.00614 (0.00056)	0.01441 (7e-04)	0.33	0	1	adjusted for multiple product types index measuring total sales volume of all auctions by each bureaucrat
Bureaucrat Success Rate	0.00188 (0.00054)	0.00763 (0.00065)	0.83	0	1	adjusted for multiple product types percentage of product types requests administered by the Bureaucrat that led to a successful contract
In-house Bureaucrat	0.02565 (0.00065)	0.03171 (0.00054)	0.41	0	1	Indicator if the Bureaucrat worked directly at the End User
No. of Auctions Run (Auc. Month)	-0.01953 (0.00068)	-0.00209 (0.00048)	45.96	1	2073	Number of auctions the Bureaucrat was running simultaneously in the same month as the auction
No. of Auctions Run (Cumulative)	-0.02563 (0.00047)	-0.01379 (0.00074)	5.02	0	11.24	Number of auctions the Bureaucrat had run cumulatively to the date of the auction
Value of Auctions Run (Auc. Month)	-0.01636 (0.00045)	-0.01502 (0.00062)	17.08	5.73	24.62	Total sales volume of the auctions the Bureaucrat was running simultaneously in the same month as the auction
Value of Auctions Run (Cumulative, bil. rubles)	-0.01058 (0.00054)	-0.00267 (0.00034)	1.03	0	126.03	Total sales volume of the auctions the Bureaucrat had run cumulatively to the date of the auction
Comparing Bidders to Pool of Potential Bidders	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description
Auctions Participated: Ratio Betw. Bidder Mean and Pool	-0.00063 (0.00037)	-0.00036 (4e-04)	2.24	0	370.72	Ratio between the mean number of total auctions among all bidders in the auction and the mean no. of auctions among members of a pool of potential bidders who supplied the same product in the same region within 3 months.

Auctions Participated: Ratio Betw. Runner-Up and Pool	-0.00083 (0.00042)	-0.00024 (0.00045)	0.13	0	1	Ratio between the number of total auctions that the auction runner-up participated in and the mean no. of auctions among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Auctions Participated: Ratio Betw. Winner and Pool	-4e-04 (0.00035)	0.00013 (0.00038)	0.1	0	1	Ratio between the number of total auctions that the auction winner participated in and the mean no. of auctions among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Auctions Participated: Runner Up's Decile in Pool	-0.00012 (0.00062)	0.00021 (8e-04)	57.88	0	100	As measured using the number of auctions participated, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Auctions Participated: Winner's Decile in Pool	-0.00048 (0.00045)	-1e-05 (5e-04)	68.67	0	100	As measured using the number of auctions participated, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Bureaucrats Worked With: Ratio Betw. Bidder Mean and Pool	-0.00083 (3e-04)	-9e-05 (0.00041)	1.54	0	120.91	Ratio between the mean number of unique bureaucrats among all bidders in the auction and the mean number of unique bureaucrats among members of a pool of potential bidders who supplied the same product
Bureaucrats Worked With: Ratio Betw. Runner-Up and Pool	-0.00025 (0.00034)	-0.00019 (0.00039)	0.1	0	1	in the same region within 3 months. Ratio between the number of unique bureaucrats that the auction runner-up worked with and the mean number of unique bureaucrats among members of a pool of potential bidders who supplied the same product
Bureaucrats Worked With: Ratio Betw. Winner and Pool	-0.00065 (0.00041)	-0.00062 (0.00049)	0.13	0	1	in the same region within 3 months. Ratio between the number of unique bureaucrats that the auction winner worked with and the mean number of unique bureaucrats among members of a pool of potential bidders who supplied the same product
Bureaucrats Worked With: Runner Up's Decile in Pool	1e-05 (0.00061)	1e-05 (0.00072)	51.73	0	100	in the same region within 3 months. As measured using the number of unique bureaucrats decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months

Bureaucrats Worked With: Winner's Decile in Pool	-0.00084 (0.00044)	-8e-05 (0.00047)	63.83	0	100	As measured using the number of unique bureaucrats decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Net Profit: Ratio Betw. Bidder Mean and Pool	8e-05 (4e-05)	-0.00045 (5e-05)	2.36	-7385.2	59926.99	Ratio between the mean net profit among all bidders in the auction and the mean net profit among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Net Profit: Ratio Betw. Runner-Up and Pool	-0.00042 (0.00012)	0.00029 (4e-05)	-0.01	- 3491.44	22.4	Ratio between the net profit of the auction runner-up and the mean net profit among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Net Profit: Ratio Betw. Winner and Pool	0.00014 (7e-05)	-0.00012 (7e-05)	0.05	- 2016.23	1094.26	Ratio between the net profit of the auction winner and the mean net profit among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Net Profit: Runner Up's Decile in Pool	9e-05 (0.00059)	-8e-05 (0.00058)	40.95	0	100	As measured using net profit, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Net Profit: Winner's Decile in Pool	-7e-04 (4e-04)	0.00039 (0.00053)	58.9	0	100	As measured using net profit, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Organizations Worked With: Ratio Betw. Bidder Mean and Pool	-9e-04 (0.00028)	4e-05 (0.00039)	1.73	0	115.09	Ratio between the mean number of unique organizations among all bidders in the auction and the mean number of unique organizations among members of a pool of potential bidders who supplied the same product
Organizations Worked With: Ratio Betw. Runner-Up and Pool	-2e-05	-0.00038	0.12	0	1	in the same region within 3 months. Ratio between the number of unique organizations that the auction runner-up worked with and the mean number of unique organizations among members of a pool of potential bidders who supplied the same product in the same region within 3 months.

	(0.00037)	(0.00042)				
Organizations Worked With: Ratio Betw. Winner and Pool	-0.00057 (4e-04)	-0.00024 (0.00047)	0.17	0	1	Ratio between the number of unique organizations that the auction winner worked with and the mean number of unique organizations among members of a pool of potential bidders who supplied the same product
Organizations Worked With: Runner Up's Decile in Pool	0.00022 (0.00063)	-0.00026 (6e-04)	54.69	0	100	in the same region within 3 months. As measured using the number of unique organizations decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Organizations Worked With: Winner's Decile in Pool	-0.00052 (0.00041)	-0.00031 (0.00049)	65.47	0	100	As measured using the number of unique organizations decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Profits per Worker: Ratio Betw. Bidder Mean and Pool	8e-05 (4e-05)	-0.00012 (4e-05)	1.46	- 10111.73	73924.35	Ratio between the mean profits per worker among all bidders in the auction and the mean profits per worker among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Profits per Worker: Ratio Betw. Runner-Up and Pool	4e-05 (8e-05)	1e-04 (9e-05)	0.03	-1615	26.6	Ratio between the profits per worker of the auction runner-up and the mean profits per worker among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Profits per Worker: Ratio Betw. Winner and Pool	1e-04 (0.00015)	-0.00051 (0.00028)	0.09	-958.37	737.16	Ratio between the profits per worker of the auction winner and the mean profits per worker among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Profits per Worker: Runner Up's Decile in Pool	9e-05 (0.00062)	-0.00011 (0.00052)	38.54	0	100	As measured using profits per worker, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months

Profits per Worker: Winner's Decile in Pool	-0.00062 (0.00046)	5e-05 (0.00057)	57.04	0	100	As measured using profits per worker, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Revenue: Ratio Betw. Bidder Mean and Pool	-0.00039 (0.00034)	1e-04 (0.00048)	1.54	0	538.22	Ratio between the mean revenue among all bidders in the auction and the mean revenue among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Revenue: Ratio Betw. Runner-Up and Pool	-9e-05 (0.00029)	-6e-05 (0.00031)	0.03	0	1	Ratio between the revenue of the auction runner-up and the mean revenue among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Revenue: Ratio Betw. Winner and Pool	-0.00037 (0.00039)	-0.00028 (0.00044)	0.06	0	1	Ratio between the revenue of the auction winner and the mean revenue among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Revenue: Runner Up's Decile in Pool	-0.00056 (0.00059)	1e-05 (7e-04)	43.46	0	100	As measured using revenue, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Revenue: Winner's Decile in Pool	-0.00061 (0.00039)	0.00023 (5e-04)	59.23	0	100	As measured using revenue, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Sales per Worker: Ratio Betw. Bidder Mean and Pool	-8e-05 (0.00038)	-0.00059 (0.00047)	1.19	0	197.12	Ratio between the mean sales per worker among all bidders in the auction and the mean sales profits per worker among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Sales per Worker: Ratio Betw. Runner-Up and Pool	-0.00018 (0.00039)	-0.00026 (0.00036)	0.08	0	1	Ratio between the sales per worker of the auction runner-up and the mean profits per worker among members of a pool of potential bidders who supplied the same product in the same region within 3 months.

Sales per Worker: Ratio Betw. Winner and Pool	-1e-04 (0.00044)	-0.00061 (0.00041)	0.11	0	1	Ratio between the sales per worker of the auction winner and the mean sales per worker among members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Sales per Worker: Runner Up's Decile in Pool	3e-04 (0.00065)	-0.00081 (0.00068)	42.58	0	100	As measured using sales per worker, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Sales per Worker: Winner's Decile in Pool	-0.00114 (0.00049)	7e-05 (0.00057)	58.63	0	100	As measured using sales per worker, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
Total Contracts (Count): Ratio Betw. Bidder Mean and Pool	2e-05 (0.00027)	-0.00038 (0.00034)	1.86	0	84.39	Ratio between the mean number of contracts won by all bidders in the auction and the mean number of contracts won by members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Total Contracts (Count): Ratio Betw. Runner-Up and Pool	0.00016 (0.00028)	-0.00024 (0.00043)	0.06	0	1	Ratio between the total number of contracts won by the auction runner-up and the mean number of contracts won by members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Total Contracts (Count): Ratio Betw. Winner and Pool	-0.00036 (0.00033)	-0.00071 (0.00037)	0.15	0	1	Ratio between the total number of contracts won by the auction winner and the mean number of contracts won by members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Total Contracts (Count): Runner Up's Decile in Pool	-3e-05 (0.00053)	-0.00049 (0.00061)	43.56	0	100	As measured using total number of contracts won, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Total Contracts (Count): Winner's Decile in Pool	-0.00013	-0.00098	60.65	0	100	As measured using total number of contracts won, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months

	(0.00047)	(5e-04)				
Total Contracts (Sum): Ratio Betw. Bidder Mean and Pool	-5e-05 (0.00025)	-0.00016 (0.00031)	1.15	0	260.49	Ratio between the mean volume of contracts won by all bidders in the auction and the mean volume of contracts won by members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Total Contracts (Sum): Ratio Betw. Runner-Up and Pool	-5e-05 (0.00023)	-0.00017 (0.00034)	0.03	0	1	Ratio between the total volume of contracts won by the auction runner-up and the mean volume of contracts won by members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Total Contracts (Sum): Ratio Betw. Winner and Pool	0.00012 (3e-04)	-0.00068 (0.00041)	0.08	0	1	Ratio between the total volume of contracts won by the auction winner and the mean volume of contracts won by members of a pool of potential bidders who supplied the same product in the same region within 3 months.
Total Contracts (Sum): Runner Up's Decile in Pool	-6e-05 (0.00049)	-0.00037 (0.00055)	39.26	0	100	As measured using total volume of contracts won, decile of the runner-up within the pool of potential bidders who supplied the same product in the same region within 3 months
Total Contracts (Sum): Winner's Decile in Pool	0.00016 (0.00042)	-0.00129 (0.00049)	56.89	0	100	As measured using total volume of contracts won, decile of the winner within the pool of potential bidders who supplied the same product in the same region within 3 months
End Users	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description
1[End User works in Agriculture]	0.00187 (0.00037)	-0.0024 (0.00033)	0	0	1	End User works in the agricultural sector
1[End User works in Culture]	0.00695 (0.00064)	0.00433 (0.00123)	0.01	0	1	End User works on cultural affairs
1[End User works in Education]	0.01488 (0.00071)	0.00225 (0.00063)	0.23	0	1	End User works in education
1[End User works in Emergency Services]	-0.00098 (0.00041)	0.00024 (6e-04)	0.01	0	1	End User works in emergency services

1[End User works in Environment]	0.00549 (0.00074)	0.0102 (0.00141)	0.01	0	1	End User works in the environmental sector
1[End User works in Forestry]	0.00047 (3e-04)	-0.00014 (0.00035)	0	0	1	End User works in the forestry sector
1[End User works in Health]	-0.0225 (0.00062)	-0.02386 (0.00054)	0.56	0	1	End User works in the health care sector
1[End User works in Housing]	-0.00218 (0.00066)	0.00503 (0.00066)	0.01	0	1	End User works in the housing sector
1[End User works in Internal Affairs]	0.00445 (5e-04)	0.02509 (0.00076)	0.01	0	1	End User works in internal affairs (police, justice, etc.)
1[End User works in Labor]	-0.00047 (0.00025)	-0.00103 (0.00033)	0	0	1	End User works in the labor sector (re-training, unemployment assistance, etc.)
1[End User works in Mining]	-0.00022 (4e-05)	-0.00171 (3e-05)	0	0	1	End User works in the mining sector
1[End User works in Natural Resources]	0.00027 (0.00011)	-9e-05 (0.00013)	0	0	1	End User works in the natural resources sector
1[End User works in News]	7e-05 (0.00013)	0.00032 (0.00018)	0	0	1	End User works in news and journalism
1[End User works in Social Policy]	-0.00787 (0.00028)	-0.00603 (0.00034)	0.05	0	1	End User works on social policy (welfare, pensions, etc.)
1[End User works in Sport]	0.00072 (0.00046)	0.00409 (0.00038)	0	0	1	End User works in the sport and recreational sector
1[End User works in Television]	-0.00061 (0.00015)	0.00021 (0.00019)	0	0	1	End User works in television and mass communications
1[End User works in Transportation]	0.007 (6e-04)	-0.00273 (0.00055)	0.01	0	1	End User works in the transportation sector
1[End User works in Veterinary Affairs]	-0.00137 (0.00022)	-6e-05 (0.00033)	0	0	1	End User works in veterinary affairs
1[End User works in Youth Services]	-0.00164 (0.00036)	0.00045 (0.00036)	0	0	1	End User works in youth services
Autonomous Organization	-0.0066 (0.00039)	0.00088 (0.00046)	0.03	0	1	End User is a non-commercial organization created by the government that enjoys more financial autonomy
Budget Organization	0.00033 (0.00029)	-0.00229 (0.00048)	0.01	0	1	Non-commercial organization with less financial autonomy and stricter budget control from government owner

Distance from Regional Capital	0.04038 (5e-04)	-0.02901 (8e-04)	6.98	0	16.16	Distance between the End User and the capital of the region where it is located (log kilometers)
End User Performance Score	-0.00786 (0.00081)	-0.00939 (0.00049)	51.6	0	209	Total performance score for the End User from independent surveys and evaluations by the Federal Treasury
End User Success Rate	-0.00027 (7e-04)	0.01308 (0.00048)	0.82	0	1	Percentage of requests administered for the End User that led to a successful contract
Federal Organization	0.04687 (7e-04)	0.04562 (0.00089)	0.12	0	1	End User receives funds from the federal government and operates on the federal level
Government Agency	-0.00124 (0.00021)	0.0015 (0.00024)	0	0	1	End User is classified as a separate government agency, operating more independent of government oversight
Municipal Organization	0.01528 (0.00063)	-0.01969 (0.00055)	0.27	0	1	End User receives funds from the municipal government and operates on the municipal level
Other Government Body	0.00604 (0.00039)	-0.00038 (0.00047)	0.95	0	1	End User has a much less common legal classification, such as a natural monopoly, audit agency, etc.
Regional Organization	-0.04381 (0.00057)	-0.01324 (0.00062)	0.61	0	1	End User receives funds from the regional government and operates on the regional level
Regions	PwCorr BurFE	PwCorr OrgFE	Mean	Min	Max	Description
Corruption Score (Democracy Index)	0.00015 (0.00088)	0.00156 (0.00093)	2.72	1	5	Regional corruption score measured on scale of 1-5 (higher values indicating less corruption)
Distance from Moscow	4e-04 (0.00061)	6e-05 (6e-04)	6.6	0	9.38	from Petrov and Titkov (2013) Distance between the region where the End User is located and Moscow
Gross Regional Product	-0.00175 (0.00076)	0.00024 (0.00086)	13.4	10.18	20.55	Gross regional product, averaged over the analysis period (log)
Higher Education Rate	-0.00043 (6e-04)	0.00077 (0.00067)	28.37	0	47.9	Higher Education rate, averaged over the analysis period (log)
Public Perceptions of Corruption	-0.00017 (0.00099)	0.00021 (0.00088)	16.91	5	35	Public perception of the severity of corruption as measured by a
Region has Border or Port	-0.00023 (0.00088)	0.00116 (0.00082)	0.5	0	1	Indicator if the region in which the auction was held had a border or port
Region has Natural Resources	3e-04 (0.00103)	0.00092 (0.001)	0.27	0	1	Indicator if the region in which the auction was held had natural resources such as oil or gas
Regional Number of Bribery Cases	-2e-04 (0.00076)	-1e-05 (0.00063)	48.24	0	161.69	Number of bribery cases filed by officials in the region in which the auction was held

Regional Number of Corruption Cases	-0.00097 (7e-04)	-0.00012 (0.00058)	119.59	0	446	Number of corruption cases filed by officials in the region in which the auction was held
Regional Number of Corruption Convictions	8e-05 (0.00086)	-0.00073 (0.00084)	12.93	0	34.56	Number of corruption convictions secured by officials in the region in which the auction was held
Regional Number of Major Corruption Convictions	-9e-05 (0.00082)	-0.00056 (0.00094)	12.19	0	28.42	Number of major corruption convictions secured by officials in the region in which the auction was held
Total Democracy Index	-0.00098 (0.00089)	-0.00061 (0.00089)	31.68	16	43	Total democracy score measured on a scale of 1-50 (higher values indicating more democracy)
Requests	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description and Titkov (2013)
Auction Admit Rate	-0.00123 (0.00053)	0.00064 (0.00063)	0.94	0.04	1	Percentage of supplier applicants admitted to auction
Deposit (ths. Rubles)	0.00082 (0.00083)	0.00402 (0.00142)	3	0	11141.39	Amount bidders are required to deposit before entering auction
No. of Applicants	-0.00839 (0.00046)	-0.00618 (0.00055)	3.18	1	80	Number of suppliers that submitted applications to participate in the auction
No. of Products Procured	-0.00113 (0.00041)	-0.00851 (0.00078)	35.14	1	6472	Number of products overall
Number of 2-Digit Product Codes	0.00393 (0.00067)	-0.00349 (8e-04)	1.91	1	18	Number of unique products (as measured by their two-digit codes)
Number of Request Revisions	-0.0012 (0.00035)	-0.00297 (0.00041)	0.05	0	18	Number of revisions that the Bureaucrat made to the contract before it was finalized
Reservation Price (bil. rubles)	0.00937 (0.00066)	0.00925 (0.00057)	13.17	0	22.33	Amount of Reservation price in billions of rubles
Winners	PwCorr - BurFE	PwCorr - OrgFE	Mean	Min	Max	Description
1[Delay in Contract Implementation]	0.00174 (0.00052)	-0.0013 (0.00044)	0.07	0	1	Indicator if the contract implementation was delayed
1[Early Delivery]	-0.00529 (0.00049)	0.00191 (6e-04)	0.18	0	1	Indicator if the contract implementation was early
1[Supp. Sold Product Often (No. Auctions)]	-0.00228 (0.00063)	-0.0021 (0.00081)	0.49	0	1	Indicator if the main product was also the most common product supplied overall by the Supplier (no. auctions)
1[Supp. Sold Product Often (Volume)]	-0.0025 (0.00071)	-0.00286 (0.00068)	0.42	0	1	Indicator if the main product was also the most common product supplied overall by the Supplier (volume)

1[Supplier Above Median Profit]	0.00383 (0.00065)	0.00054 (0.00072)	0.61	0	1	Indicator if the Supplier has above-median profit relative to the other suppliers in the dataset
1[Supplier Above Median Revenue]	0.00368 (7e-04)	0.00131 (0.00077)	0.66	0	1	Indicator if the Supplier has above-median revenue relative to the other suppliers in the dataset
1[Supplier from Same Region]	-0.00174 (0.00062)	-0.0068 (0.00058)	0.7	0	1	Indicator if the Supplier is located in the same region as the End User
1[Supplier has Foreign Ownership]	-5e-04 (0.00052)	0.00235 (5e-04)	0.01	0	1	Indicator if the Supplier has foreign ownership
1[Supplier is Exporter]	0.00053 (0.00049)	0.00481 (0.00044)	0.02	0	1	Indicator if the Supplier has exporting activities
1[Supplier is Federal Government Agency]	0.00229 (0.00062)	0.00284 (0.00066)	0	0	1	Indicator if the Supplier is registered as a federal government agency
1[Supplier is Importer]	0.00384 (0.00046)	-0.00022 (0.00067)	0.13	0	1	Indicator if the Supplier has importing activities
1[Supplier is NGO]	0.00055 (0.00031)	1e-05 (0.00038)	0	0	1	Indicator if the Supplier is a nongovernmental organization
1[Supplier is Regional Government Agency]	-0.00358 (0.00043)	-0.00042 (0.00043)	0.01	0	1	Indicator if the Supplier is registered as a regional government agency
Contract was Terminated	0.00069 (0.00044)	2e-05 (0.00037)	0.02	0	1	Indicator if the contract was terminated before being completely implemented
Days Contract Implementation Delayed	0.00045 (4e-04)	0.00022 (0.00037)	0.25	0	73	Number of days early that the contract was implemented
Days Contract Implementation Early	-0.00078 (0.00065)	0.00048 (0.00045)	0.6	0	29	Number of days late that the contract was implemented
Number of Contract Revisions	-0.00014 (0.00043)	0.00762 (0.00069)	1.28	1	58	Number of revisions that the Bureaucrat made to the request before it was finalized
Problem with Contract Implementation	0.00043 (0.00072)	-0.00013 (4e-04)	0	0	1	Indicator if a problem arose during the contract's implementation
Supplier - Export Products	0.0021 (0.00047)	-0.00069 (0.00041)	0.7	0	290	Number of unique products the Supplier exports
Supplier - Import Products	0.00336 (0.00047)	0.00496 (0.00053)	3.44	0	893	Number of unique products the Supplier imports
Supplier Age (year)	-0.00145 (0.00092)	-0.00162 (9e-04)	2.24	0	5.44	Number of years Supplier has registered as a legal entity (log)

Supplier All Contracts / Revenue	0.00334 (0.00047)	0.00027 (0.00049)	0.2	0	9.99	Ratio of Supplier's total contract volume to revenue
Supplier Exports / Revenue	0.00019 (3e-05)	9e-05 (3e-05)	0.01	0	2540.02	Ratio of Supplier's total export volume to revenue
Supplier Imports / Revenue	0.00178 (0.00038)	-0.00398 (0.00057)	0.01	0	1	Ratio of Supplier's total import volume to revenue
Supplier No. Employees (log)	0.00213 (0.00082)	0.00045 (0.00068)	3	0	13.04	Number of employees working at the Supplier (log)
Supplier No. of Auctions Participated in (Auc. Month)	-0.00737 (0.00058)	0.00683 (0.00065)	37.3	1	3884	Number of auctions the Supplier was participating in simultaneously in the same month as the auction
Supplier No. of Auctions Participated in (Cumulative)	-0.00626 (0.00056)	0.00406 (0.00057)	525.04	0	79131	Number of auctions the Supplier had participated in cumulatively to the date of the auction
Supplier Number of Countries Exported to	1e-04 (0.00039)	0.00214 (0.00036)	0.21	0	104	Number of unique countries that the Supplier exported to
Supplier Product HHI Index (Auctions)	-0.00895 (0.00062)	-0.00256 (0.00079)	0.54	0	1	Hirschmann-Herfindahl index measuring number of auctions (count) won by supplier across two-digit product types
Supplier Product HHI Index (Volume)	-0.00319 (0.00051)	8e-04 (0.00058)	0.59	0	1	Hirschmann-Herfindahl index measuring sales volume of all auctions won by supplier across two-digit product types
Supplier Profit	-0.00029 (0.00061)	0.00473 (0.00054)	2.52	-11.63	13.61	Supplier net profit
Supplier Profit / Worker	-0.00371 (0.00051)	0.00578 (0.00059)	0.67	-7.71	8.43	Supplier profits per worker
Supplier Revenue (log)	0.00403 (0.00119)	0.00936 (0.00115)	5.66	0	15.87	Supplier revenue (log)
Supplier SOE Contracts / Revenue	-0.00102 (0.00047)	-7e-05 (0.00036)	0.03	0	9.59	Ratio of Supplier's total volume of contracts with state-owned enterprises to revenue
Supplier Total Assets (log)	-0.00107 (0.00078)	0.00413 (7e-04)	4.01	0	15.93	Supplier total assets (log)
Supplier Value of Auctions Won (Cumulative)	0.00266 (0.00052)	0.00189 (0.00058)	16.25	3	24.55	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)	-0.00472	0.00763	0.72	0	385	Total sales volume of the auctions the Supplier had participated in cumulatively to the date of the auction

	(0.00045)	(0.00044)				
Supplier is Individual Entrepreneur	-0.00051 (0.00039)	-0.00246 (4e-04)	0.13	0	1	Indicator if the Supplier is registered as an individual entrepreneur
Supplier is Registered with Tax Authorities	0.00076 (0.00045)	0.00094 (0.00047)	0.76	0	1	Indicator if the Supplier is registered with the tax authorities

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. 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.

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.