Individuals and Organizations as Sources of State Effectiveness*

Michael Carlos Best† Jonas Hjort‡ David Szakonyi§

December 21, 2018

Abstract

Policymakers do not implement states' policies—bureaucrats do. How important are bureaucrats in determining the productivity of the state enterprise? To what extent do the tradeoffs between different policies—and hence optimal policy design—depend on the effectiveness of the bureaucracy tasked with implementation? We investigate these questions in the context of public procurement. Using data on 16 million purchases in Russia during 2011–2016, we first 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. Our estimates imply that ineffective bureaucracies massively reduce public sector output: moving the least effective quartile of procurers to 75th percentile effectiveness would save the Russian government USD 13 billion each year. To explore the implications of bureaucratic effectiveness for policy design, we analyze Russia's adoption of a ubiquitous procurement policy—bid preferences for domestic suppliers. Using a generalized difference-in-differences strategy, we estimate the impact of the policy separately for procurers of different effectiveness. Consistent with a simple endogenous-entry auction model, we find that bid preferences save the government 17.5 percent when implemented by the *least* effective quartile of bureaucrats, but only 0.7 percent when implemented by the most effective quartile. Overall, our results demonstrate that the often overlooked bureaucratic apparatus is critical for state effectiveness, which helps to explain why many policies work well only in some settings; and that policy designed with bureaucratic context in mind can partially offset the cost of bureaucratic ineffectiveness.

JEL codes: H11, O17, O2

---

*We are especially grateful to Ben Olken for his very helpful discussion of the paper at the AEA 2018 meetings. We also thank John M. Abowd, Daron Acemoglu, Tim Besley, Nick Bloom, Richard Blundell, Raj Chetty, Allan Collard-Wexler, Francesco Decarolis, Andres Drenik, Pascaline Dupas, Ben Faber, Fred Finan, Matthew Gentzkow, Josh Gottlieb, Caroline Hoxby, Amit Khandelwal, Brad Larsen, David Margolis, Torsten Persson, Andrea Prat, Imran Rasul, Jimmy Roberts, Jonah Rockoff, Tobias Salz, Orie Shelef, Andrei Shleifer, David Silver, Eric Verhoogen, Guo Xu, and Danny Yagan for valuable comments; Andrei Yakovlev for institutional guidance; Georgiy Syunyaev and Andrey Druzhaev for outstanding research assistance; and seminar participants at the AEA 2018 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 based 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, 2010; Acemoglu *et al.* , 2015; Page & Pande, 2018). A key feature of states is that policy implementation is delegated to a bureaucracy—the state's middle management tier. Some bureaucracies are equipped with the skills and organizational capacity to implement complex policy directives, while others are not; some bureaucracies do not share policymakers' priorities, while others do. In sum, the individuals and organizations who make up the bureaucratic apparatus may play an important role in determining the productivity of the state.

Our goals in this paper are twofold. First, to quantify the importance of the bureaucracy for public sector output. Second, to explore how the tradeoffs between different policies depend on the effectiveness of the bureaucracy in charge of implementation.[1] These questions are challenging to address in part because bureaucracies produce a wide array of outputs, many of which cannot be measured in public-sector wide, administrative data. However, one particular 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. Using an empirical specification derived from the model, we estimate that over 40 percent of the variation in prices paid is attributable to the procurers who manage procurement processes—individual procurement officers and the public organizations paying for and using the goods—and that individuals and organizations each contribute roughly half to this. Differences in effectiveness of such magnitude may have far-reaching implications for policy design. To provide a concrete example, we study the introduction of a bid preference regime common throughout the world. Under Russia's bid preference policy, contract winners offering goods manufactured abroad are paid only 85 percent of their bid. Our model predicts that the impact of such a differential will vary systematically with procurement managers' effectiveness. Consistent with the model, we estimate that cost savings and the increase in the likelihood of domestic producers winning the contract are both considerably larger when the policy is implemented by *less* effective bureaucrats. The magnitude of the differences in the policy's impact we estimate suggest that the gains from tailoring policy to the effectiveness of the implementing bureaucracy can be large.

[1] These two goals are closely related, as states can increase bureaucratic effectiveness either by directly improving the productivity of the bureaucracy or by optimizing the policies the bureaucracy is asked to implement. The answers depend on the state's "production function", an object that remains almost entirely unknown. Existing evidence does offer clues, however. First, a growing body of research has documented dramatic differences in abilities and practices across individual public sector workers and across the organizations they work for (see e.g. Bloom *et al.* , 2015; Finan *et al.* , 2017; Khan *et al.* , 2016, 2018). What remains unclear is the extent to which such differences influence public sector output, and how this varies with the policies in use (for evidence that front-line public sector workers such as teachers and health and community workers do, at least to some extent, influence public sector output, see the excellent overview of the literature by Finan *et al.* (2017)). Second, empirically, many policies appear to work well in some countries or regions and poorly in others (Rodrik, 2009). For example, value-added taxes generate the intended tax compliance in most developed countries, but rarely do so in developing countries (Bird & Gendron, 2007); the NREGA employment guarantee scheme supports poor workers and helps complete important infrastructure projects in some Indian states, but is largely unused in others (Gulzar & Pasquale, 2017); and the postal services in Algeria, Barbados, and Uruguay comply with the policy of returning incorrectly addressed letters to sender, but the ones in Cambodia, Russia, and Tajikistan do not (Chong *et al.* , 2014).

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

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 and less diverse participants, and pay higher quality-adjusted prices.

Introducing bid preferences favoring locally manufactured goods that are, on average, more expensive in turn also has two impacts. First, penalizing bids by foreign bidders makes them less competitive, tending to increase prices paid. Second, such differential penalization makes participation less attractive for foreign bidders and more attractive for local bidders. When state effectiveness is high, baseline participation is high and so the participation effect is small and procurement outcomes worsen. However, when state effectiveness is low, baseline participation is low and so the likelihood that a local bidder who enters has to face a more efficient, foreign, bidder is low. As a result, the introduction of bid preferences has a large impact on the likelihood that a local bidder can win the contract, and hence on their willingness to participate. In this case the participation effect dominates and overall procurement outcomes improve.

To compare the performance of bureaucrats and organizations 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 end-user organizations (e.g. ministries, schools or hospitals) are observed working with multiple bureaucrats (procurement officers) and vice versa, providing us with thousands of quasi-experiments for identification. Event studies reveal sharp decreases in prices paid for a given good when organizations switch to more effective bureaucrats (for example, prices decrease by 18 percent when an organization switches from a bureaucrat in the worst quartile of average prices to a bureaucrat in the best quartile), and corresponding increases when

[2] Russia spends over half of its total public procurement budget on such goods. Unlike in most low- and middle-income countries, detailed administrative datasets covering all levels of government and the entire country are publicly available.

[3] Our methodology ensures that within-category quality differences are minimal, while maintaining generality by not restricting the sample to very specific types of goods. The difficulty of categorizing goods accurately so as to ensure like-for-like comparison has long dogged several literatures. In foregoing conventional methods 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.

organizations switch is to a less effective bureaucrat. The event studies also provide clear evidence supporting a causal interpretation of these effects.[4]

To aggregate the impacts of individual bureaucrats and organizations on prices into an estimate of the share of the total variation in priced paid that is explained by the bureaucratic apparatus, 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 (Neyman & Scott, 1948; Lancaster, 2000). To do so, we adapt split-sample (see e.g. Finkelstein *et al.* , 2016; Silver, 2016) and shrinkage (Kane & Staiger, 2008; Chetty *et al.* , 2014a) methods to a two-dimensional context, explicitly accounting for the covariance between the estimation error in the bureaucrat and the organization dimension (Andrews *et al.* , 2008).[5] Second, we address the missing links that arise because we estimate performance effects attributable to workers and organizations engaged in the same task. That is, the switches that identify individuals' and organizations' impacts on prices do not link *all* individuals and organizations in our setting.[6] We show that (i) the combined productivity effect of pairs of workers and organizations can nevertheless be identified, and (ii) a suitable normalization of the effects allows estimation of lower bounds on the shares of the variation in performance explained by bureaucrats and organizations.

We find that the individuals and organizations of the bureaucracy together account for more than 40 percent of the variation in prices paid, of which individuals and organizations account for roughly equal shares. These results, from the baseline policy regime that treats all suppliers' bids equally, imply that, in weak institutional contexts such as Russia, the impact on public sector output of improving micro level state effectiveness would be massive. 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 next correlate our estimates of procurer effectiveness with a rich set of indicators on each procurer's auctions (see also Lacetera *et al.* , 2016)—measures of initial entry barriers, how the auction was executed, procurer experience, etc. We find, consistent with our theoretical framework, that lowering the participation cost imposed on suppliers is a key part of what enables procurers to pay low prices: effective procurers set lower reservation prices, attract more applicants, and revise contracts less frequently.

The second part of the paper focuses on the implications of the heterogeneity in effectiveness we document in the first part for the design of policy. To shed light on how the costs and benefits of specific policies depend on the effectiveness of policy implementers in the bureaucracy, we focus on a particular

[4]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 these assumptions, as does a battery of additional tests. Studies of the wages of workers and firms in the private sector from the labor economics literature tend to find the same (see e.g. Mendes *et al.* , 2010; Card *et al.* , 2013, 2016; Goldschmidt & Schmieder, 2017; Shelef & Nugyen-Chyung, 2015; Alvarez *et al.* , 2018; Bloom *et al.* , 2018, forthcoming). In the public sector there are additional institutional reasons to expect these assumptions to hold (see Section 2).

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

[6]Worker mobility is often high enough that almost all workers and firms belong to the biggest "connected set" if the outcome being compared across job spells is wages, as in the existing literature. This paper is to our knowledge the first to estimate performance effects attributable to workers and organizations engaged in the same task. Workers and organizations performing the same task form a less connected network.

policy change in Russia—the introduction of bid preferences for domestically manufactured goods. Bid preferences are a ubiquitous policy tool used to favor local bidders in public procurement purchases.[7] 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 show 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, in that average prices paid if anything decrease slightly.[8]

To test our model's predictions that bid preferences will have better impacts when implemented by less effective bureaucracies, 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 treatment 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.).[9] This dependence of policy impact on state effectiveness may be part of the reason why some policies work well in some countries or regions and poorly in others. It also suggests that tailoring the design of policy to the effectiveness of the implementing bureaucracy may be very advantageous, and can partly offset the costs of a less effective bureaucracy.

This paper contributes to two main strands of literature on the causes and consequences of state effectiveness. We develop a new approach to measuring workers' and organizations' task-specific productivity and use it to demonstrate the extent to which state effectiveness is embodied in the bureaucratic apparatus. Our performance measure is free of the limitations that arise from comparing workers and/or organizations (e.g. firms) (i) engaged in different activities and/or (ii) based on wages and profits.[10] In this sense our paper is most closely related to recent work by Khan *et al.* (2016, 2018) analyzing how the incentives of bureaucrats matter for tax revenue in Pakistan (see also Bertrand *et al.* (2016) on how bureaucrat incentives matter for perceived performance and aggregate outcomes in India), and to the lit-

[7]The WTO's Government Procurement Agreement seeks to restrict countries' ability to use procurement in this way. However, the agreement has only 10 signatories to date, presumably due to the appeal of favoring local bidders in procurement.

[8]The 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. (For example, a five percent bid preference for small businesses in Californian road construction procurement is estimated to increase average costs by 1-4 percent (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.

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

[10]The seminal work of Abowd *et al.* (1999, 2002) spawned a large empirical literature using employer–employee matched datasets to address a range of important questions in labor economics (see, among many others, the papers cited in footnote 4, and also Bertrand & Schoar (2003) and the literature that followed on CEO effects). Wages do not necessarily reflect productivity (Eeckhout & Kircher, 2011; 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) and Cardoso *et al.* (2016) 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).

eratures on management practices in organizations and individual characteristics in the public sector.[11] Since we compare workers and organizations pursuing a single objective, we avoid concerns about multitasking and unobserved dimensions of performance. Moreover, our empirical approach allows us to quantify, to our knowledge for the first time, the importance of the bureaucracy for public sector output *relative to (all) other contributors*.[12]

Second, our paper contributes to an emerging body of evidence on how public policy design should be tailored to context (see e.g. Laffont, 2005; Duflo *et al.* , forthcoming; Best *et al.* , 2015; Hansman *et al.* , 2017). We do so by focusing on the often overlooked fact that policy implementation is delegated to bureaucracies. 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 systematically affects the relative costs and benefits of different policies (see also Vivalt, 2017; Dehejia *et al.* , 2016).[13] 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; Deaton, 2010; Heckman, 2010, 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 Russian 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 "buy local" 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, the Russian government created an extremely decentralized system to perform public procurement, which now makes up about 10 percent of Russia's non-resource GDP.[14] Each government entity

[11] In addition to Bertrand *et al.* (2016); Khan *et al.* (2016, 2018), see also Burgess *et al.* (2012); Duflo *et al.* (2013); Callen *et al.* (2016a,b); Deserranno (forthcoming) for important evidence on how performance incentives affect public sector workers' performance; Rasul & Rogger (2018); Bloom *et al.* (2015); Janke *et al.* (2018) on management practices in public sector organizations; and Dal Bo *et al.* (2013); Duflo *et al.* (2013); Ashraf *et al.* (2014a,b); Hanna & Wang (2017); Callen *et al.* (2016a,b); Deserranno (forthcoming); Bertrand *et al.* (2016) on the characteristics of front-line public sector workers and public goods provision. 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); Decarolis *et al.* (2016); Coviello & Gagliarducci (2017); 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.

[12] In this sense our paper is related to Yao & Zhang (2015). They estimate the share of the variance in cities' economic growth in China attributable to mayors. Their study belongs to an important body of work analyzing how public sector leaders and politicians matter for aggregate economic outcomes (Jones & Olken, 2005; Bertrand *et al.* , 2016; Easterly & Pennings, 2017). Most recently, Xu (forthcoming) convincingly shows that governors of colonies in the British Empire were more likely to be promoted if they were connected to their superior, despite raising less revenue and investing less.

[13] The 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.* , forthcoming; Allcott, 2015).

[14] The Soviet Union, like other socialist states (see e.g. Bai & Jia, 2016), operated a centralized bureaucracy (see e.g. Cheremukhin *et al.* , 2016). Since 1991, the Russian bureaucracy has become very decentralized (Enikolopov & Zhuravskaya, 2007).

has the legal authority to make its own purchases and there are no centralized purchases (such as framework contracts). Conversely, the legal framework for procurement is centralized. Competitive bidding for all purchases above USD 35,000 is mandatory and the rules and regulations governing tender processes at federal, regional, and municipal levels of government were harmonized under Federal Law No. 94-FZ in 2005 (Yakovlev *et al.* , 2010; Krylova & Settles, 2011). Electronic auctions are the most common method of procurement. They were used for 53.5 percent of Russian procurement during our data period (2011–2016, see Online Appendix Table OA.4) and so our analysis will restrict attention to these electronic auctions in order to study bureaucrats and organizations each performing the exact same task.

Electronic 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 1 traces out the steps involved in a procurement process together with the number of purchases in 2011–2016 that followed each path from announcement to contract.

Each purchase starts with an auction announcement. The announcement contains technical details on the item(s) to be purchased, a maximum allowable 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 electronic auction. Purchasing bureaucrats oversee this entire process, collecting information on technical specifications from clients and directing subordinates to draw up procurement plans. In order to participate in an auction, suppliers must first obtain accreditation from the platform. This requires not being in a state of bankruptcy, not being sanctioned under administrative law, not having substantial unpaid taxes, and not being listed in the registry of suppliers who have committed violations of procurement rules during the last two years. Once accredited, suppliers must prepare a formal application, consisting of two parts. The first part describes the good 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, is designated to oversee the auction. The commission receives and evaluates the anonymized first part of each application from the platform before the auction is held. The purchasing bureaucrat takes the lead in directing the review work of the commission. Applications to participate are denied if the supplier is not accredited, cannot pay the security deposit, or if its proposal does not comply with the requested item specifications. In the event that only one supplier is approved to participate, the auction is declared "not held", the commission receives the second part of the approved supplier's application, and a contract is drawn up with that supplier at the maximum allowable price. This is relatively common; in 1.4 million cases, or 22 percent of purchases, there is only one eligible participant. If there are no approved applicants, the purchase is cancelled. This occurs in 13 percent of purchases.

If more than one supplier is approved, the auction is held. Approved suppliers remain anonymous and are each assigned a participant number. All participants log in to the online 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 made available to all auction participants. The auction continues until ten minutes have passed since the most recent qualifying bid.

Following the conclusion of the auction, the commission receives and reviews the second part of the applications. These contain the identifying information for the participants, but do not allow for suppliers to be linked to bids. The commission checks the materials to make sure the suppliers' accreditations, licenses, names, tax ID numbers, registration, and documents confirming participation in the auction are correct. Among the set of bidders deemed to be in accordance with the rules, the contract is signed with the participant who submitted the lowest bid. Purchasing bureaucrats then monitor contract fulfillment.

2.2 The role of bureaucrats and organizations in procurement

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, who we will refer to as the *procurers*, are tasked with acquiring the good 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. As such, under any given set of central rules, the price paid should be considered the main measure of how effective procurers have been at implementing the government's procurement policy.

Each regional authority sets rules dictating that the organizations under its jurisdiction use either a bureaucrat who is an employee of the organization or a bureaucrat who is an employee of an external public agency—agencies whose bureaucrats conduct procurement auctions with and for multiple organizations—for a given type of purchase. The type of purchase is defined by the maximum price of the contract and the nature of the item.[15] Bureaucrats can thus either be "in-house" or "external". This means that instances in which a given bureaucrat is observed working with more than one organization (and vice versa) occur for two different reasons in Russian public procurement. The first is that bureaucrats change employers—from working for one organization (or external procurement agency) to another, as organizations hire and separate from bureaucrats. The other is that external bureaucrats may conduct purchases with multiple organizations, and a given organization may end up using 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.[16]

[15] External procurement agencies can be organized by a given authority (for example an education or health ministry/department), at the federal, regional, or municipal level. Part of the motivation for allowing the creation of such public agencies was to allow different organizations purchasing the same or similar goods to join forces so as to achieve lower per-unit prices. In practice, the decentralized management of procurement in Russia and coordination required to co-purchase means that joint purchases are very rare. Note that we control for the factors that authorities with an external procurement agency use to determine whether an item can be purchased by an in-house bureaucrat or must be purchased by external bureaucrats—the type of good and/or maximum allowable price of the contract—in our empirical analysis below.

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

The education and labor markets for procurement officers are decentralized. Individuals interested in working in public procurement seek out educational and employment opportunities as in the private sector.[17] In particular, the Russian government does not educate bureaucrats, nor does it operate a centralized civil service administration to recruit, train, and assign public servants to postings (Barabashev & Straussman, 2007). In all cases we are aware of, procurement bureaucrats are paid a flat salary.

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 a maximum allowable price. The organization and bureaucrat then together designate the commission to oversee the auction process. The bureaucrat is on the committee, except in special circumstances, and manages its workflow. With the help of the committee, she or he is in charge of first stage review of applications, the auction itself, and second stage review. Procurement officers act as professional managers of a team of specialists who provide consultations to the organization and collect information needed to design the tender. At the end of the process, the organization signs the contract with the winner and signs for delivery. The same or a similar division of labor applies when in-house bureaucrats, i.e. the procurement managers working directly for the organization, are used, and in purchases conducted before 2014.

2.3 Preferences for domestically manufactured goods

As part of the 2005 law, the Russian government established a system to provide for special treatment of—a "preference" for—some types of goods in electronic auctions. Certain categories of goods manufactured in Russia received a 15 percent bid preference for parts of each year from 2011 through 2016.[18] If among the first stage offers there was at least one bidder offering foreign-made goods and at least one bidder offering locally manufactured goods, then if the winning bidder had offered a foreign good, the contract she was offered to sign would be for 85 percent of her final bid.

The application of the preference regime was determined as follows. Each year from 2011 to 2014 a list of good categories for which a preference for domestic goods was to apply was drawn up.[19] The presidential order defining the list was passed in May or June and remained in effect through 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). Preferences were not applied during the first period of each year from 2011–2014. Organizations filing procurement requests for goods on the list were required to publicly inform potential suppliers that the preference would be applied. Preferenced goods spanned many categories, including automobiles, clocks, various food

[17] Examples of private academies offering trainings in the procurement sector include ArtAleks `http://artaleks.ru/` and the Granit Center `http://www.granit.ru/`. Interviews with experts and a review of recent vacancies posted to open online job boards revealed that the primary requirements are a legal education, management experience, and knowledge of the existing laws governing procurement.

[18] The country of origin of a good was defined as that where the good was "completely produced", or where it underwent "significant reprocessing".

[19] Preferences were first given to domestic manufacturers in 2008 to stimulate the economy during the period of 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.

products, medical equipment, pharmaceuticals, textiles and furs (see Online Appendix Table OA.5).

Our analysis of the role of the bureaucratic apparatus in driving variation in procurement performance in sections 4 and 4.5 restricts attention to the standard policy regime without preferences. In section 5.1 we analyze the average impact of the preference regime, while in section 5.2 we show that variation in procurer effectiveness drives heterogeneity in the impact of the preference regime.

2.4 Building a dataset of comparable procurement purchases

A key innovation of the 2005 law was the creation of a centralized official procurement website (`http://zakupki.gov.ru/`), launched on January 1, 2011, which provides information to the public and suppliers about all purchases. This website is our main source of data on procurement. As mentioned above, in order to focus on purchases made under a consistent set of rules, we focus exclusively on electronic auctions. We use data on the universe of electronic auction requests, review protocols, auction protocols, and contracts from January 1, 2011 through December 31, 2016. In all, we have information on 6.5 million auction announcements for the purchase of 21 million items. In order to be able to compare purchases made by different individuals and organizations, we must hold constant the item that is being purchased. Purchases of services and works contacts are highly idiosyncratic, so we remove them from our sample, resulting in a sample of 15 million purchased items.

Table 1 shows summary statistics of our procurement data. Column (1) summarizes all items purchased without the application of the bid preference regime. Our Analysis Sample—which, as explained in Section 4.2, leaves out "isolated" bureaucrats, organizations, and goods—is summarized in Column (2).[20] While the organizations in our Analysis Sample are more likely to be federal or regional, and less likely to fall under Internal Affairs or agriculture, the purchases they conduct are of very similar size and quantity to those in the full sample, reassuring us that the sample we use for analysis is fairly representative. Column (3) summarizes our sample when purchases to which the bid preference regime applied, and which were conducted by procurers in the Analysis Sample, are included. This sample shows similar differences from the full sample in types of organizations represented—for example, fewer federal organizations are included—but purchases in the Analysis Sample incl. Bid Preferences are very similar in size and quantity to those in the Analysis Sample.

In order to make comparisons of procurement performance across buyers, we must hold constant the nature of the good being purchased. A great deal of previous research in economics has faced this challenge, but existing approaches typically achieve within-category homogeneity at the cost of losing generality.[21] To avoid doing so, we use the text of the final contracts, in which the precise nature of the good

[20]Since our empirical approach relies on measuring price changes associated with organizations switching bureaucrats and vice versa, we focus on an analysis sample that removes bureaucrats, organizations, and goods that are not connected to switchers, or which make fewer than five purchases. We also remove sets of organizations connected by bureaucrats switching between them that contain fewer than three bureaucrats or organizations. Section 4.2 and Online Appendix OA.2 discuss the implications of the connectedness of the bureaucrats and the organizations for our empirical approach.

[21]Broadly, three approaches have been taken: using hedonic regressions to estimate consumers' demand for and/or suppliers' costs of producing good attributes when rich attribute data is available(Griliches, 1971; Rosen, 1974; Epple, 1987); using product codes provided by e.g. customs agencies to partition goods (Rauch, 1999; Schott, 2004); or restricting attention to products that are by nature especially homogeneous (Syverson, 2004; Hortacsu & Syverson, 2007; Foster *et al.* , 2009).

being purchased is laid out, to classify purchases into narrow product categories within which quality differences are likely to be negligible. The contracts contain detailed descriptions, but in unstructured text fields, so we apply text analysis machine learning methods to assign purchases to homogeneous categories (see also Gentzkow & Shapiro, 2010; Hansen *et al.* , 2018; Hoberg & Phillips, 2016).

Our method proceeds in three steps. First, we transform the good descriptions in the contracts into vectors of word tokens. Second, we develop a transfer learning procedure. The procedure uses good descriptions and their corresponding 10-digit Harmonized System product codes in data on the universe of Russian customs declarations to train a classification algorithm that assigns good codes to vectors of word tokens. We then apply this algorithm 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 created in the second step. Details are in Online Appendix OA.3.[22]

To complement this approach to achieving comparability between purchases without restricting the sample to very specific goods, we collect additional data on purchases of pharmaceuticals, a type of good that is homogeneous by nature (Bronnenberg *et al.* , 2015). Russia's government regulates the pharmaceutical market to ensure that certain drugs are available and affordable, compelling suppliers to register these drugs in a List of Vital and Essential Medicinal Drugs (LVEMD). This list includes information on each drug's International Nonproprietary Name (INN)[23]; the name and location of the manufacturer; date of registration; and maximum price for sale on the Russian market. We use fuzzy string matching to combine the contract data on procured medicines with corresponding entries in LVEMD using each drug's international brand (trademark) name, active ingredient (INN), dosage (mg, g, mkg), active units (IU), concentration (mg/ml, mkg/ml), volume (ml), and units (tablets, packages). This allows us to construct a barcode-level classification for pharmaceuticals purchases that we can use as an alternative to our text classification categories, and to identify the manufacturer (and thus country of origin) for each pharmaceutical procured.[24] The pharmaceuticals subsamples are summarized in columns (4)–(6) of Table 1, which are defined analogously to columns (1)–(3). The pharmaceuticals subsamples are quite similar to the All Products samples (and to each other) in most dimensions, with a somewhat higher average reservation price and somewhat lower average unit price than in the full sample.

Finally, to study the firms that participate in any stage of the procurement process we use data from the Bureau Van Dijk's *Ruslana* database. Ruslana covers the vast majority of registered firms in Russia that file financial information.[25]

[22]Online Appendix OA.3 also analyzes the sensitivity of our main findings to the choices made when developing our text analysis methodology. As Figure OA.1 and Tables OA.11 and OA.12 show, the findings are remarkably robust.

[23]INN is a globally recognized term to denote the chemical substance of the medicine, see `http://www.who.int/medicines/services/inn/en/`.

[24]We restrict the Pharmaceuticals Subsample to purchases of drugs we can match to LVEMD. Cases in which we are unable to match a drug to the LVEMD can arise both because the medicine is not classified by the Russian government as "essential" (i.e., covered by the LVEMD) and because insufficient information on dosage and quantity is available in the procurement contract.

[25]All firms are by law required to submit accounting data on an active basis. All statistics are standardized by the Russian Ministry of Finance and provided to agencies such as BVD for dissemination to end-clients.

2.5 Corruption

Both public procurement and Russia are associated with widespread corruption (Transparency International, 2016; Szakonyi, 2017). Corruption in procurement can take many forms. Most of these will result in a higher purchase price, and as such be captured in the price-based measure of bureaucratic effectiveness we develop in Section 4. If such forms of corruption—e.g., collusion—are associated with specific procurers, then our empirical procedure will assign a lower effectiveness score to those procurers.[26]

There are two forms of corruption or incompetence that could undermine our estimates of bureaucratic effectiveness. First, if our controls for the nature of the good being purchased are inaccurate, our estimates will conflate true performance differences with differences in the quality of the goods being purchased. We address this through the text-based good classification methodology and barcode-level classification of pharmaceuticals described in the previous sub-section, and a battery of tests that relax the within-category homogeneity assumption in Sub-section 4.4.

Second, procurers who achieve low prices may in theory purchase items that are simply not delivered. Our contract execution dataset is unusual, however, in that it includes information on whether the organization paying for the items signed for delivery. Non-delivery is very rare.[27] Russian procurement laws do not allow for any form of renegotiation of cost of delivery for the off-the-shelf goods we focus on, as commonly occurs e.g. in works contracts (see e.g. Bajari *et al.* , 2014; Decarolis, 2014; Decarolis & Palumbo, 2015).

Given this, we believe that our price-based procurer effectiveness estimates capture what governments mandate procurers to target: the price paid for goods of specified quality that are satisfactorily delivered. In the theoretical framework and empirical analysis below, we remain entirely agnostic about the extent to which some procurers pay higher prices than others because they are prone to forms of corruption that manifest themselves in the prices paid and the extent to which they do so because they are of lower intrinsic ability.[28]

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.

[26] It is also possible that some procurers systematically see their auctions won by suppliers that subsequently do not sign the contract, either because the winners choose not to do so or because they are deemed by the procuring commission to offer sub-standard goods. In such cases the contract goes to the second-lowest bidder. Since we observe both the bids and the contract signed, we also observe the instances in which the contract is not signed with the lowest bidder. Such instances are rare, accounting for under one percent of purchases (see Figure 1). More importantly, the outcome we focus on when estimating procurer effectiveness is the price ultimately *paid* for the item. As such, the consequences of auction winners not signing the contract will be captured by our effectiveness measures.

[27] Less than one percent of the auctions in our sample suffered from "bad execution". The data also include information on early and late delivery. These are correlated with estimated procurer effectiveness, but can only explain a tiny fraction of the dispersion in price effectiveness we estimate. Results available from the authors upon request.

[28] In Italy, Bandiera *et al.* (2009) find that 83 percent of waste in Italian public procurement purchases is due to low bureaucratic ability rather than corruption.

We also show how the introduction of bid preferences is expected to differentially affect procurement by bureaucracies with different levels of state effectiveness. The model makes predictions for the patterns of heterogeneity in the impact of the introduction of bid preferences that we test for in Section 5.2.

3.1 Procurement with heterogeneous state effectiveness 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 foreign or a local supplier through a second-price descending auction. Bureaucratic effectiveness affects the prices the government is able to achieve in two ways: directly, by increasing the cost to suppliers of fulfilling a government contract; and indirectly, by imposing participation costs on sellers wishing to bid for contracts.

The production cost of the item to be procured for a firm with the lowest possible markdown is $e^{\mathbf{X}'\boldsymbol{\beta}}$, where **X** are observable attributes of the good being purchased. When writing tender documents, bureaucrats add requirements α_M and organizations add requirements ψ_M. These requirements change the cost of complying with the contract from the production cost to $M = e^{\mathbf{X}'\boldsymbol{\beta}+\alpha_M+\psi_M}$. These requirements may include the precise date and place of delivery, the size of the order, and any other requirements that directly affect the cost to firms of fulfilling the contract. Similarly, bureaucrats and organizations can add 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 costs to suppliers of bidding for the contract, though not directly of fulfilling the contract.

In the first stage of the procurement process, two firms—one local and one foreign—observe the specifications $\{\mathbf{X}, \alpha_M, \alpha_c, \psi_M, \psi_c\}$ and then decide whether to pay their participation cost c_i to learn the markdown they are able to offer m_i and enter the auction.[29] The foreign firm $i = F$ and the local firm $i = L$ differ in both their expected markdowns and their participation costs. Markdowns $m_i \geq 1$ are independent and Pareto distributed with Pareto parameters δ_F and δ_L, making the firms' costs of fulfiling the contract M/m_i. Foreign firms have higher expected markdowns ($\delta_F < \delta_L$) but face higher participation costs: $c_i = \frac{M}{1+\delta_i} - \frac{M}{1+\delta_L}\sqrt{1-\alpha_c-\psi_c}$.[30]

In the second stage of the procurement process, if only one supplier chose to enter the auction, she is awarded the contract at a price M. If neither supplier chose to enter, the bureaucracy finds an outside supplier and awards her the contract at a price of M.[31] 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; Krishna, 2010).

[29] Note we assume that firms do not know their markdown 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 markdown 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.

[30] This functional form makes the expressions for profits and prices tractable. In particular, the first term ensures that there is no pure strategy equilibrium in which both suppliers enter with certainty, and the second term ensures that expected log prices are linear in α_c and ψ_j but the qualitative results only require the participation costs to be increasing in α_c and ψ_c.

[31] A 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.

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.1. 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. Denoting the bidding strategy of supplier i with markdown m by $b_i(m)$, we have $b_F(m) = b_L(m) = M/m$ (see e.g. Milgrom, 2004; Krishna, 2010). 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 bidder i has fulfillment cost 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 M yielding expected profits of $M - \mathbb{E}[M/m_i]$. 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

$$\begin{aligned} q_j\mathbb{E}[\pi_i] + (1-q_j)(M - \mathbb{E}[M/m_i]) - c_i = 0 \\ \iff q_j = \frac{M - \mathbb{E}[M/m_i] - c_i}{M - \mathbb{E}[M/m_i] - \mathbb{E}[\pi_i]}, \end{aligned} \tag{1}$$

where $i, j \in \{F, L\}$, $i \neq j$. We can summarize the equilibrium of the auction game in the following proposition:

Proposition 1. *In the Nash equilibrium of the auction, the bidders, $i \in \{F, L\}$ enter with probabilities*

$$q_i = \sqrt{\kappa(1 - \alpha_c - \psi_c)}, \tag{2}$$

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

$$\begin{aligned} \mathbb{E}[\log(p)] &= \log(M) - \tfrac{q_F q_L}{\delta_F+\delta_L} = \mathbf{X}'\boldsymbol{\beta} + \alpha_M + \psi_M + \tfrac{\kappa}{\delta_F+\delta_L}(\alpha_c + \psi_c - 1) \\ &= \mathbf{X}'\boldsymbol{\beta} - \tfrac{\kappa}{\delta_F+\delta_L} + \tilde{\alpha} + \tilde{\psi}, \end{aligned} \tag{3}$$

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

1. *Bureaucracies that impose higher contract fulfillment costs α_M, ψ_M 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.*

Proof. See Online Appendix OA.1.1. □

Equation (3) shows how prices vary with with the fulfillment and participation 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: introducing bid preferences

We now study the impact of introducing bid preferences favoring the local bidder L in this simple model.[32] The winner of the auction is still the bidder with the lowest bid, but if that bidder is the foreign bidder, the contract price will only be $p = \gamma b_L$, where $\gamma < 1$, while if the local firm submits the lowest bid the contract price will be the undiscounted $p = b_F$. Otherwise the auction protocol is unchanged.

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 M/m_F: $b_F = M/\gamma m_F$. However, when her shaded bid would have no chance of winning ($m_F < 1/\gamma$), 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

$$\mathbb{E}\left[\pi_F|b_F,\gamma\right] = \mathbb{E}\left[\gamma\left(b_L - b_F\right)|b_L > b_F\right]\mathbb{P}\left(b_L > b_F\right) \tag{4}$$

$$\mathbb{E}\left[\pi_L|b_L,\gamma\right] = \mathbb{E}\left[b_F - b_L|M \geq b_F > b_L\right]\mathbb{P}\left(M \geq b_F > b_L\right) + \mathbb{P}\left(m_F < 1/\gamma\right)\left(M - b_L\right). \tag{5}$$

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.

The effect on the expected price in an auction depends on the strength of two offsetting effects. First, the direct effect of the penalty lowers prices: when the foreign bidder wins, she is paid less. Second, the preference regime affects prices indirectly through the fact that the foreign bidder is less likely to win, in which case the government pays more to the local bidder. We assume that the latter effect dominates.[33] That is, we make

Assumption 1. $\gamma^{\delta_F}\left[1 - \log\left(\gamma^{\delta_L}\right)\right] > 1$.

Effectively, Assumption 1 requires the bid penalty to be sufficiently small (γ sufficiently large). The preference regime also changes the bidders' participation decisions. Participation decisions depend on the expected profitability of entering the auction. Bidder F has to bid more aggressively to have a chance of winning, discouraging her from entering, while bidder L has a better chance of winning, encouraging her to enter more often.The net effect on participation depends on which effect is stronger. A sufficient condition for participation to increase in the mixed strategy equilibrium is the following:

[32]In 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.

[33]Appendix OA.1.2 derives the impact of introducing preferences when Assumption 1 fails. The difference is only that in the first case in Proposition 2 prices fall instead of rising.

Assumption 2. $\gamma^{\delta_F}\left[1+\frac{\delta_L}{\delta_F+\delta_L}\left(1-\gamma^{1+\delta_F}\right)\right]<1$

Assumption 2 effectively requires the bid penalty to be sufficiently strong (γ sufficiently small) to increase overall participation. Assumptions 1 and 2 together require γ to be in a certain range. The overall effect of the preference regime depends on the combination of participation and bidding effects. We can summarize the overall effects in the following proposition:

Proposition 2. *When assumptions 1 and 2 hold, the introduction of bid preferences γ has different effects on three groups of bureaucracies differing in their effectiveness.*

1. *For bureaucracies with $\alpha_c+\psi_c\leq\underline{c}$, prices rise, the expected number of bidders is unchanged, and the probability that bidder L wins the contract at auction increases;*

2. *For bureaucracies with $\underline{c}<\alpha_c+\psi_c\leq\bar{c}$, prices rise, the expected number of bidders falls, and the probability that bidder L wins the contract at auction decreases;*

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

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

$$\underline{c}=1-\left(\frac{1+\delta_L}{1+\delta_F}\left(1-\gamma^{1+\delta_F}\right)+\frac{1+\delta_L}{1+\delta_F+\delta_L}\gamma^{1+\delta_F}\right)^2 \tag{6}$$

$$\bar{c}=1-\left(\frac{1+\delta_L}{1+\delta_F+\delta_L}\gamma^{\delta_F}\right)^2 \tag{7}$$

Proof. See Online Appendix OA.1.2. □

When bureaucracies impose very low participation costs on potential bidders ($\alpha_c+\psi_c\leq\underline{c}$) the preference policy does not deter the firms from entering the auction, but it changes bidding behavior in the auction. The local bidder is more likely to win, and because of Assumption 1, the direct effect of the preferences lowering prices is smaller than the indirect effect through the higher probability that the local bidder wins raising prices.

The preference regime increases expected profits from the auction for the local bidder and lowers them for the foreign bidder. For bureaucracies that impose intermediate participation costs on potential bidders ($\underline{c}<\alpha_c+\psi_c\leq\bar{c}$), expected profits from the auction net of participation costs are positive for local bidders and negative for foreign bidders. As a result, the mixed strategy equilibrium of Proposition 1 breaks down; the local bidder always enters; and the foreign bidder does not find it profitable to also enter. Since only the local bidder enters, the auction does not take place and the local firm gets the contract at the maximum price M.

Finally, when bureaucracies impose high participation costs ($\bar{c}<\alpha_c+\psi_c$) neither bidder is willing to enter with certainty and so the equilibrium still involves mixed strategies, as in Proposition 1. Assumption 2 implies that the increase in bidder L's willingness to enter is larger than the decrease in bidder F's willingness to enter. This increase in the probability of both bidders entering and the auction taking place lowers 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. The entry effect is also strong enough that the increase in the probability that the local bidder wins the contract at auction is larger in this case than in the case of buyers who impose very low participation costs on potential bidders ($\alpha_c + \psi_c \leq \underline{c}$).

Proposition 2 makes three predictions about heterogeneity in the impact of bid preferences we should expect to see in our empirical analysis. First, we should see that bureaucracies that pay higher prices when there are no bid preferences—a pattern that Proposition 1 shows is a proxy for higher participation costs—experience price *decreases*, while bureaucracies that pay lower prices absent the bid preferences experience price *increases*. Second, we should see that the average number of participants in procurement processes increases 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 Section 5.2

4 How Important is a Good Bureaucracy?

In this section we estimate the share of the variation in procurement effectiveness that can be attributed to the individuals and organizations in the bureaucracy. To do so, we extend the method pioneered by Abowd *et al.* (1999). Our method exploits *switchers*—bureaucrats who make purchases with multiple organizations, and organizations who make purchases with multiple bureaucrats—to identify the shares of the sample-wide variation in performance that can be attributed to the specific agents who implement procurement policy.

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 procurement auction. To do so, we use an event study analysis tracking organizations that switch bureaucrats. This happens frequently in Russia. Column (1) of Table 2 displays the number of events we observe where an organization switches bureaucrat by quartile of bureaucrat effectiveness. We observe, for example, 1,736 events in which an organization switches from a bureaucrat of highest quartile effectiveness—as defined below—to a bureaucrat of lowest quartile effectiveness, and 1,356 events in which an organization switches from a lowest to a highest quartile bureaucrat. In total, we observe around 65,000 events in which organizations switch bureaucrats. Column (2) of Table 2 shows that the average number of purchases observed for each bureaucrat-organization involved in a switch is also large.

To construct the event study, 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, each employment spell (event time ≤ 0) is matched with the earliest future spell (event time > 0) involving the same organization but a different bureaucrat. This change of bureaucrats then constitutes an event. 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).[34] In Appendix A.1 we show that the patterns we see in the resulting event studies are very similar when we change a series of choices made in constructing the event studies.

Figure 2 shows the time profile of prices achieved around the time that organizations switch between bureaucrats of different effectiveness. The horizontal axis displays event time, i.e. purchase weeks. The vertical axis displays the average quality-adjusted prices paid by the relevant bureaucrat-organization pair in a given week. Table 2 highlights that the number of switches used to construct each quartile-to-quartile plot in Figure 2, 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.[35]

Four key findings emerge from Figure 2. First, the figure shows that quality-adjusted prices paid change sharply, and in the expected direction, at the point in time 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 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.

We also construct analogous event study figures for organizations and bureaucrats switching from purchasing one type of *good* to another. The results are presented in Online Appendix Figure OA.2. Each event study shows the same general patterns as in Figure 2.

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.

4.2 Variance decomposition method

In this sub-section we first show how to aggregate causal effects of specific bureaucrats and organizations estimated through switches into estimates of the share of sample-wide variation in procurement performance bureaucrats as a whole and organizations as a whole explain. These estimates are the starting point for our *sampling error-corrected* predictions of the impact of specific bureaucrats or organizations on

[34] We quality-adjust prices by regressing them on log quantity, good fixed effects, month fixed effects, interactions between 2-digit HS product categories, years, regions, and lot size (as explained in more detail in the next sub-section).

[35] The table also displays the average number of calendar weeks between each purchase week on the x-axis of Figure 2

prices paid. We use these predictions to examine the mechanisms through which procurers affect prices in Section 4.5 and how bureaucratic effectiveness impacts the way policy rules map into public sector output in Section 5.

We adapt and extend the method to study wage dispersion in the private sector pioneered by Abowd *et al.* (1999). We specify 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 that is due to the organization $\tilde{\psi}_j$. As the model 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

$$p_i = \mathbf{X}_i\boldsymbol{\beta} + \tilde{\alpha}_{b(i,j)} + \tilde{\psi}_j + \varepsilon_i \tag{8}$$

To control flexibly for the item being purchased, we include in $\mathbf{X}_i$ log quantity, good fixed effects, month fixed effects, and interactions between 2-digit HS product categories, years, regions, and lot size.[36]

As we showed in Sub-section 4.1, switches—bureaucrats making purchases on behalf of multiple organizations and organizations using multiple bureaucrats to make purchases—provide a compelling means of identifying the impact of procurers on prices. However, such switches do not connect all bureaucrats and organizations that conduct procurement in Russia. As Abowd *et al.* (2002) show, individual and organization effects are only identified *within* sets of organizations connected by individuals moving between them.[37]

To proceed, we normalize the bureaucrat and organization effects to have mean zero in each connected set and augment (8) to include an intercept, $\gamma_{s(b,j)}$, specific to each connected set:

$$p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i \tag{9}$$

In Online Appendix OA.2, we show that while the $\tilde{\alpha}$s and $\tilde{\psi}$s in equation (8) are not identified, the αs, ψs and γs in equation (9) are. These are related to the underlying bureaucrat and organization effects as follows: $\alpha_b = \tilde{\alpha}_b - \overline{\alpha}_{s(b)}$, $\psi_j = \tilde{\psi}_j - \overline{\psi}_{s(j)}$, and $\gamma_{s(b,j)} = \overline{\alpha}_{s(b,j)} + \overline{\psi}_{s(b,j)}$, where $\overline{\alpha}_{s(b)}$ is the mean bureaucrat effect in the connected set containing bureaucrat b, and similarly $\overline{\psi}_{s(j)}$ is the mean organization effect in organization j's connected set.[38]

[36] By 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 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".

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

[38] Faced 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.*

We can use equation (9) to decompose the variance of prices into its constituent parts using

$$\begin{aligned}\mathsf{Var}(p_i) = \mathsf{Var}\left(\alpha_{b(i,j)}\right) + \mathsf{Var}\left(\psi_j\right) + \mathsf{Var}\left(\gamma_{s(b,j)}\right) + 2\mathsf{Cov}\left(\alpha_{b(i,j)}, \psi_j\right) + \mathsf{Var}\left(\mathbf{X}_i\boldsymbol{\beta}\right) \\ + 2\mathsf{Cov}\left(\alpha_{b(i,j)} + \psi_j, \gamma_{s(b,j)} + \mathbf{X}_i\boldsymbol{\beta}\right) + 2\mathsf{Cov}\left(\gamma_{s(b,j)}, \mathbf{X}_i\boldsymbol{\beta}\right) + \mathsf{Var}\left(\varepsilon_i\right)\end{aligned} \tag{10}$$

all of which can be identified. Since $\mathsf{Var}\left(\alpha_{b(i,j)}\right)$ and $\mathsf{Var}\left(\psi_j\right)$ are variances within connected sets, they are lower bounds on the true variances of bureaucrat and organization effects.[39] However, we can combine our estimates to capture the variance in prices that is attributable to the bureaucrats and the organizations *jointly* as follows:

$$\mathsf{Var}\left(\tilde{\alpha}_b + \tilde{\psi}_j\right) \equiv \mathbb{E}\left[\mathsf{Var}\left(\tilde{\alpha}_b + \tilde{\psi}_j | s(b,j)\right)\right] + \mathsf{Var}\left(\mathbb{E}\left[\tilde{\alpha}_b + \tilde{\psi}_j | s(b,j)\right]\right) = \mathsf{Var}\left(\alpha_b + \psi_j\right) + \mathsf{Var}\left(\gamma_{s(b,j)}\right) \tag{11}$$

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 (Klepikova, 2017). 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. In Appendix Table OA.6 we show that our results are robust to using only the largest set of connected organizations. Table 1 compares the full sample and the Analysis Sample, and Online Appendix Table OA.3 compares the Analysis Sample to its largest connected set. The three samples are broadly similar on a host of observable characteristics, as discussed in Section 2.

We can obtain unbiased estimates of bureaucrat and organization effects using ordinary least squares under the assumption that the residuals ε_i in equation (9) 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 (9) 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 equation (9). Suppose that, in addition, the bureaucrat-organization matches that form

(2016) who study the largest male and female connected sets in Portuguese data.

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

are correlated with the omitted complementarities, for example because organizations seek to recruit bureaucrats who specialize in particular goods they require. Under such sorting-on-match-effects, an estimate of a bureaucrat effect from equation (9) will recover a mixture of the true effect and the average complementarity of the bureaucrat with the organizations he/she is matched to.

However, if sorting on match effects occurs, we would expect the event study in Figure 2 to show that organizations switching from bureaucrats who on average pay high prices to bureaucrats with low average prices enjoy larger price decreases than the price increases suffered by organizations switching in the opposite direction. Under such a scenario we should see 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 2.[40] The striking symmetry of the event study evidence indicates that omitted complementarities between bureaucrats and organizations are unlikely to bias our estimates. In Section 4.4 we explore the possibility that the log-linear model is misspecified further.

We use a large sample of public procurers, but nevertheless, our estimates need not be consistently estimated, even if they are unbiased. This is because we seek to estimate a large number of bureaucrat and organization effects. Consistency of the estimated fixed effects requires that the number of observations *on each group* tends to infinity (Neyman & Scott, 1948; 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 may be compounded if the network of bureaucrats and organizations features too few switches. In finite samples, such *limited mobility bias* often results in a spurious negative correlation between the two dimensions of estimated fixed effects (Andrews *et al.* , 2008). In our data, each connected set is unusually densely connected—we observe bureaucrats working with 5.2 organizations on average, and organizations with 4.8 bureaucrats.[41] However, limited mobility bias cannot be ruled out a priori.

We address these sampling error issues in three ways. First, to estimate standard errors for our variance decomposition, we bootstrap so that we can take into account the patterns of correlation in the residuals.[42]

[40] If anything, Figure 2 showed slightly *smaller* price decreases when organizations switch to lower average-price bureaucrats than when organizations switch to higher average-price bureaucrats.

[41] This is in part because switches in Russian public procurement arise not only through traditional labor market churn, but also because some bureaucrats conduct purchases for multiple organizations at the same point in time, and vice versa, as discussed in Sub-section 2.2.

[42] Each bootstrap randomly samples from the residuals of equation (9) and reestimates the equation. The standard errors are then computed as the standard deviation of the bootstrap estimates. Specifically, we construct partial residuals $\epsilon_i = p_i - \mathbf{X}_i\hat{\boldsymbol{\beta}}$ and randomly reassign bureaucrats and organizations to each observation, preserving 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 approach drastically speeds up computation and makes bootstrapping feasible with our large dataset, but has limitations. (In particular, the procedure imposes clustering at the bureaucrat-organization level in the standard errors. Moreover, since we use the partial residuals ϵ_i rather than reestimating the full model on each iteration, we do not account for correlation between bureaucrat and organization assignment and $\mathbf{X}$).

Second, we take a non-parametric, split-sample approach to estimating the variance components in (10), 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 (9) separately on each sample, yielding two estimates ($k = 1, 2$) for each bureaucrat ($\hat{\alpha}_b^k$), organization ($\hat{\psi}_j^k$), and connected set ($\hat{\gamma}_s^k$) effect. Both estimates are estimated with error, but the errors in the two estimates should be uncorrelated, so we can create split-sample estimates of the variance decomposition terms as follows: $\widehat{\mathsf{Var}}^{SS}(\alpha_b) = \mathsf{Cov}\left(\hat{\alpha}_b^1, \hat{\alpha}_b^2\right)$, $\widehat{\mathsf{Var}}^{SS}(\psi_j) = \mathsf{Cov}\left(\hat{\psi}_j^1, \hat{\psi}_j^2\right)$, $\widehat{\mathsf{Var}}^{SS}(\gamma_s) = \mathsf{Cov}\left(\hat{\gamma}_s^1, \hat{\gamma}_s^2\right)$, and $\widehat{\mathsf{Var}}^{SS}(\alpha_b + \psi_j) = \mathsf{Cov}\left(\hat{\alpha}_b^1 + \hat{\psi}_j^1, \hat{\alpha}_b^2 + \hat{\psi}_j^2\right)$.

Third, we adopt two shrinkage approaches to create predictions of each bureaucrat's and each organization's effect that minimize the mean-squared prediction error. 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 (9) yields estimates of the variance of the sampling error for each bureaucrat and organization effect, s_b^2 and s_j^2. With these we estimate the signal variance of the bureaucrat effects as $\hat{\sigma}_\alpha^2 = \mathsf{Var}(\hat{\alpha}) - \mathbb{E}_b\left[s_b^2\right]$ and analogously for the organizations.

Shrinking pulls each estimate toward zero (their mean), and more noisily estimated effects are shrunk more. Our first approach shrinks the bureaucrat and organization estimates separately, as is common in studies of teacher value-added (see e.g. Kane & Staiger, 2008; Chetty *et al.* , 2014a).[43] To address the limited mobility bias potentially arising from the finite number of switches available for identification, we extend the shrinkage approach used in existing work to explicitly account for the correlation between the estimation errors induced by limited mobility bias. We exploit the fact that our bootstrap also provides estimates of the covariance of all the estimation errors. Formally, we seek the linear combination of the full vector of fixed effects that minimizes the expected mean-squared error of the predictions. Denoting the vector of estimated bureaucrat and organization fixed effects by $\hat{\boldsymbol{\theta}}$ and the matrix of weights by $\boldsymbol{\Lambda}$, the objective is $\min_{\boldsymbol{\Lambda}} \mathbb{E}\left[\left(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)'\left(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)\right]$, which has solution $\boldsymbol{\Lambda}^* = \mathbb{E}\left[\boldsymbol{\theta}\hat{\boldsymbol{\theta}}'\right]\left(\mathbb{E}\left[\hat{\boldsymbol{\theta}}\hat{\boldsymbol{\theta}}'\right]\right)^{-1}$. Replacing the expectations with their sample analogues yields the shrinkage matrix $\hat{\boldsymbol{\Lambda}}^* = diag\left(\hat{\sigma}_\alpha^2, \hat{\sigma}_\psi^2\right)\left(diag\left(\hat{\sigma}_\alpha^2, \hat{\sigma}_\psi^2\right) + \boldsymbol{\Sigma}\right)^{-1}$, where $\boldsymbol{\Sigma}$ is the covariance matrix of the bootstrap estimates and $diag\left(\hat{\sigma}_\alpha^2, \hat{\sigma}_\psi^2\right)$ is the diagonal matrix with $\hat{\sigma}_\alpha^2$ in entries corresponding to entries for bureaucrats in $\boldsymbol{\theta}$ and $\hat{\sigma}_\psi^2$ in entries corresponding to organizations. We label this method "covariance shrinkage". It yields our preferred estimates of the price variance decomposition in equation (10).[44]

4.3 Results

Table 3 shows the results of implementing our variance decomposition method to estimate the share of the variation in procurement performance attributable to the bureaucrats and organizations responsible

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

[44]We 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.

for administering purchases. The first column shows estimates of the standard deviations using the raw fixed effects estimates from equation (9), while estimates from 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.

Three key findings emerge. First, bureaucrats and organizations are each important determinants of variation in 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. The same is found in many studies applying the AKM method to private sector wages.[45] However, our covariance shrinkage approach yields an estimate of the correlation of 0.33, implying that more effective bureaucrats tend to work with more effective organizations and vice versa.[46] This more plausible estimate suggests that covariance shrinking successfully addresses finite sample biases. As a result, the covariance shrunk estimates of share of the variation in performance explained by bureaucrats and organizations—24 and 26 percent 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 Section 4.2, the estimates of the variation in bureaucrat and organization effects are to be interpreted as estimates *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 policy performance Russia-wide. These estimates are shown in row 8 of Table 3. 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 3 have correspondingly large implications for the scope of potential savings from improving the effectiveness of the bureaucracy. To illustrate the magnitude, we can con-

[45]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)).

[46]Of course, such assortative matching does not violate the no-sorting-on-match-effects assumption discussed in sub-sections 4.1 and 4.4.

sider simple counterfactual bureaucracies in which bureaucrats and/or organizations are moved from one percentile of the effectiveness distribution to another, for example because of changes in recruiting, training of existing bureaucrats, or improved organizational management. As seen in Figure 3, our estimates imply that moving all bureaucrats in the lowest quartile of effectiveness to 75th percentile-effectiveness would save the Russian government 3.3 percent of its total, annual procurement expenses. Figure 3 also shows that moving all bureaucrats *and* organizations below 25th percentile-effectiveness to 75th percentile-effectiveness would save the government 10.7 percent of procurement expenditures. Since annual procurement expenses average USD 86 billion, this implies savings of USD 13 billion each year, or 0.9 percent of non-resource GDP (see Online Appendix Table OA.4)—roughly one fifth, for example, of the total amount spent on health care by the Russian government at federal, regional, and municipal level combined in both 2013 and 2014.

4.4 Robustness

We interpret the results in the previous sub-section as capturing the total, causal contribution of bureaucrats and organizations to the Russian state's effectiveness at minimizing the price paid for each specific good it procures. There are two potential challenges to this interpretation. First, 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. Second, the estimated procurer effects can only be interpreted causally if bureaucrats and organizations do not sort based on unmodeled match effects.

Like-for-like comparison To probe these assumptions, 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; Hortacsu & Syverson, 2007; Foster *et al.* , 2009; Bronnenberg *et al.* , 2015). We extract the active ingredient, dosage, and packaging of each drug purchased, as described in Sub-section 2.4, and use these characteristics to assign purchases to barcode-level bins. These bins are used to define good categories and thus to determine which purchases to compare in the pharmaceuticals sample. We then make the same connectivity restrictions as in the full sample to create an analysis sample of pharmaceuticals purchases. As seen in columns (4) and (5) of Table 1, the pharmaceuticals analysis sample is in many ways similar to the full pharmaceuticals sample.

Table 4 presents the results of re-estimating (9) on the pharmaceuticals sample. Naturally, since the sample is more homogenous by construction 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 should happen, since we model the fulfillment costs imposed by bureaucrats and organizations on suppliers as proportional costs.

Second, 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 code to.[47] This result, seen in Column (6) of Table 5, is further evidence that procurers our estimation procedure labels *effective* do not systematically purchase lower quality goods than procurers the procedure labels *ineffective*.

Third, we show that our results are robust to (i) gradually restricting the analysis sample to include more and more homogeneous categories of goods, and—related—to (ii) restricting attention to the most homogeneous types of goods in the 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 (9) on successive subsamples. The first five columns of Table 5 show the results. Column (5) includes all observations for which the Sutton (1998) measure is available.[48] 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 largely unchanged across the columns. In Online Appendix Table OA.7 we repeat this exercise using an alternative measure of scope for quality differentiation developed by Khandelwal (2010) and find the same result.[49]

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.

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 2.

Second, in Appendix A.2, we examine patterns in the size of residuals across the bivariate distribution of the estimated bureaucrat and organization effects. If match effects omitted from (9) 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 A.2. In the bottom panel we repeat this exercise, but now we plot the residuals from a model that is analogous to (9) 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 (9) with fixed effects for each bureaucrat-organization *pair* added in Appendix A.2. 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

[47] The algorithm developed in Step 2 of the procedure outlined in Sub-section 2.4 and Online Appendix OA.3 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.

[48] We are able to match 70 percent of the items assigned an 10-digit HS code in Step 2 of the text analysis method with the Sutton (1998) measure. We thank Eric Verhoogen for sharing the Sutton measure data with us.

[49] Another 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.

et al. , 2013).

In sum, the six tests of the good homogeneity and no-sorting-on-match-effects assumptions discussed in this sub-section provide strong evidence that our empirical model captures the *causal* impact of bureaucrats and organizations on prices paid for each type of good.

4.5 What do Effective Bureaucrats do Differently?

Having established that the state's effectiveness at procuring inputs is partly due to individual bureaucrats and organizations, an important question naturally arises: what distinguishes effective from ineffective procurers? That is, how does bureaucratic effectiveness manifest itself in the process through which inputs are purchased, or attributes of the buyers? To address this question, we leverage the richness of Russian procurement data. As discussed in Section 2, the data we use 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, which covers firms filing financial information with the Ministry of Finance. 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 bureaucracy.

The dataset we start from contains 113 potential explanatory variables, listed in Online Appendix Table OA.8. To avoid overfitting and for parsimony, we use standard regularization techniques to select the variables to focus on. The regularization techniques we choose account for small firms not being covered by the *Ruslana* data and the strong correlation between some of our variables.[50]

Having selected which process indicators to focus on, we relate our covariance shrunk estimates of the ultimate effectiveness of the procurers involved from Sub-section 4.2 to these indicators. Figures 4 and 5 show the results. The left panel of each figure shows regression coefficients from a series of bivariate regressions of the bureaucrat effect (in Figure 4) and the organization effect (in Figure 5) on each of the selected observables alone. The right panel shows 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 4 and 5 need not be causal, in

[50]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. The correlation between our variables can be seen in Online Appendix Figure OA.3. 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. To account for correlation between the explanatory variables without making arbitrary choices of which to include, we use an elastic net procedure for regularization. The elastic net penalty function is a weighted average between a linear penalty (LASSO) and a quadratic penalty (Ridge regression). Intuitively, multicollinearity causes the coefficients of correlated variables to have highly variable and negatively correlated estimates. Ridge regression deals with this by shrinking large coefficients towards zero. Having stabilized the coefficients with the ridge penalty, the elastic net mixing penalty can determine which variables to select (see, for example Friedman *et al.* (2013) for discussion). Our regularization procedure puts a weight of 0.3 on the ridge penalty, but our results are not particularly sensitive to using other weights. See Online Appendix Figures OA.4-OA.9.

part because we do not observe everything different procurers do differently.

Four key findings emerge. First, Figure 4 shows that effective bureaucrats are more experienced than their ineffective peers, making more purchases and buying a greater range of goods. Second, while experience is less important for organizations' performance, Figure 5 reveals significant heterogeneity in performance across organizations by the function they perform, possibly related to variation in management. Third, more effective bureaucrats attract more applicants, and then go on to admit more, and more diverse, bidders into their auctions. Fourth, organizations that set higher reservation prices and revise contracts more often are less effective. These last two findings 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 on sellers bidding for such contracts and consequently attract fewer participants. We conclude from these findings that a key part of what makes procurers effective is their ability to reduce the barriers of entry to participate in procurement auctions.

5 Policy Design with a Heterogeneous Bureaucracy

In our empirical analysis so far, we have held the policy environment constant. We showed that fully 40 percent of the variation in procurement performance in the baseline policy regime governing procurement in Russia is attributable to differences in the effectiveness of the agents—individuals and organizations—charged with implementing policy, and that the returns to increasing the capacity of poorly performing policy implementers are correspondingly large. This raises two subsequent questions. First, how specific are these findings to this particular policy regime? Would these same policy implementers affect procurement performance differently if they were implementing a different set of policy rules? Second, if so, might policy outcomes be improved if different policies were used for bureaucracies of different levels of effectiveness? These questions are important because states have two (potential) ways to increase bureaucratic effectiveness—directly improving the productivity of the bureaucracy, and optimizing the policies the bureaucracy is directed to implement—and in many settings the former may be infeasible or costly.

To investigate, we study the impact of a particular policy change that took place in Russia during our data period—the introduction of bid preferences favoring locally manufactured goods. We estimate both the average impact of the policy and, crucially, how variation in the impact of the policy depends on the effectiveness of the implementing bureaucracy. By doing so we provide a cleanly identified example of the policy design implications of micro level sources of state effectiveness, and evidence on the extent to which context-specific policy design can potentially be used to offset low bureaucratic effectiveness.

5.1 Average impact of bid preferences for locally manufactured goods

We first estimate the overall impact of Russia's bid preferences for locally manufactured goods. Many governments use such "buy local" procurement policies—Russia's version is described in Sub-section 2.3—to attempt to steer demand towards local firms. Evidence on the overall impact of bid preferences is thus important in its own right. Moreover, how prices respond to preference programs of the form we study is theoretically ambiguous.[51]

Russia's preferences policy came into effect only in May or June each year during 2011-2014, and during our whole study period of 2011-2016, the policy covered only a subset of goods—a subset that varied from year to year.[52] We exploit this variation in a generalized difference-in-differences strategy. We estimate regressions of the form:

$$p_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \delta \text{Preferenced}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt} \tag{12}$$

where p_{igt} is the price paid for purchase i of good g in month t. Preferenced_{gt} is a dummy indicating that g is on the preferences list in the year month t falls within, and PolicyActive_t is a dummy indicating that the year's list of preferenced goods has been published and the policy activated. $\mathbf{X}_{igt}$ 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 are Intent to Treat (ITT) effects, and hence directly relevant for policy design.

To estimate (12), we expand the Analysis Sample to also include purchases where bid preferences apply, and which were managed by bureaucrats and organizations in the Analysis Sample. We run the regression separately for the full sample of products and the pharmaceuticals subsample, where we observe the country of origin of the goods being supplied. The samples are summarized in columns (3) and (6) of Table 1. Since all pharmaceuticals are on the preference list, we cannot use variation across products in whether the policy applies for identification in that sample. Instead, we exploit the facts that the preference policy only applies when an auction features both foreign and domestically manufactured goods and that not all pharmaceuticals are manufactured both domestically and abroad. In the pharmaceuticals subsample, we thus redefine Preferenced_{gt} as equal to one if the drug purchased is made—by at least one supplier—both in Russia and abroad.

Table 6 shows the results of estimating (12). In columns (1)–(4) we see that, in both samples, the preferences policy on average 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. This suggest that there is no adverse overall impact on prices or participation of distorting competition in the auctions. 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 and is paid less than its bid.

Our estimates capture the causal effect of the policy under the assumption that the trend in prices

[51] See e.g. McAfee & McMillan (1989); Marion (2007); Krasnokutskaya & Seim (2011); Athey *et al.* (2013); Bhattacharya *et al.* (2014). Both the entry and bidding (conditional on entry) behavior of favored and non-favored firms are expected to respond endogenously.

[52] Preferenced goods spanned many categories, including automobiles, clocks, various food products, medical equipment, pharmaceuticals, and textile and furs (see Online Appendix Table OA.5 for the full list).

paid for preferenced goods would have mirrored that for unpreferenced goods had the policy not been implemented. This is a priori plausible in our setting since the policy switches on and off multiple times. To investigate the parallel trends assumption directly, we estimate an event study version of (12) using data in a window starting three months before and ending four months after each year's preference list is published:

$$p_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \sum_{s=-3}^{4} \delta_s \text{Preferenced}_{gt} \times \mathbf{1}\left\{t - \text{ListMonth}_t = s\right\} + \varepsilon_{igt} \tag{13}$$

where terms are defined as above, and ListMonth_t is the month closest to month t in which a preference list is published. Panel A of Figure 6 shows the event study coefficients δ_s from estimation of (13). They are all close to zero and statistically indistinguishable from zero in the months leading up to the publication of the preference list, reassuring us that our difference-in-differences approach captures the causal effect of the preference policy.

In the full sample we do not observe goods' country of origin and so we cannot assess whether the preferences policy achieves the government's goal of channeling demand to domestic manufacturers. However, in the pharmaceuticals sample, we do observe where goods are manufactured. As seen in Column (5) of Table 6, we find that the preference policy increases the likelihood that auctions are won by suppliers whose goods are domestically manufactured.[53] The preferences policy thus achieves the government's goal of purchasing more Russian-made goods, at least in the pharmaceuticals sector.

Shifting demand towards domestic manufacturers using preferences in procurement comes at no direct cost to the government—in fact, bid preferences if anything *lower* prices paid in Russian procurement, as we saw in columns (1) and (3) of Table 6. These findings contrasts with the results from studies of similar preference policies in the U.S., which suggest that prices increase when the government introduces bid preferences for a subset of firms (see e.g. Marion, 2007; Krasnokutskaya & Seim, 2011). In this sense, our quasi-experimentally identified estimates of the average impact of Russia's "buy local" policy point toward the possibility that industrial policies of this form are more successful in countries like Russia, where bureaucrats and organizations are on average likely *less* effective than in richer countries. This foreshadows our findings in the next sub-section, where we analyze how the impact of the policy depends on the effectiveness of the policy implementers.

5.2 Bureaucracy-driven heterogeneity in the impact of policy change

In this section we turn to our analysis of how the tradeoffs between policies depend on the effectiveness of those who implement policy—individuals and organizations in the bureaucratic apparatus. To do so, we estimate differential impacts of the introduction of bid preferences when implemented by bureaucracies of different levels of effectiveness.

[53] 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 7, where we test the model's prediction for how bureaucratic effectiveness affects the probability of a domestic producer winning the contract when an auction takes place. 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).

Proposition 2 in Section 3 summarizes how we expect the effect of bid preferences to vary with the entry costs bureaucracies impose on suppliers. As Proposition 1 showed, these entry costs are key expected drivers of the differences in the prices ultimately achieved by different bureaucracies—their effectiveness. In Section 4 we calculated predictions of the price impacts of each bureaucrat and each organization (under the policy regime without bid preferences). We thus look for heterogeneity in the impact of the introduction of bid preferences when administered by bureaucracies of low versus high effectiveness.

To do so, we extend (12) to estimate heterogeneous treatment effects as follows:

$$\begin{aligned} p_{igt} = {} & \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta \text{Preferenced}_{gt} \times \text{PolicyActive}_t + \gamma \text{Preferenced}_{gt} \times \hat{\alpha}_b \\ & + \zeta \text{Preferenced}_{gt} \times \hat{\psi}_j + \eta \text{PolicyActive}_t \times \hat{\alpha}_b + \theta \text{PolicyActive}_t \times \hat{\psi}_j \\ & + \pi \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \nu \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt} \end{aligned} \tag{14}$$

The parameters of interest are π, the heterogeneity of the treatment effect by bureaucrat effectiveness, and ν, the heterogeneity of the treatment effect by organization effectiveness.

Table 7 shows the results.[54] We see that the small negative average price effect found in Table 6 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 paid higher prices (had a higher $\hat{\alpha}_b$) when there were no bid preferences, both in the full sample and the pharmaceuticals sample. Similarly, prices drop more for organizations who paid higher prices (had a higher $\hat{\psi}_j$). However, this effect is more muted in the full sample, and not statistically significant in the pharmaceuticals sample. This shows that heterogeneity in bureaucrats' type is more important than heterogeneity across organizations for the impact of policy changes, and conversely, suggests that while there are significant differences across policies in how impactful policy changes are when implemented by different bureaucrats, we may not see big differences in impact across implementing organizations.

Figure 7 presents estimates from extending equation (14) to capture heterogeneity in treatment effects less parametrically—separately by deciles of bureaucrat effectiveness $\hat{\alpha}_b$ and organization effectiveness $\hat{\psi}_j$. On the horizontal axis we plot the average effectiveness within the relevant decile, and on the vertical axis the corresponding treatment effect estimate. The figure reveals that the average treatment effect combines a price *increase* when the preferences policy is administered by effective bureaucrats and a larger price *decrease* among ineffective bureaucrats, consistent with our theoretical framework. The estimated treatment effects decrease in magnitude throughout the observed range, and the decline appears relatively linear. An important take-away from Figure 7 is that heterogeneity in how the preferences policy affects procurement outcomes is not concentrated among especially effective or ineffective bureaucracies, but seen throughout the distribution of effectiveness.

The estimated heterogeneity in policy impact does not appear to be driven by potential confounds like mean reversion or differences in seasonality across different types of bureaucrats and organizations. Panels B and C of Figure 6 show event studies of the introduction of bid preferences separately for the

[54] In Table 7 we use the covariance shrunk estimates of the bureaucrat and organization effects. Online Appendix Table OA.10 uses the raw fixed effect estimates and shows very similar results.

lowest and highest quartile bureaucrats (Panel B) and organizations (Panel C). The figure shows no discernible trend in prices before the introduction of the bid preferences. In particular, prices do not diverge in the two groups before the policy is implemented. After bid preferences are introduced, the average price achieved by effective bureaucrats begins to increase, while that achieved by ineffective bureaucrats begins to decrease, with the difference increasing over time. These patterns provide compelling evidence that the estimates in Table 7 capture the causal differential of interest.[55]

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 7 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 increase in auctions administered by ineffective procurers.[56]

Finally, Proposition 2 predicts differential impacts of the policy on the likelihood that an auction is won by a bidder offering locally manufactured goods. We can test this prediction in the pharmaceuticals sample, where we observe the origin of the winning product. In columns (1) and (3) we saw significant heterogeneity in the estimated impact on prices by bureaucrat effectiveness, but not by organizations' effectiveness, and so we focus on bureaucrat heterogeneity in the impact on the origin of the goods. Column (5) of Table 7 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.

In combination with the differential change in price under bid preferences shown in columns (1) and (3), 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 outlined in Section 3. For organizational effectiveness, the policy design implications of our findings in Table 7 are less pronounced.

Two simple calculations help to illustrate the implications of these findings for policy design. First, our estimates suggest that the Russian bid preference policy saved the government 17.5 percent when it was implemented by the least effective quartile of bureaucrats, but only saved the government 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.

Second, our estimates show that the variation in performance attributable to a heterogeneous bureaucracy will depend on the nature of the policy being implemented. When the policy regime does not feature bid preferences, the standard deviation of the impact of bureaucrats and organizations on prices is 0.525. However, when these same bureaucrats are asked to implement a policy featuring bid prefer-

[55]In Panel C of Figure 6, which studies organizations, the change in the relative performance of effective and ineffective organizations after preferences turn on is less visually clear, consistent with the weaker effect for organizations in Table 7.

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

ences, our estimates suggest that this goes down to 0.461, a reduction of 12 percent. This illustrates 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.

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 state enterprise 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. Studying the impact of a "buy local" policy that provides bid preferences for locally manufactured goods, we find that participation increases and prices decrease when the policy is implemented by less effective bureaucrats, while performance deteriorates 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. For example, our estimates imply that if the least effective quartile of bureaucrats and organizations had 75th percentile effectiveness, the Russian government would save around USD 13 billion each year. The large magnitude of the procurer effects we estimate suggests that the middle, bureaucratic tier of the state is a hugely important, though often overlooked, determinant of state effectiveness.

Second, our findings imply that the nature of the policy regime in place determines the extent to which differences in the effectiveness of procurers manifest themselves in differences in public sector output. In turn, this suggests that there are likely large gains from tailoring the design of policy to the effectiveness of the bureaucracy tasked with implementing it. In particular, policies that are suboptimal when state effectiveness is high may become second-best optimal when state effectiveness is low. For example, our estimates imply that the Russian bid preference policy saved the government 17.5 percent of annual procurement expenses when it was implemented by the least effective quartile of procurers, but only 0.7 percent when it was implemented by the most effective quartile of procurers. Such dependence of policies' impact on state effectiveness may be part of the reason why many policies work well in some countries or regions and poorly in others.

A final implication is methodological. This paper shows that in order to extrapolate an average treatment effect of a public policy estimated in one setting to another setting, knowledge of differences in policy implementer effectiveness across the two settings is essential. We demonstrate how bureaucratic effectiveness can be estimated in baseline data and then used in the estimation of heterogeneous treatment effects to guide such extrapolation.

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*.

ACEMOGLU, DARON, GARCIA-JIMENO, CAMILO, & ROBINSON, JAMES A. 2015. State Capacity and Economic Development: A Network Approach. *American Economic Review*.

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

ALVAREZ, JORGE, ENGBOM, NIKLAS, & MOSER, CHRISTIAN. 2018. Firms and the Decline of Earnings Inequality in Brazil. *American Economic Journal: Macroeconomics*, **10**(1), 149–189.

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.

ASHRAF, NAVA, BANDIERA, ORIANA, & LEE, SCOTT S. 2014a. *Do-gooders and Go-getters: Career Incentives, Selection, and Performance in Public Service Delivery*. mimeo, Harvard.

ASHRAF, NAVA, BANDIERA, ORIANA, & JACK, KELSEY. 2014b. No Margin, No Mission? A Field Experiment on Incentives for Public Services Delivery. *Journal of Public Economics*.

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

BAI, YING, & JIA, RUIXUE. 2016. Elite Recruitment and Political Stability: The Impact of the Abolition of China's Civil Service Exam. *Econometrica*.

BAJARI, PATRICK, HOUGHTON, STEPHANIE, & TADELIS, STEVE. 2014. Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs. *American Economic Review*.

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. 2016. *The Costs of Bureaucratic Rigidity: Evidence from the Indian Administrative Service*. Working Paper.

BESLEY, T., & PERSSON, T. 2010. State capacity, conflict and development. *Econometrica*, **78**, 1–34.

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

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

BHATTACHARYA, VIVEK, ROBERTS, JAMES W., & SWEETING, ANDREW. 2014. Regulating bidder participation in auctions. *RAND Journal of Economics*.

BIRD, RICHARD, & GENDRON, PIERRE-PASCAL. 2007. *The VAT in developing and transitional countries*. Cambridge University Press.

BLOOM, NICHOLAS, PROPPER, CAROL, SEILER, STEPHAN, & REENEN, JOHN VAN. 2015. The Impact of Competition on Management Practices in Public Hospitals. *The Review of Economic Studies*.

BLOOM, NICHOLAS, BENDER, STEFAN, CARD, DAVID, REENEN, JOHN VAN, & WOLTER, STEPHANIE. 2018. Management practices, workforce selection and productivity. *Journal of Labor Economics*.

BLOOM, NICHOLAS, SONG, JAE, PRICE, DAVID, GUVENEN, FATIH, & VON WACHTER, TILL. forthcoming. Firming up inequality. *Quarterly Journal of Economics*.

BOLD, TESSA, KIMENYI, MWANGI, MWABU, GERMANO, NG'ANG'A, ALICE, & SANDEFUR, JUSTIN. forthcoming. Scaling Up What Works: Experimental Evidence on External Validity in Kenyan Education. *Journal of Public Economics*.

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*.

BURGESS, ROBIN, HANSEN, MATTHEW, OLKEN, BENJAMIN, POTAPOV, PETER, & SIEBER, STEFANIE. 2012. The political economy of deforestation in the tropics. *Quarterly Journal of Economics*, **127**, 1707–1754.

CALLEN, MICHAEL, GULZAR, SAAD, HASANAIN, ALI, KHAN, YASIR, & REZAEE, ARMAN. 2016a. *Personalities and Public Employee Performance: Experimental Evidence from Pakistan*. mimeo UCSD.

CALLEN, MICHAEL, GULZAR, SAAD, HASANAIN, ALI, & KHAN, YASIR. 2016b. *The Political Economy of Public Employee Absence: Experimental Evidence from Pakistan*. mimeo UCSD.

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*.

CARDOSO, ANA RUTE, GUIMARÃES, PAULO, & PORTUGAL, PEDRO. 2016. What Drives the Gender Wage Gap? A Look at the Role of Firm and Job-Title Heterogeneity. *Oxford Economic Papers*, **68**, 506–524.

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.

CHEREMUKHIN, ANTON, GOLOSOV, MIKHAIL, GURIEV, SERGEI, & TSYVINSKI, ALEH. 2016. The Industrialization and Economic Development of Russia through the Lens of a Neoclassical Growth Model. *The Review of Economic Studies*.

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

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

CHONG, ALBERTO, LA PORTA, RAFAEL, DE SILANES, FLORENCIO LOPEZ, & SHLEIFER, ANDREI. 2014. Letter Grading Government Efficiency. *Journal of the European Economic Association*.

COVIELLO, DECIO, & GAGLIARDUCCI, STEFANO. 2017. Tenure in Office and Public Procurement. *American Economic Journal: Economic Policy*.

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*.

DEATON, ANGUS. 2010. Instruments, Randomization and Learning about Development. *Journal of Economic Literature*, 424–455.

DECAROLIS, FRANCESCO. 2014. Awarding Price, Contract Performance and Bids Screening: Evidence from Procurement Auctions. *American Economic Journal: Economic Policy*.

DECAROLIS, FRANCESCO, & PALUMBO, GIULIANA. 2015. Renegotiation of Public Contracts: An Empirical Analysis. *Economic Letters*.

DECAROLIS, FRANCESCO, SPAGNOLO, GIANCARLO, & PACINI, RICCARDO. 2016. *Past Performance and Procurement Outcomes*. NBER Working Paper no. 22814.

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. 2016. *From Local to Global: External Validity in a Fertility Natural Experiment*. NBER Working Paper No. 21459.

DESERRANNO, ERIKA. forthcoming. Financial Incentives as Signals: Experimental Evidence from the Recruitment of Village Promoters in Uganda. *American Economic Journal: Applied Economics*.

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. forthcoming. The Value of Discretion in the Enforcement of Regulation: Experimental Evidence and Structural Estimates from Environmental Inspections in India. *Econometrica*.

EASTERLY, WILLIAM, & PENNINGS, STEVEN. 2017. *Shrinking dictators: how much economic growth can we attribute to national leaders?* mimeo, NYU.

EECKHOUT, JAN, & KIRCHER, PHILIPP. 2011. Identifying Sorting–In Theory. *The Review of Economic Studies*.

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

EPPLE, DENNIS. 1987. Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products. *Journal of Political Economy*, **95**, 59–80.

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*.

FOSTER, LUCIA, HALTIWANGER, JOHN, & SYVERSON, CHAD. 2009. Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability. *American Economic Review*, **98**.

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

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.

GENTZKOW, MATTHEW, & SHAPIRO, JESSE. 2010. What Drives Media Slant? Evidence From US Daily Newspapers. *Econometrica*, **78**, 35–71.

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

GOLDSCHMIDT, DEBORAH, & SCHMIEDER, JOHANNES. 2017. The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure. *Quarterly Journal of Economics*.

GRILICHES, ZVI (ed). 1971. *Price Indexes and Quality Change*. Cambridge: Harvard University Press.

GULZAR, SAAD, & PASQUALE, BENJAMIN. 2017. Politicians, Bureaucrats, and Development: Evidence from India. *American Political Science Review*, **111**, 162–183.

HANNA, REMA, & WANG, SHING-YI. 2017. Dishonesty and Selection into Public Service: Evidence from India. *American Economic Journal: Economic Policy*.

HANSEN, STEPHEN, MCMAHON, MICHAEL, & PRAT, ANDREA. 2018. Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach. *Quarterly Journal of Economics*.

HANSMAN, CHRISTOPHER, HJORT, JONAS, & LEÓN, GIANMARCO. 2017. 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.

HECKMAN, JAMES J. 2010. Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy. *Journal of Economic Literature*, **48**, 356–98.

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

HORTACSU, ALI, & SYVERSON, CHAD. 2007. Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices. *Journal of Political Economy*.

JANKE, KATHARINA, PROPPER, CAROL, & SADUN, RAFFAELLA. 2018. *The Impact of CEOs in the Public Sector: Evidence from the English NHS*. HBS Working Paper no. 18-075.

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.

KLEPIKOVA, E. A. 2017. Job Mobility of Public and Private Sector Workers. *Sociological Research*.

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

KRISHNA, VIJAY. 2010. *Auction Theory (second edition)*. Elsevier.

KRYLOVA, DINA V., & SETTLES, ALEXANDER. 2011. *Procurement Reform in Russia: Implications for the Fight Against Corruption*. Economic Reform Feature Service.

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

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

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

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

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

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

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

MENDES, RUTE, VAN DEN BERG, GERARD J., & LINDEBOOM, MAARTEN. 2010. An Empirical Assessment of Assortative Matching in the Labor Market. *Labour Economics*, **17**, 919–29.

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

NEYMAN, JERZY, & SCOTT, ELIZABETH L. 1948. Consistent Estimates Based on Partially Consistent Observations. *Econometrica*, **16**, 1–32.

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

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

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.

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

RODRIK, DANI. 2009. Diagnostics before Prescription. *Journal of Economic Perspectives*.

ROSEN, SHERWIN. 1974. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, **82**, 34–55.

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.

SCHOTT, PETER K. 2004. Across-Product Versus Within-Product Specialization in International Trade. *Quarterly Journal of Economics*, **119**, 647–678.

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

SHELEF, ORIE, & NUGYEN-CHYUNG, AMY. 2015. *Competing for Labor through Contracts: Selection, Matching, Firm Organization and Investments*. mimeo, Stanford.

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. 2017. Businesspeople in Elected Office: Identifying Private Benefits from Firm-Level Returns. *American Political Science Review*.

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

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

VIVALT, EVA. 2017. *How Much Can We Generalize from Impact Evaluations?* Mimeo: Australian National University.

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

XU, GUO. forthcoming. The Costs of Patronage: Evidence from the British Empire. *American Economic Review*.

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*.

YAO, YANG, & ZHANG, MUYANG J. 2015. Subnational leaders and economic growth: evidence from Chinese cities. *J Econ Growth*.

Figure 1: Procurement Process Flow-Chart

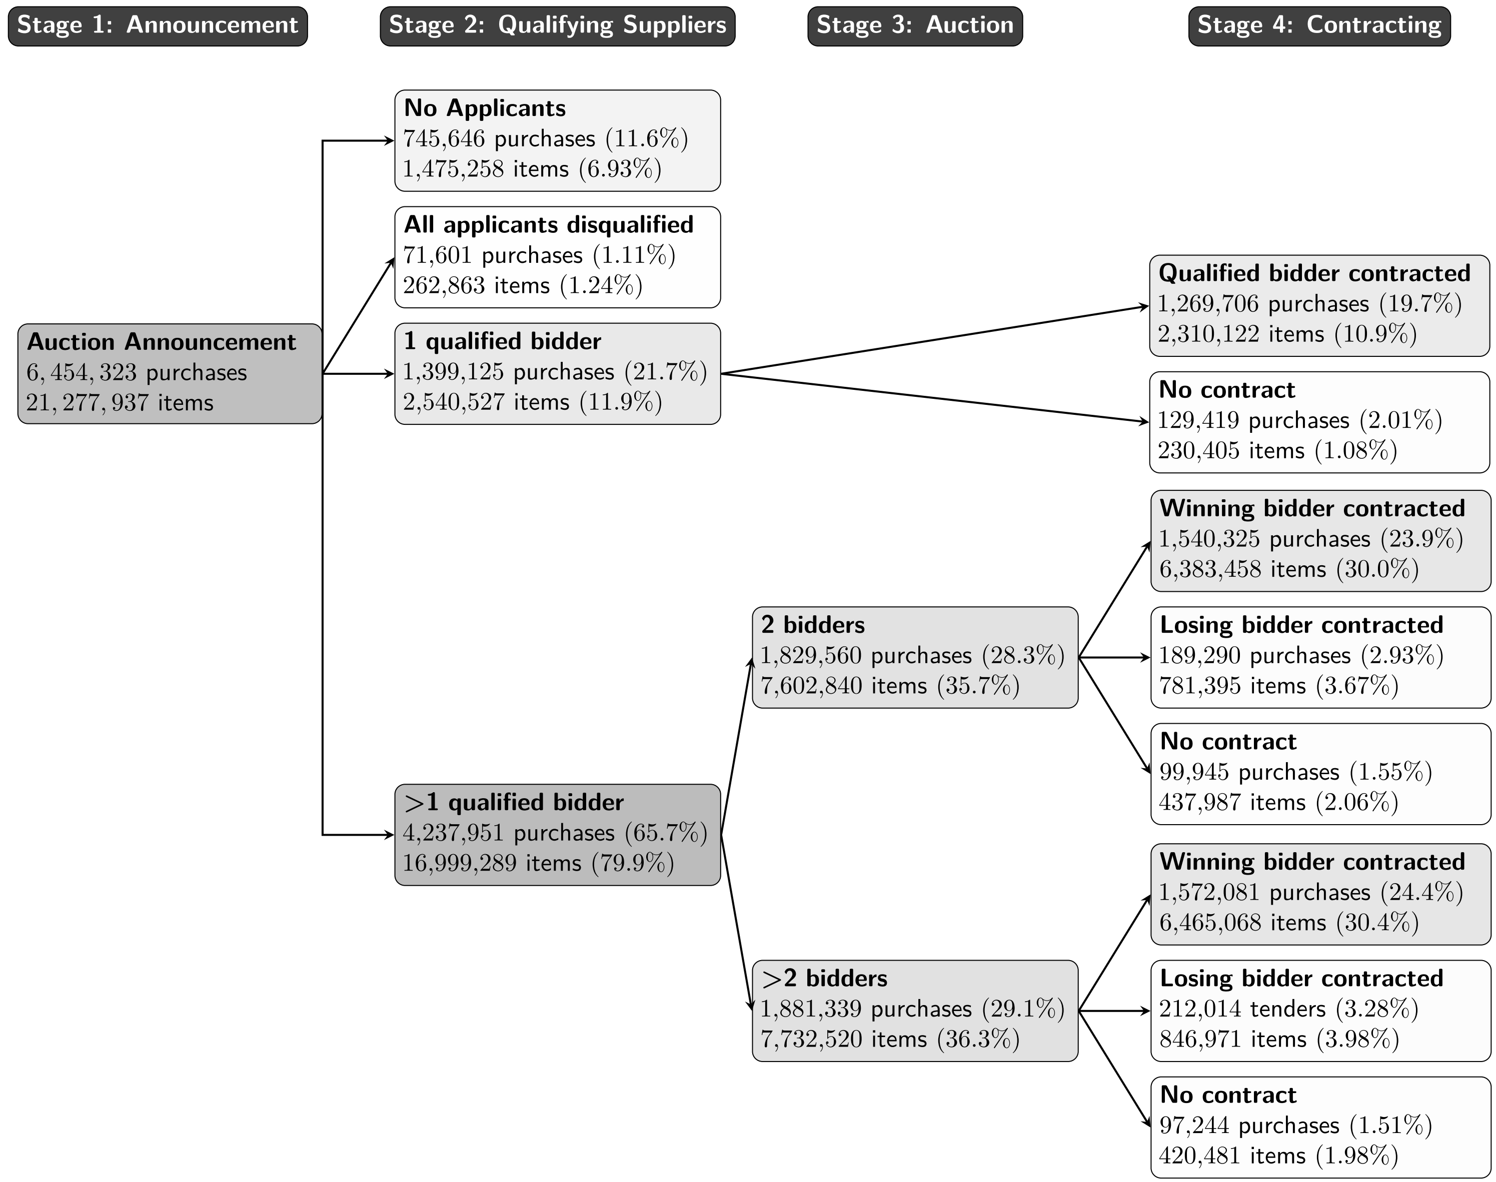

This figure lays out the stages of the process public procurement purchases of off-the-shelf goods through electronic auctions follow in Russia. Numbers are based on all purchases made under laws 94 and 44 in 2011-2016. The stages are described in detail in Sub-section 2.1.

FIGURE 2: EVENT STUDY OF PROCUREMENT PRICES AROUND TIMES ORGANIZATIONS SWITCH BUREAUCRATS

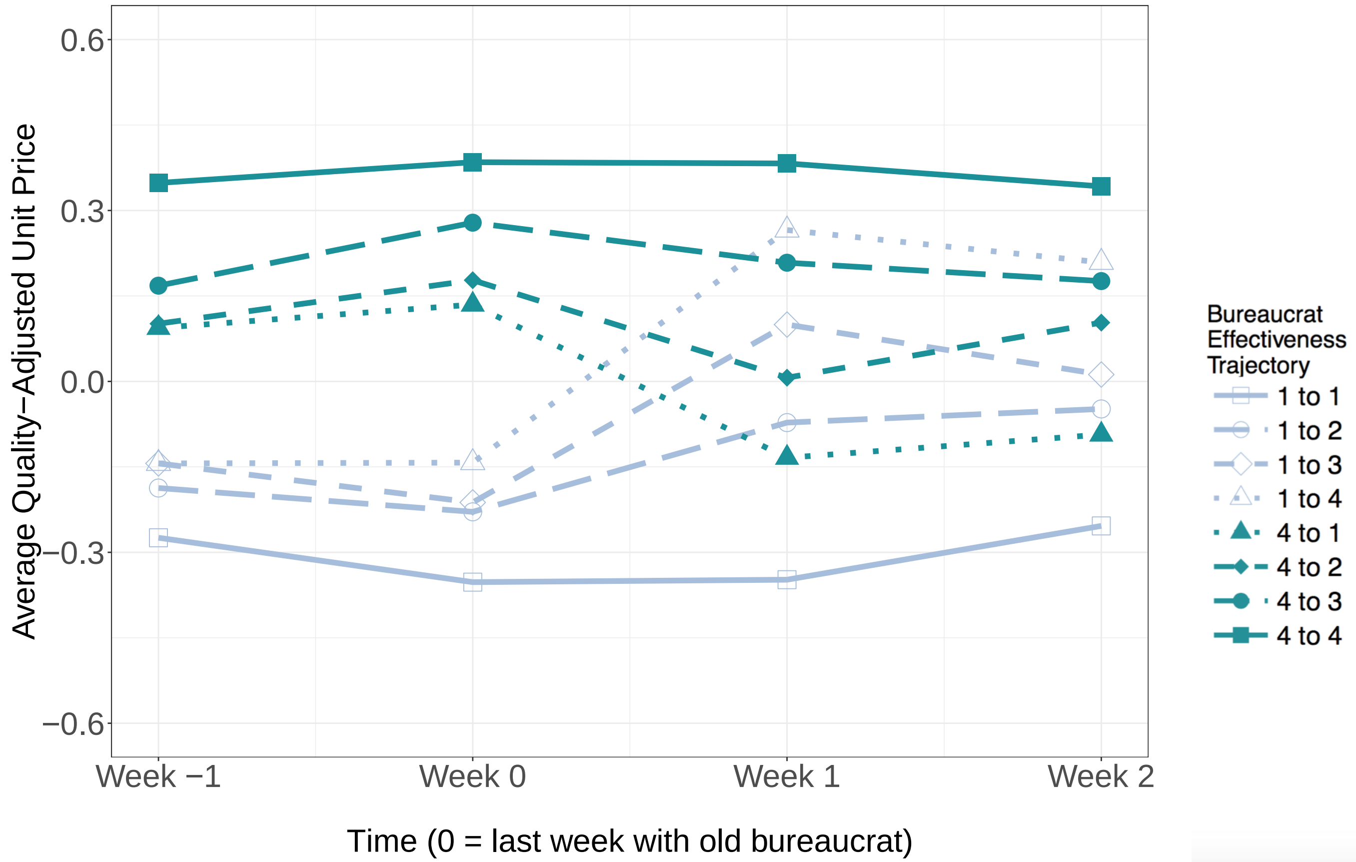

The figure shows time trends in prices around the time that organizations switch which bureaucrat makes purchases on their behalf. The horizontal axis measures fortnight on which bureaucrat-organization pairs work together, with time 0 being the last fortnight on which the organization works with the 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 3: CRUDE COUNTERFACTUALS

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

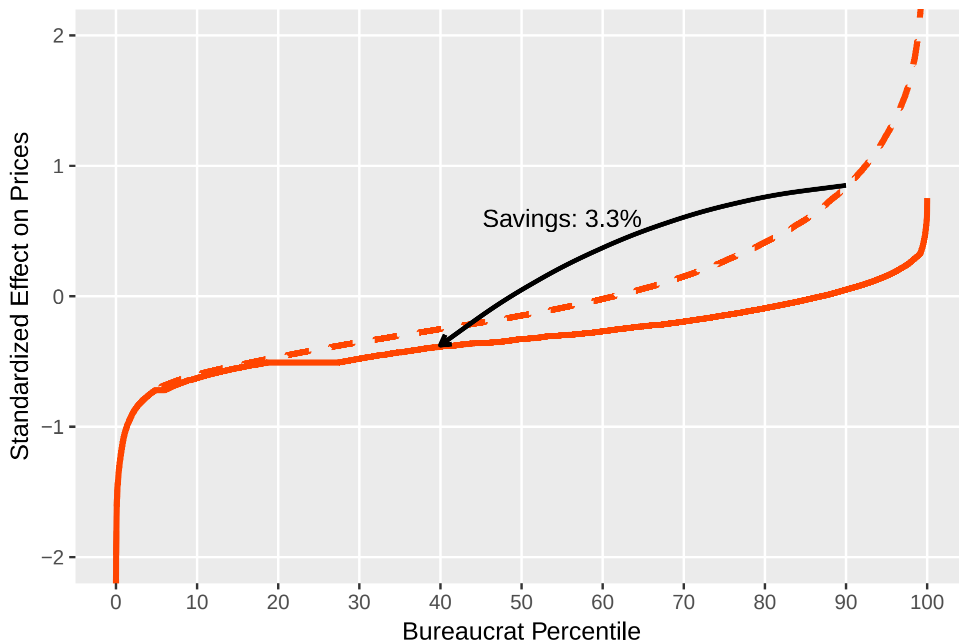

Panel B: Moving Least Effective 25% of Bureaucrats and Organizations to 75th Percentile Effectiveness

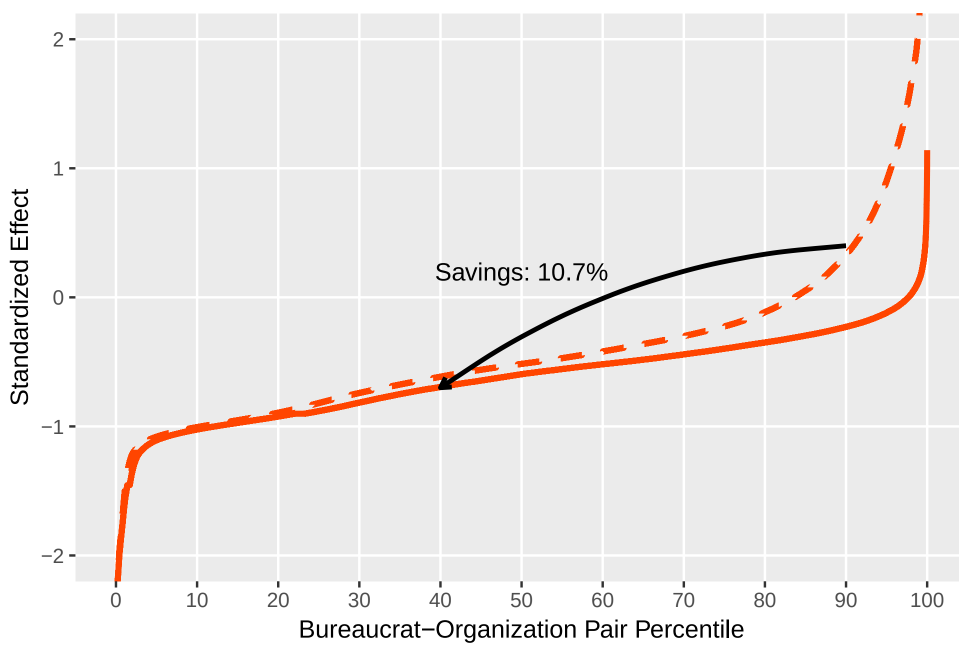

The figure shows the impact of two counterfactual scenarios on the distribution of our estimated price effects. Panel A considers moving all bureaucrats above the 75th percentile of their connected set's distribution of covariance shrunken price effects down to their connected set's 25th percentile. 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 4: CORRELATES OF BUREAUCRAT EFFECTIVENESS

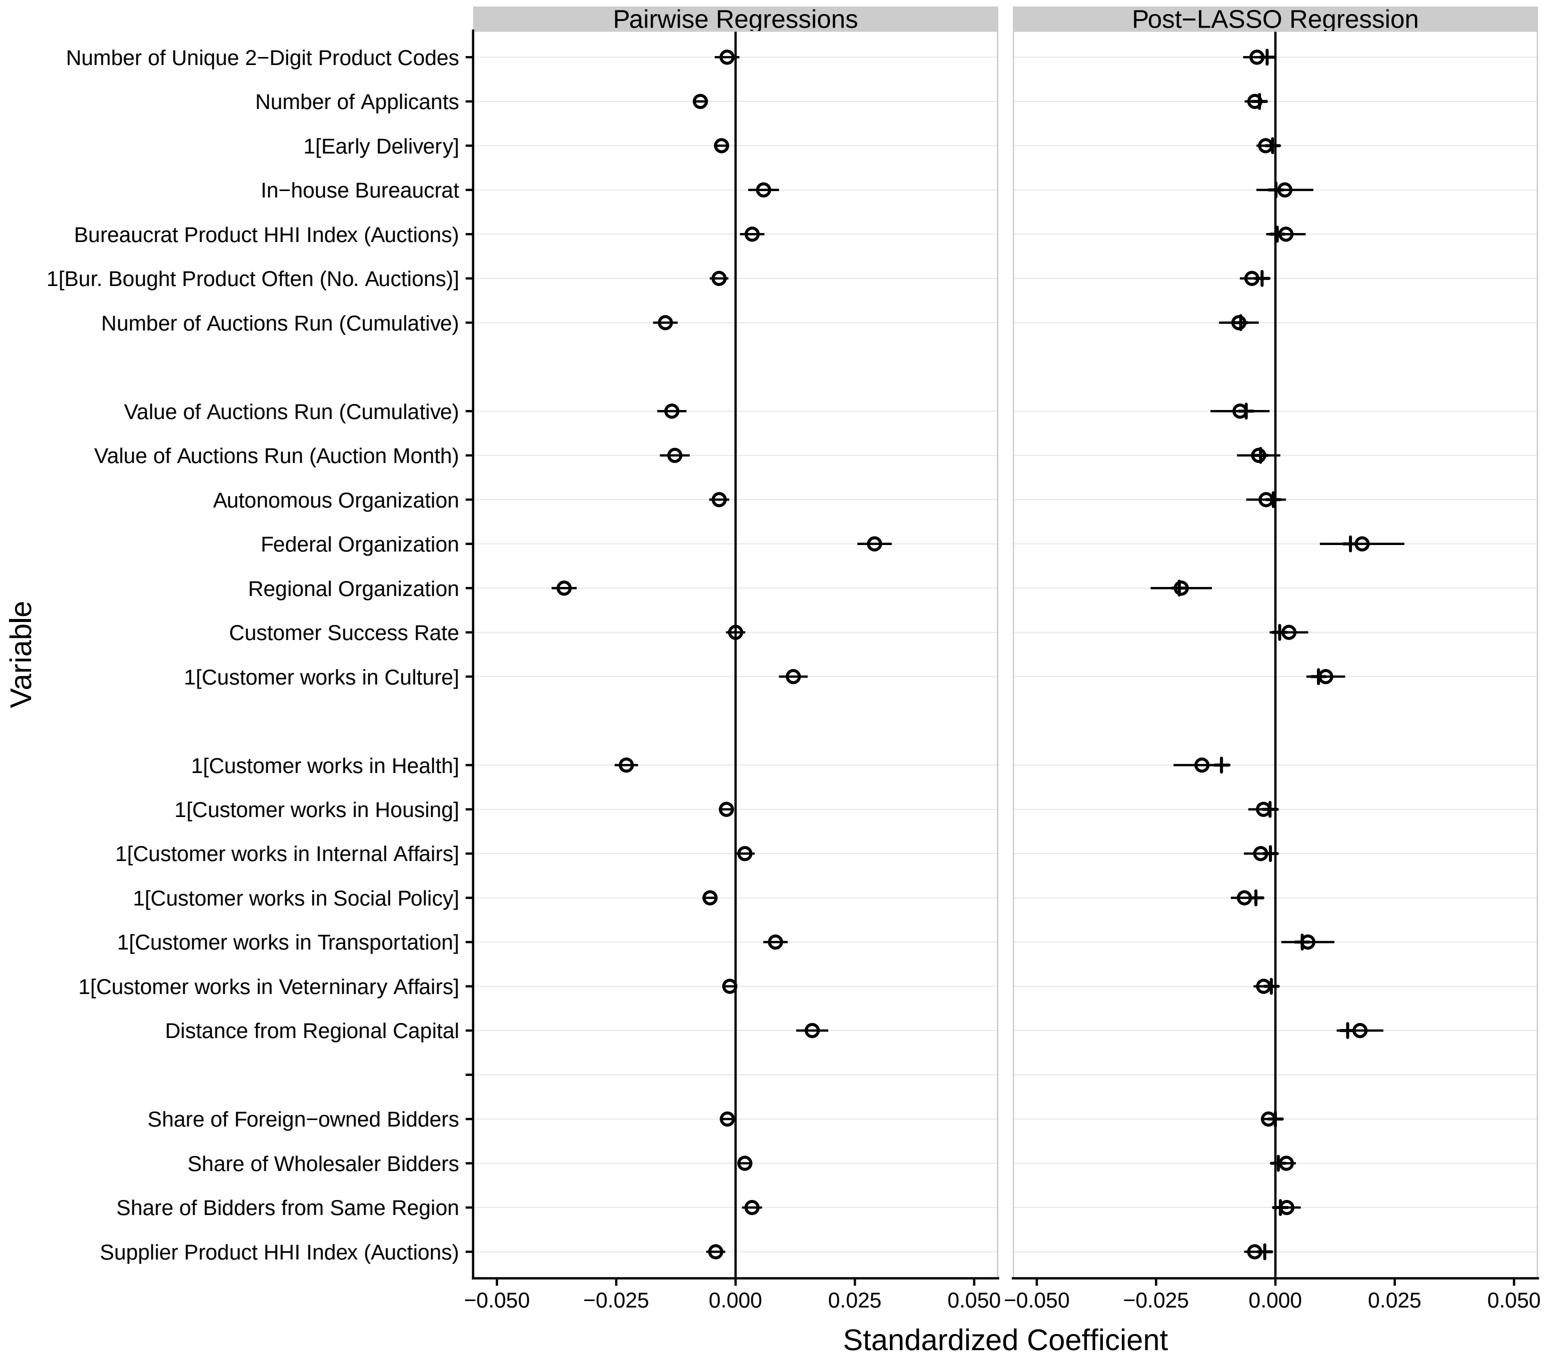

Figure 5: Correlates of Organization Effectiveness

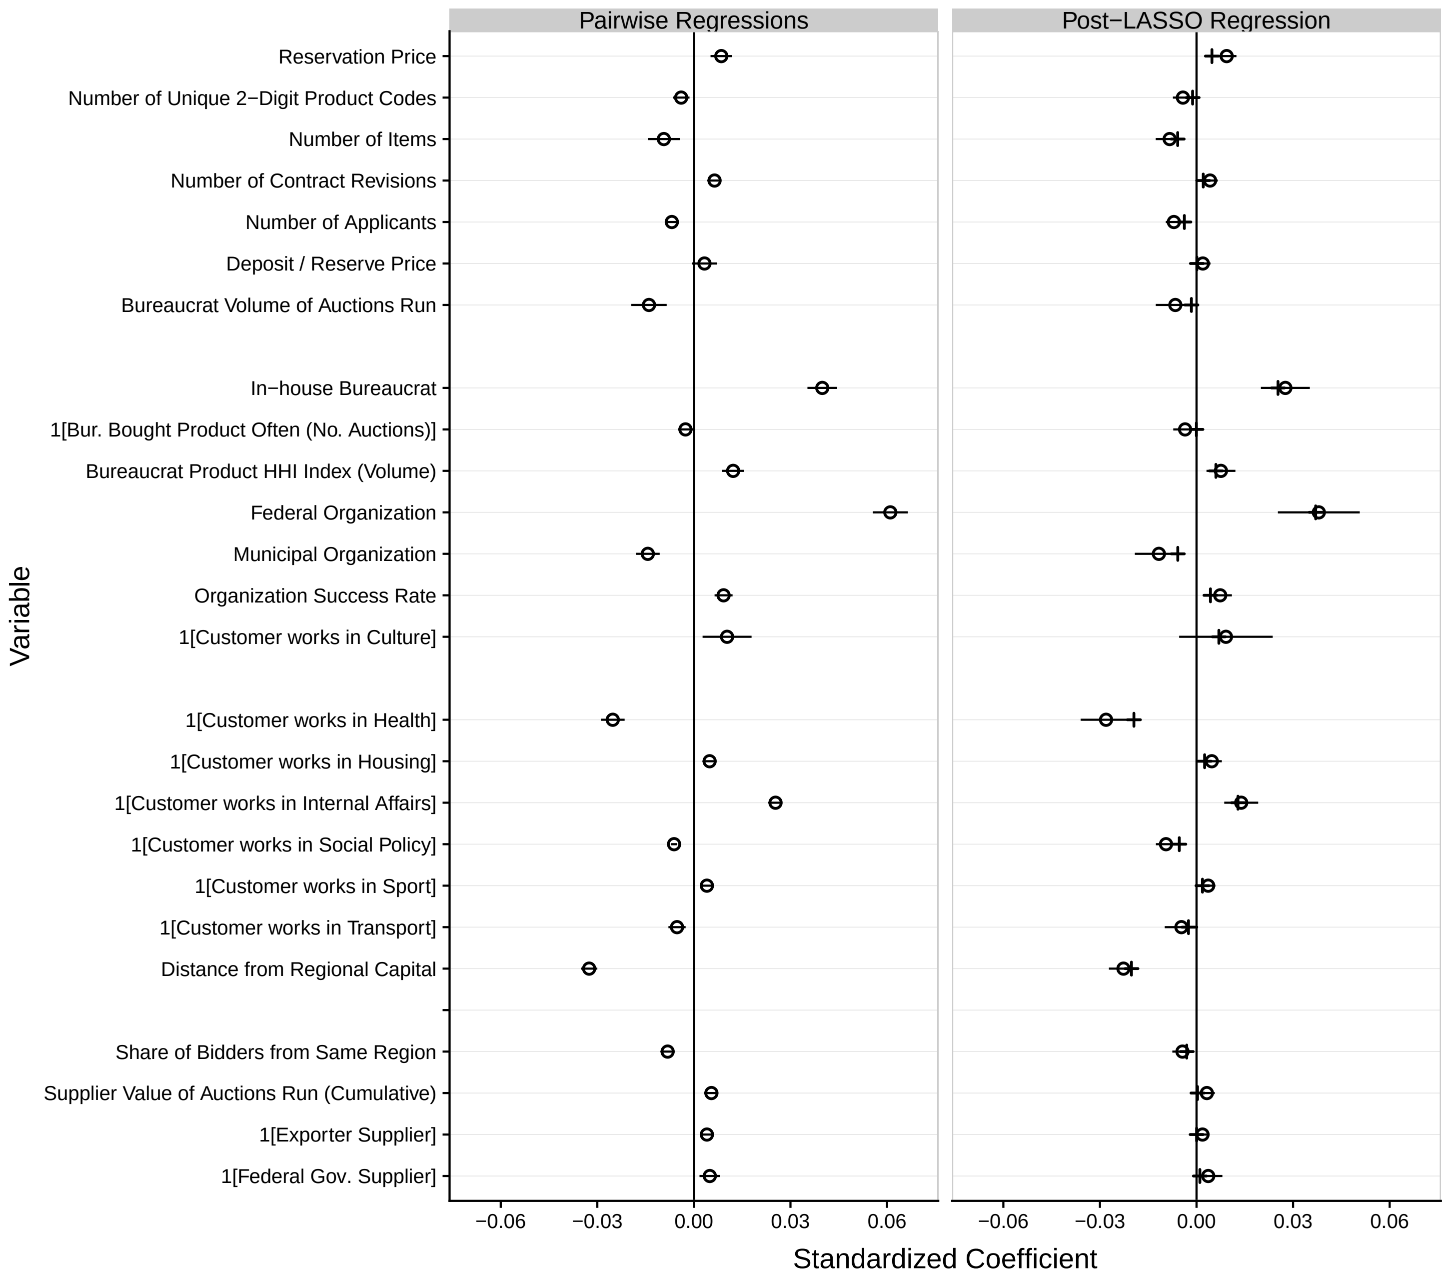

The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The circles in right

Figure 6: Effect of Bid Preferences for Domestic Manufacturers: Graphical Evidence

Panel A: Event Study of Average Policy Impact

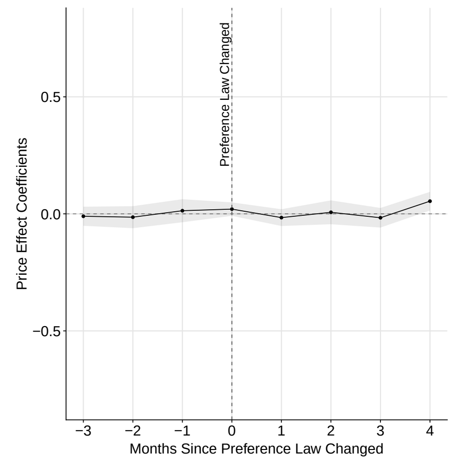

Panel B: Bureaucrat Heterogeneity

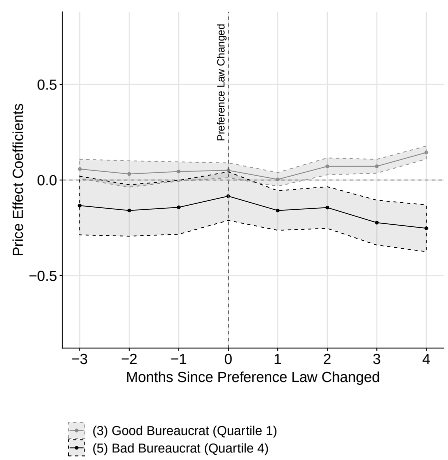

Panel C: Organization Heterogeneity

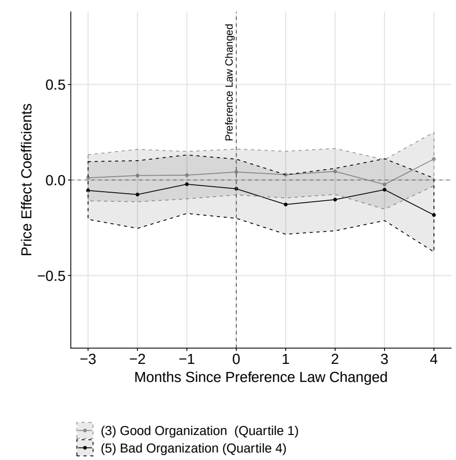

The figure shows a graphical analysis of the preferences policy over the period of study. The x-axis is measured in the number of months preceding or following the activation of the annual preferences laws in 2011, 2012, 2013, and 2014. The dotted vertical lines indicates when the policy was became active. The y-axis in each plot shows the month-specific coefficients from estimation of equation (13): $p_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \sum_{s=-3}^{4} \delta_s \text{Preferenced}_{gt} \times \mathbf{1}\left\{t - \text{ListMonth}_t = s\right\} + \varepsilon_{igt}$, where p_{igt} is the price paid for purchase i of good g in month t. Preferenced_{gt} is a dummy indicating that g is on the preferences list in the year month t falls within, and ListMonth_t is the month closest to month t in which a preference list is published. $\mathbf{X}_{igt}$ are the same controls we use in Section 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. The top panel shows the average policy impact for the full sample. The bottom panels shows the event studies separately for the highest and lowest quartiles of bureaucrat effectiveness (Panel B) and organization effectiveness (Panel C).

FIGURE 7: NON-PARAMETRIC HETEROGENEITY OF EFFECT OF BID PREFERENCES FOR DOMESTIC PRODUCERS BY BUREAUCRAT AND ORGANIZATION EFFECTIVENESS

PANEL A: HETEROGENEITY IN EFFECT ON PRICE BY BUREAUCRAT EFFECTIVENESS

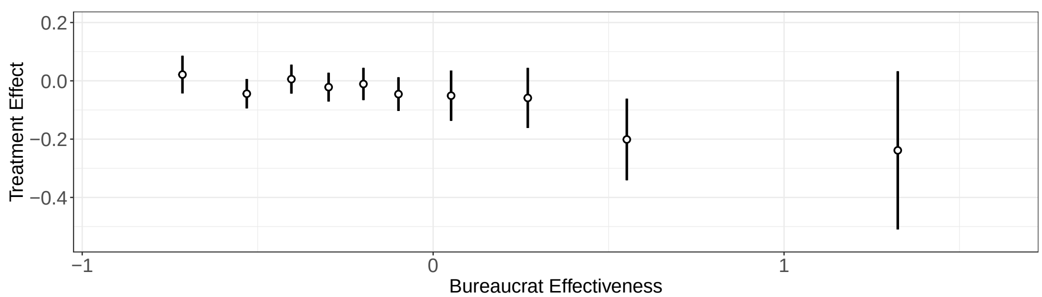

PANEL B: HETEROGENEITY IN EFFECT ON PRICE BY ORGANIZATION EFFECTIVENESS

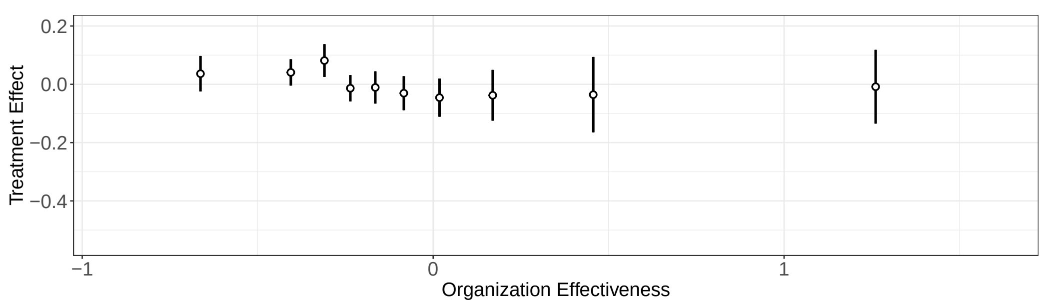

The figure shows results from a non-parametric extension of the triple-differences equation (14):

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \gamma\text{Preferenced}_{gt} \times \hat{\alpha}_b + \zeta\text{Preferenced}_{gt} \times \hat{\psi}_j + \eta\text{PolicyActive}_t \times \hat{\alpha}_b$$
$$+ \theta\text{PolicyActive}_t \times \hat{\psi}_j + \pi\text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \zeta\text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt}$$

Bureaucrat and Organization effects are in this figure binned into deciles and the decile dummies interacted with the treatment indicators PolicyActive_t and Preferenced_{gt} and their interaction. On the horizontal axis we plot the average effectiveness within the relevant decile. On the vertical axis we show the corresponding treatment effect estimate.

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	(6) 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 unpreferenced auctions. Analysis Sample denotes all unpreferenced auctions in connected sets that fulfill three restrictions: singleton bureaucrat-organization, bureaucrat-good, and organization-good pairs are removed; each procurer (bureaucrats and organizations) implements a minimum of five purchases; and connected sets have at least three bureaucrats and organizations. With Bid Preferences denotes all preferenced auctions that fulfill the same three restrictions. Organizations working in Education include schools, universities, pre-schools, and youth organizations. Organizations working in Internal Affairs include police, emergency services, local administration, taxes, and transportation. Organizations working in Agriculture or the Environment include environmental protection funds, agricultural departments and nature promotion agencies. The Other category includes funds, monitoring agencies, and land cadasters, among many others. All sums are measured in billions of US dollars at an exchange rate of 43 rubles to 1 US dollar.

Table 2: Event Studies Summary Statistics

Origin/destination Quartile*	Number of Moves (1)	Number of Observations (2)	Mean Log Residuals of Bureaucrat Movers: Week -1 (3)	Week 0 (4)	Week 1 (5)	Week 2 (6)	Mean Weeks Betw. Cols: (3)-(4) (7)	(4)-(5) (8)	(5)-(6) (9)
1 to 1	5,605	240,974	-0.274	-0.359	-0.351	-0.254	12.518	28.253	12.119
1 to 2	5,442	224,305	-0.187	-0.224	-0.092	-0.048	12.431	23.619	12.538
1 to 3	3,393	136,756	-0.144	-0.192	0.061	0.012	13.780	28.949	13.366
1 to 4	1,736	70,098	-0.144	-0.139	0.245	0.209	14.171	35.299	16.808
2 to 1	5,604	229,646	-0.043	-0.094	-0.211	-0.168	13.056	27.703	13.575
2 to 2	9,659	484,044	-0.034	-0.050	-0.010	-0.033	12.200	25.723	11.923
2 to 3	6,122	277,266	-0.035	-0.043	0.087	0.047	12.789	28.965	14.853
2 to 4	2,243	87,754	0.066	0.010	0.246	0.186	11.985	35.071	15.127
3 to 1	3,173	132,081	0.015	0.003	-0.115	-0.159	15.436	26.144	12.123
3 to 2	5,822	262,335	0.001	0.043	0.017	0.008	13.601	24.306	12.060
3 to 3	5,608	239,169	0.030	0.089	0.117	0.112	15.337	27.186	14.642
3 to 4	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
4 to 2	1,728	73,812	0.101	0.163	0.013	0.103	15.639	29.743	12.104
4 to 3	2,159	90,069	0.168	0.287	0.212	0.176	15.145	29.638	12.552
4 to 4	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 (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. This regression regresses the price achieved in an auction on log quantity, good fixed effects, month fixed effects, interactions between 2-digit HS product categories, years, regions, and lot size, as explained in detail in Sub-section 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 half-year as the bureaucrat in question.

TABLE 3: 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 (10). 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 (9): $p_i = \mathbf{X}_i \boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates. Columns 2 and 4 show standard errors of the estimates in columns 1 and 3, respectively, estimated by bootstrapping 100 times. Column 5 uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2, and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2/(\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column 1, and analogously for the organization effects. Column 6 shows our preferred estimates, which form predictions of the bureaucrat and organization effects that minimize the expected sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates. Formally, the covariance shrinkage predictors solve $\min_{\boldsymbol{\Lambda}} \mathbb{E}\left[\left(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)' \left(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)\right]$ where $\hat{\boldsymbol{\theta}}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 4.2.

Table 4: Robustness to Restricting to Pharmaceuticals Subsample with Barcode Information

	Fixed Effects (1)	(s.e.) (2)	Split Sample (3)	(s.e.) (4)	Shrinkage (5)	Covariance Shrinkage (6)
(1) s.d. of Bureaucrat Effects (across burs)	0.209	(0.0074)	0.216	(0.00728)	0.103	0.0632
(2) s.d. of Organization Effects (across orgs)	0.165	(0.00627)	0.173	(0.00634)	0.078	0.0481
(3) s.d. of Connected Set Effects (across CS)	0.228	(0.0116)	0.152	(0.0123)	0.191	0.191
(4) s.d. of Bureaucrat Effects (across items)	0.124	(0.0122)	0.132	(0.0105)	0.0885	0.0515
(5) s.d. of Organization Effects (across items)	0.129	(0.0131)	0.137	(0.0108)	0.0815	0.0416
(6) s.d. of Connected Set Effects (across items)	0.152	(0.00841)	0.129	(0.00794)	0.137	0.137
(7) Bur-Org Effect Correlation (across items)	-0.333	(0.0783)	-0.164	(0.0439)	-0.264	-0.0886
(8) s.d. of Total Bur + Org Effects (across items)	0.190	(0.00779)	0.179	(0.00679)	0.163	0.147
(9) s.d. of log unit price	1.915		1.915		1.915	1.915
(10) s.d. of log unit price | good, month	0.319		0.319		0.319	0.319
(11) Adjusted R-squared	0.997		0.997		0.997	0.997
(12) Number of Bureaucrats	2,501		2,501		2,501	2,501
(13) Number of Organizations	1,880		1,880		1,880	1,880
(14) Number of Bureaucrat-Organization Pairs	8,112		8,112		8,112	8,112
(15) Number of Connected Sets	129		129		129	129
(16) Number of Observations	181,493		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 (10). The sample used is the Pharmaceuticals-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 items, effectively weighting the estimates in rows 1–3 by the number of purchases they conduct. Column 1 uses the fixed effect estimates from equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates. Columns 2 and 4 show standard errors of the estimates in columns 1 and 3, respectively, estimated by bootstrapping 100 times. Column 5 uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2, and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2/(\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column 1, and analogously for the organization effects. Column 6 shows our preferred estimates, which form predictions of the bureaucrat and organization effects that minimize the expected sum of the mean-squared errors of the predictions and take into account the covariance of the estimation errors, estimated from the bootstrapped estimates. Formally, the covariance shrinkage predictors solve $\min_{\boldsymbol{\Lambda}} \mathbb{E}\left[\left(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)'\left(\boldsymbol{\theta} - \boldsymbol{\Lambda}\hat{\boldsymbol{\theta}}\right)\right]$ where $\hat{\boldsymbol{\theta}}$ is the vector of estimated bureaucrat and organization fixed effects. All methods are described fully in Section 4.2.

Table 5: 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 (10) (see notes to Table 3 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 6: Bid Preferences for Locals Increase Domestic Winners With Limited Impact on Prices or Participation

	All Products		Pharmaceuticals		
	Log Price	No. Bidders	Log Price	No. Bidders	Domestic Winner
	(1)	(2)	(3)	(4)	(5)
log Standardized Quantity	-0.308^{***}	0.043^{***}	-0.030^{***}	0.012^{***}	0.003^{***}
	(0.016)	(0.003)	(0.003)	(0.004)	(0.001)
Preferenced * Policy Active	-0.019^{*}	-0.059^{*}	-0.026	-0.025	0.049^{***}
	(0.012)	(0.035)	(0.017)	(0.027)	(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

*** p<0.01, ** p<0.05, * p<0.1 This table estimates the Intent to Treat (ITT) from equation (12): $p_{igt} = \mathbf{X}_{igt}\boldsymbol{\beta} + \mu_g + \lambda_t + \delta \text{Preferenced}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt}$. The With Bid Preferences samples summarized in columns (3) and (6) of Table 1 are used, i.e. the 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 Preferenced = 1 if the type of good appears on the list of goods covered by the preferences policy for that year. In the Pharmaceuticals sample, Preferenced = 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. 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 7: BID PREFERENCES ARE MORE EFFECTIVE WHEN IMPLEMENTED BY LESS EFFECTIVE BUREAUCRATS

	All Products		Pharmaceuticals		
	Log Price	No. Bidders	Log Price	No. Bidders	Domestic Winner
	(1)	(2)	(3)	(4)	(5)
log Standardized Quantity	-0.310^{***}	0.042^{***}	-0.027^{***}	0.008^{**}	0.003^{***}
	(0.016)	(0.003)	(0.003)	(0.003)	(0.001)
Bureaucrat FE * Preferenced * Policy Active	-0.120^{***}	0.070^{**}	-0.351^{***}	1.004^{***}	0.162^{**}
	(0.025)	(0.034)	(0.102)	(0.225)	(0.066)
Organization FE * Preferenced * Policy Active	-0.083^{***}	0.014	0.029	0.476^{*}	-0.119^{**}
	(0.030)	(0.028)	(0.116)	(0.249)	(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

*** $p<0.01$, ** $p<0.05$, * $p<0.1$ This table estimates the triple-difference from equation (14): $p_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta \text{Preferenced}_{gt} \times \text{PolicyActive}_t + \gamma \text{Preferenced}_{gt} \times \hat{\alpha}_b + \zeta \text{Preferenced}_{gt} \times \hat{\psi}_j + \eta \text{PolicyActive}_t \times \hat{\alpha}_b + \theta \text{PolicyActive}_t \times \hat{\psi}_j + \pi \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \nu \text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt}$. 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 Preferenced = 1 if the type of good appears on the list of goods covered by the preferences policy for that year. In the Pharmaceuticals sample, Preferenced = 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 covariance-shrunk, estimated bureaucrat and organization effects 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.

A Appendix

A.1 Robustness of Event Study Design

In Appendix Figure A.1, 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 2, 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 2.

A.2 Probing the log-linearity assumption

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 A.2 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 (9) 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 (9) 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.

A.3 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.* (2014b) 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.[57] 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.

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

Figure A.1: Robustness of Event Studies to Design Choices

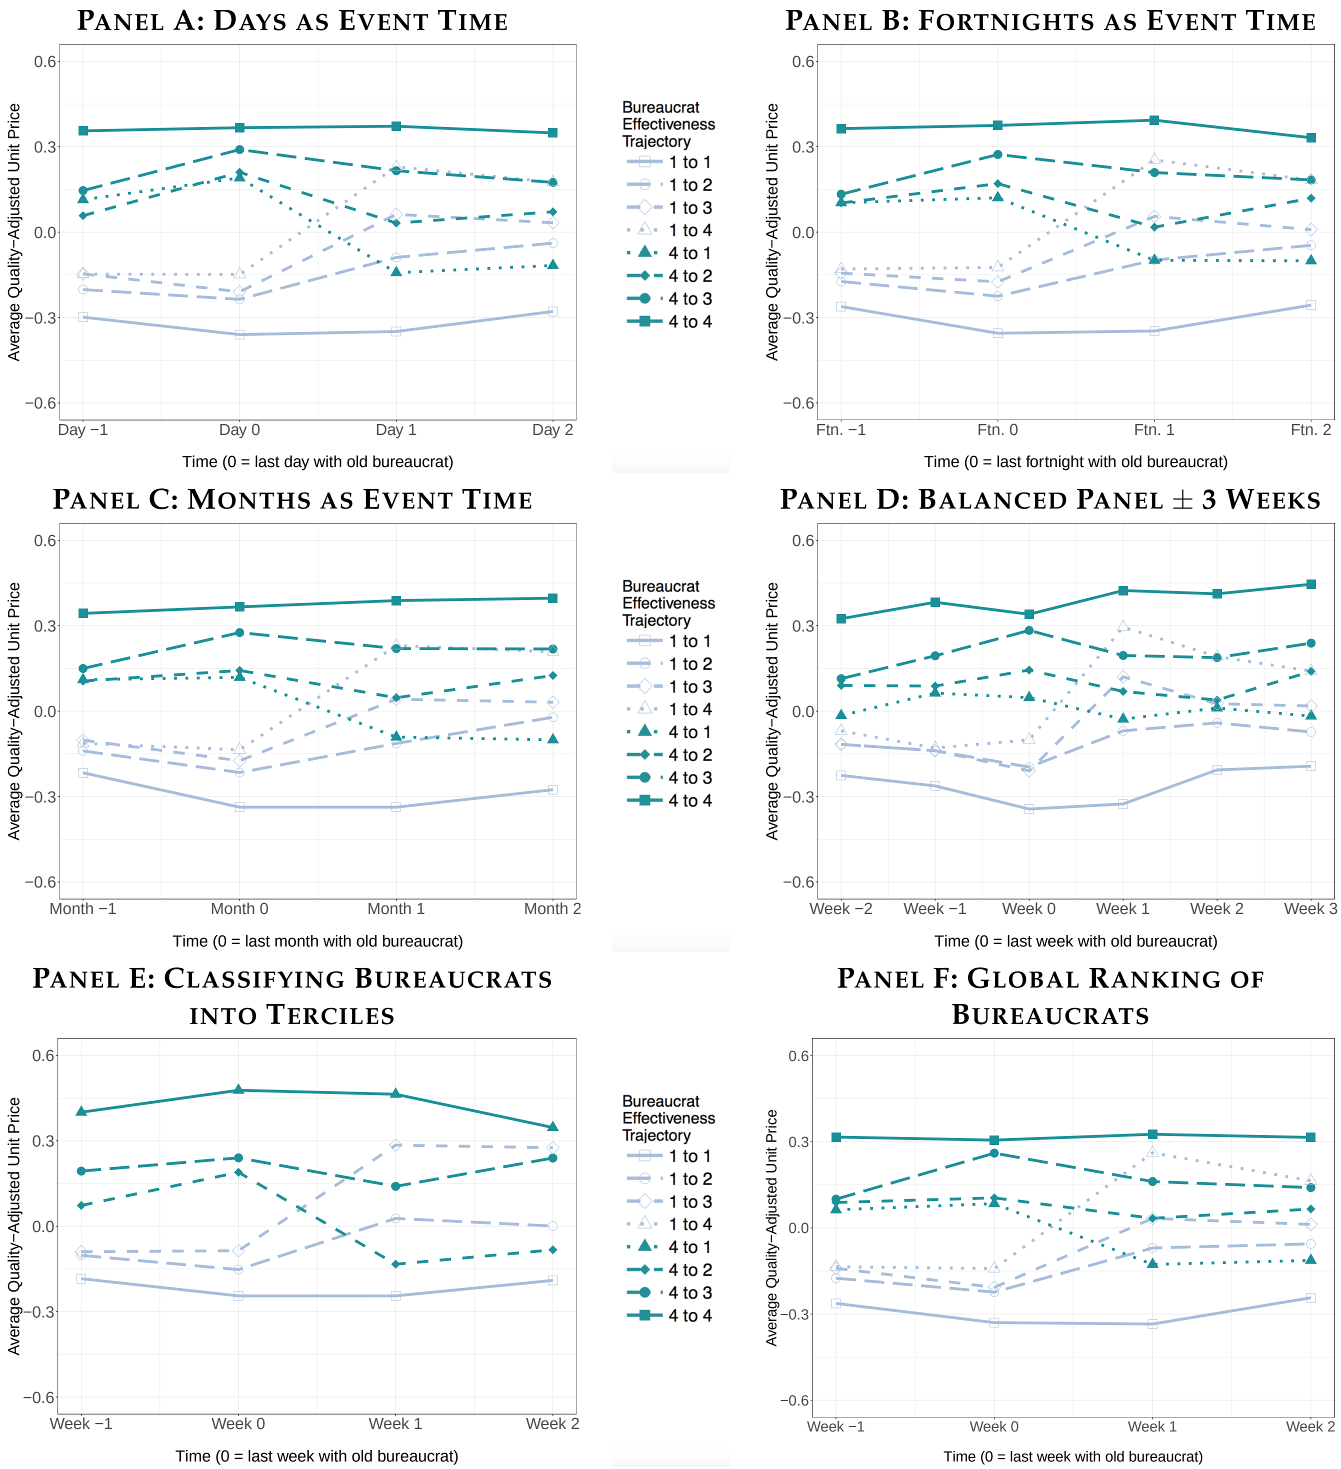

Each panel in the figure is analogous to Figure 2 (see notes to that figure for details of construction), with the following changes. In Panel A, rather than requiring the bureaucrat-organization pair to work together in two separate weeks, we require the pair to work together on two separate days. In 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 A.2: Correlation of Residuals with Estimated Bureaucrat and Organization Effects

Panel A: Prices in Logs (Main Specification)

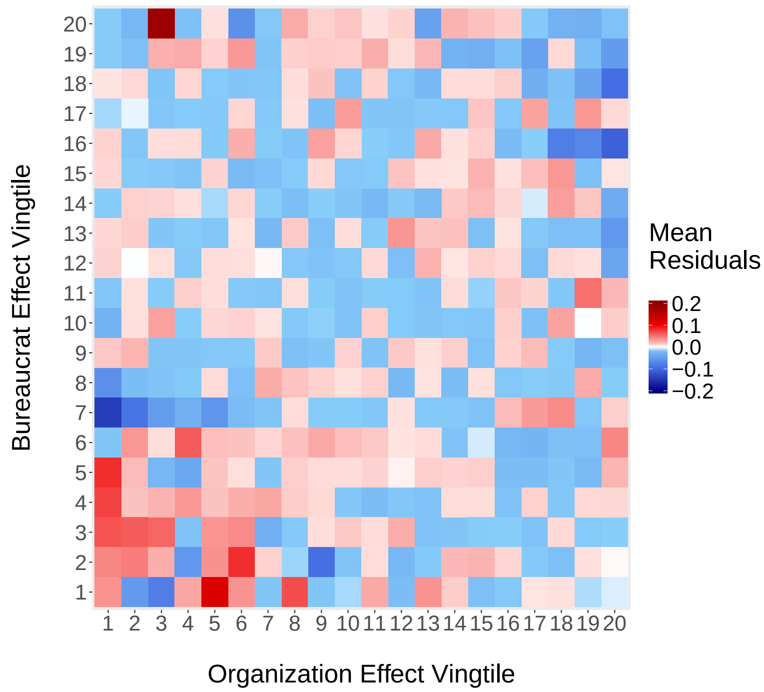

Panel B: Prices in Levels (Illustrating Misspecification)

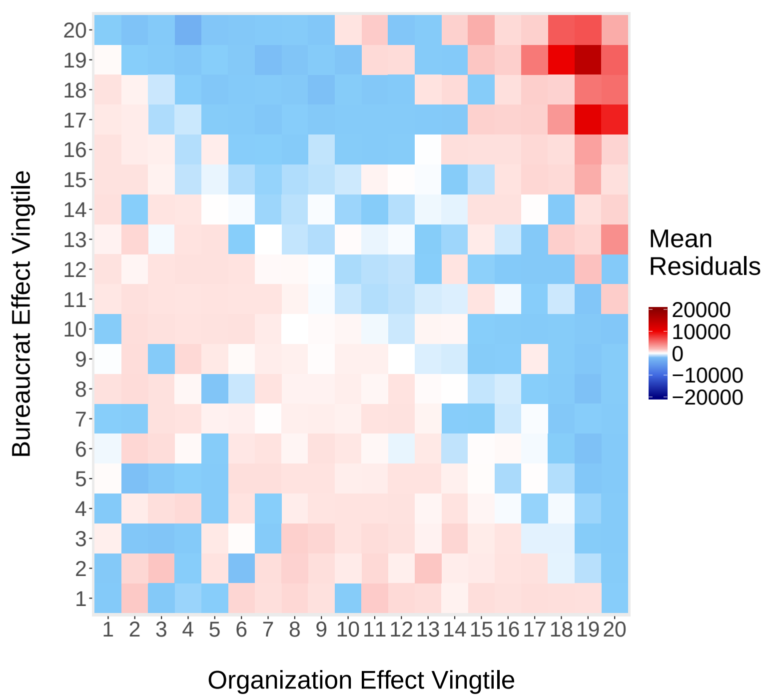

The figure presents heatmaps of averages of the residuals from the estimation of equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ — in logs (Panel A) and in levels (Panel B). The residuals are binned by vingtiles of the estimated bureaucrat effect $\hat{\alpha}_b$ and organization effect $\hat{\psi}_j$ within each connected set. The sample used is the Analysis Sample (All Products) summarized in Table 1.

OA Online Appendix (For Web Publication Only)

OA.1 Proofs of Propositions

OA.1.1 Proof of Proposition 1

Proof. We begin by demonstrating that the entry probabilities are as given by (2). For the firms to be indifferent between entering and not entering, equation (1) must hold. Solving the equation requires us to derive expressions for $\mathbb{E}[b_i]$ and $\mathbb{E}[\pi_i]$. The distribution of the bids is given by the bidding functions $b_i = M/m_i$ and the Pareto distributions of the markdowns m_i: $G_i(m_i) = 1 - m_i^{-\delta_i}$.

$$H_i(b) \equiv \mathbb{P}(b_i \leq b) = \mathbb{P}\left(m_i \geq \frac{M}{b}\right) = \left(\frac{b}{M}\right)^{\delta_i} \tag{OA.1}$$

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

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

$$\begin{aligned}
\mathbb{E}[\pi_i|b_i] &= \mathbb{E}_{b_j}[b_j - b_i|b_j > b_i]\,\mathbb{P}(b_j > b_i) = \int_{b_i}^{M} (b_j - b_i)\, dH_j(b_j) \\
&= \frac{\delta_j}{1+\delta_j} M - b_i + b_i \left(\frac{b_i}{M}\right)^{\delta_j} \frac{1}{1+\delta_j},
\end{aligned} \tag{OA.2}$$

where the final equality follows by inserting (OA.1) and integrating. Now we can derive unconditional expected profits by the law of iterated expectations:

$$\mathbb{E}[\pi_i] = \mathbb{E}_{b_i}[\mathbb{E}[\pi_i|b_i]] = \int_0^M \mathbb{E}[\pi_i|b_i]\, dH_i(b_i) = \left(\frac{1}{1+\delta_i} - \frac{1}{1+\delta_F+\delta_L}\right) M. \tag{OA.3}$$

Inserting these and the definition of the entry costs c_i into (1) and rearranging yields (2).

Turning to the expected prices, whenever neither or only one firm enters, the price is M. When both enter, the price is the higher of the two bids.

$$\mathbb{P}(p \leq x) = \mathbb{P}(\max\{b_F, b_L\} \leq x) = H_F(x)\, H_L(x) = \left(\frac{x}{M}\right)^{\delta_F+\delta_L} \tag{OA.4}$$

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

$$\begin{aligned}
\mathbb{P}(\log(p) \leq x) &= \mathbb{P}(p \leq \mathrm{e}^x) = \left(\frac{\mathrm{e}^x}{M}\right)^{\delta_F+\delta_L} \\
\mathbb{E}[\log(p) \mid \text{both enter}] &= \int_{-\infty}^{\log(M)} x \frac{\delta_F+\delta_L}{M^{\delta_F+\delta_L}} e^{(\delta_F+\delta_L)x} dx = \log(M) - \frac{1}{\delta_F+\delta_L}
\end{aligned} \tag{OA.5}$$

The expected log price is then simply $\mathbb{E}[\log(p)] = q_F q_L \mathbb{E}[\log(p) \mid \text{both enter}] + (1 - q_F q_L)\log(M)$. Inserting (OA.5) and the entry probabilities q_F and q_L yields expression (3) in the proposition.

The comparative statics on prices follow straightforwardly from equation (3). 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.1.2 Proof of Proposition 2

Proof. Consider the three cases in proposition 2 in turn.

Buyers with $\alpha_c + \psi_c \leq \underline{\mathbf{c}}$. In this case, both bidders enter the auction with certainty. Entering the auction is a best response to the other bidder entering whenever $\mathbb{E}\left[\pi_i|\gamma\right] - c_i > 0$. 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 (4) over all bids,

$$\mathbb{E}\left[\pi_F|\gamma < 1\right] = \int_0^M \mathbb{E}\left[\pi_F|b_F, \gamma < 1\right] dH_F\left(b_F|\gamma < 1\right) = \gamma^{1+\delta_F} M \left(\frac{1}{1+\delta_F} - \frac{1}{1+\delta_F+\delta_L}\right) \tag{OA.6}$$

Setting (OA.6) 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:

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

Combining these the distribution of prices is given by

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

In turn, the distribution of log prices is given by

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

making the expected log price in the auction

$$\mathbb{E}\left[\log\left(p\right)|\text{both enter}\right] = \int_{-\infty}^{\log(\gamma M)} \tfrac{\delta_L \gamma^{-\delta_L}+\delta_F\gamma^{\delta_F}}{M^{\delta_F+\delta_L}} x \mathrm{e}^{(\delta_F+\delta_L)x}\mathrm{d}x + \int_{\log(\gamma M)}^{\log(M)} \tfrac{\delta_F\gamma^{\delta_F}}{M^{\delta_F+\delta_L}} x \mathrm{e}^{(\delta_F+\delta_L)x}\mathrm{d}x$$
$$+\left[1-H_F\left(M\right)\right]\log\left(M\right)$$
$$= \log\left(M\right) - \frac{\gamma^{\delta_F}\left(1-\log\left(\gamma^{\delta_L}\right)\right)}{\delta_F+\delta_L}. \tag{OA.7}$$

Comparing (OA.7) to the expected price without preferences (OA.5), prices rise as long as assumption 1 holds.

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

$$\mathbb{P}\left(L \text{ wins}\right) = \mathbb{P}\left(b_L < b_F\right) = \int_0^M H_L\left(b_F|\gamma=1\right)dH_F\left(b_F|\gamma=1\right) = 1-\frac{\delta_L}{\delta_F+\delta_L}, \tag{OA.8}$$

while when there are preferences this increases to

$$\mathbb{P}\left(L \text{ wins}\right) = \mathbb{P}\left(b_L < b_F|\gamma<1\right) = \int_0^M H_L\left(b_F|\gamma<1\right)dH_F\left(b_F|\gamma<1\right) = 1-\gamma^{\delta_F}\frac{\delta_L}{\delta_F+\delta_L}. \tag{OA.9}$$

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 $\mathbb{E}\left[\pi_F|\gamma\right] - c_F < 0$ and $\mathbb{E}\left[\pi_L|\gamma\right] - c_L > 0$. In this case, since only L enters, the price is M 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 $\mathbb{E}\left[\pi_L|\gamma\right] - c_L = 0$. To find the upper threshold $\bar{c}$, we require an expression for $\mathbb{E}\left[\pi_L|\gamma\right]$:

$$\mathbb{E}\left[\pi_L|\gamma<1\right] = \int_0^M \mathbb{E}\left[\pi_L|b_L,\gamma<1\right]dH_L\left(b_L|\gamma<1\right) = M\left(\frac{1}{1+\delta_L} - \frac{\gamma^{\delta_F}}{1+\delta_F+\delta_L}\right). \tag{OA.10}$$

Setting (OA.10) 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: $\mathbb{E}\left[\pi_i|\gamma\right] - c_i < 0\ \forall i$ and so the equilibrium is in mixed strategies. As in proposition 1, the entry probabilities are given by

$$q_i = \frac{M - \mathbb{E}\left[M/m_j\right] - c_j}{M - \mathbb{E}\left[b_j\right] - \mathbb{E}\left[\pi_j|\gamma<1\right]}.$$

In this case the expected price is given by

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

Inserting the entry probabilities and the price equation (OA.7) and rearranging, the expected price when there are preferences is lower whenever

$$\begin{aligned} & q_F(\gamma<1)\, q_L(\gamma<1)\,(\log(M)-\mathbb{E}[\log(p)\,|\text{both enter}, \gamma<1]) \\ & -q_F(\gamma=1) q_L(\gamma=1)\,(\log(M)-\mathbb{E}[\log(p)\,|\text{both enter}, \gamma=1]) \geq 0 \\ & \iff -\log\left(\gamma^{\delta_L}\right)-\frac{\delta_L}{1+\delta_F}\left(1-\gamma^{1+\delta_F}\right) \geq 0 \end{aligned} \tag{OA.11}$$

Noting that (OA.11) holds with equality when $\gamma=1$ and that the left hand side of (OA.11) has slope $-\delta_L\left(\gamma^{-1}-\gamma^{\delta_F}\right)<0\ \forall \gamma<1$ shows that (OA.11) holds for all $\gamma<1$. Participation in the auction is $\mathbb{E}[N]=q_F+q_L$. When there are no preferences

$$\mathbb{E}[N|\gamma=1]=q_F(\gamma=1)+q_L(\gamma=1)=2\frac{1+\delta_F+\delta_L}{1+\delta_L}\sqrt{1-\alpha_c-\psi_c}, \tag{OA.12}$$

while with preferences participation is

$$\begin{aligned} \mathbb{E}[N|\gamma<1] &= q_F(\gamma<1)+q_L(\gamma<1) \\ &= \left(\frac{1}{\gamma^{\delta_F}}+\frac{1}{\gamma^{1+\delta_F}+\left(1-\gamma^{1+\delta_F}\right)\frac{1+\delta_F+\delta_L}{1+\delta_F}}\right)\frac{1+\delta_F+\delta_L}{1+\delta_L}\sqrt{1-\alpha_c-\psi_c}. \end{aligned} \tag{OA.13}$$

Comparing (OA.12) to (OA.13) shows that participation increases whenever

$$\frac{1}{\gamma^{\delta_F}}+\frac{1}{1+\frac{\delta_L}{\delta_F+\delta_L}\left(1-\gamma^{1+\delta_F}\right)}>2 \tag{OA.14}$$

Equation (OA.14) is implied by assumption 2. 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

$$\begin{aligned} \frac{q_F(\gamma=1)}{q_F(\gamma<1)} &= \frac{M-\mathbb{E}[M/m_L]-\mathbb{E}[\pi_L|\gamma<1]}{M-\mathbb{E}[M/m_L]-\mathbb{E}[\pi_L|\gamma=1]}=\gamma^{\delta_F} \\ \frac{q_L(\gamma=1)}{q_L(\gamma<1)} &= \frac{M-\mathbb{E}[M/m_F]-\mathbb{E}[\pi_F|\gamma<1]}{M-\mathbb{E}[M/m_F]-\mathbb{E}[\pi_F|\gamma=1]}=1+\frac{\delta_L}{1+\delta_F}\left(1-\gamma^{1+\delta_F}\right) \end{aligned}$$

Combining these two components shows that the statement is correct as long as assumption 2 holds. □

OA.2 Identification of Bureaucrat and Organization Effects with Multiple Connected Sets

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

To see this explicitly, write the model as

$$\mathbf{p} = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\alpha} + \mathbf{F}\boldsymbol{\psi} \tag{OA.15}$$

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

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

$$\begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\alpha}}_{OLS} \\ \hat{\boldsymbol{\psi}}_{OLS} \end{bmatrix} = \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \mathbf{p} \tag{OA.16}$$

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

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

Specifically, we augment the model to be

$$\mathbf{p} = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\tilde{\boldsymbol{\alpha}} + \mathbf{F}\tilde{\boldsymbol{\psi}} + \mathbf{S}\boldsymbol{\gamma} \tag{OA.17}$$

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

$$\begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} & \mathbf{S} \end{bmatrix} \\ \begin{bmatrix} \mathbf{S_b} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S_o} & \mathbf{0} \end{bmatrix} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\alpha}} \\ \hat{\boldsymbol{\psi}} \\ \hat{\boldsymbol{\gamma}} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \mathbf{p} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \tag{OA.18}$$

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.15), then $\hat{\boldsymbol{\alpha}}$, $\hat{\psi}$, and $\hat{\gamma}$, the estimators of $\tilde{\boldsymbol{\alpha}}$, $\tilde{\psi}$ and γ in the augmented model (OA.17) that solve the augmented normal equations (OA.18) (i) are uniquely identified, and (ii) are related to the true bureaucrat and organization effects $\boldsymbol{\alpha}$ and ψ by*

$$\begin{bmatrix} \hat{\boldsymbol{\alpha}} \\ \hat{\psi} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\alpha} - \mathbf{S_b}'\overline{\boldsymbol{\alpha}} \\ \psi - \mathbf{S_o}'\overline{\psi} \\ \overline{\boldsymbol{\alpha}} + \overline{\psi} \end{bmatrix} \tag{OA.19}$$

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

Proof. We will prove each part of the result separately. To see uniqueness, first note that the standard normal equations for (OA.17) only has rank $N_b + N_o - N_s$. To see this, we note that $\mathbf{BS_b}' = \mathbf{FS_o}' = \mathbf{S}$ and so $2N_s$ columns of the $N \times (N_b + N_o + N_s)$ matrix $[\mathbf{B}\ \mathbf{F}\ \mathbf{S}]$ are collinear. However, the $2N_s$ restrictions $\mathbf{S_b}\hat{\alpha} = 0$ and $\mathbf{S_o}\hat{\psi} = 0$ are independent of the standard normal equations, so the first matrix in (OA.18) has rank $N_b + N_o + N_s$ and hence the solution to (OA.18) is unique.

To see the second part, it suffices to show that (OA.19) solves (OA.18). First, substitute the estimators out of (OA.18) using (OA.19) and substitute in the true model using (OA.15) to rewrite (OA.18) as

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

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

OA.3 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.3.1 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 Toolkit[58]

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

We refer to the result as the *tokenized* sentence. For example the product description "NV-Print Cartridge for the Canon LBP 2010B Printer" would be broken into the following tokens: [cartridge, NV-Print, printer, Canon, LBP, 3010B]. [59] Similarly, the product description "sodium bicarbonate - solution for infusion 5%,200ml" would result in the following tokens: [sodium, bicarbonate, solution, infusion, 5%, 200ml].[60]

OA.3.2 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

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

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

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

of token indicators and indicators for each possible bi-gram (word-pair), denoted by $\mathbf{x}_i \in \mathcal{X}_C$.[61] Each sentence also has a corresponding good classification $g_i \in \mathcal{G}_C$, so we can represent our customs data as the pair $\{\mathbf{X}_C, \mathbf{g}_C\}$ and we seek to find a classifier $\hat{g}_C(\mathbf{x}) : \mathcal{X}_C \to \mathcal{H}_C$ that assigns every text vector **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:

$$\hat{g}_C(\mathbf{x}_i) = \arg\max_{g \in \mathcal{G}_C} p_g(\mathbf{x}_i) \tag{OA.20}$$

Each binary classifier is a logistic regression solving

$$\min_{\mathbf{w}_g, a_g} \frac{1}{N_C} \sum_{i=1}^{N_C} \frac{1}{\ln 2} \ln\left(1 + \mathrm{e}^{-y_{gi} \cdot (\mathbf{w}_g \cdot \mathbf{x}_i + a_g)}\right) \tag{OA.21}$$

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.20). We implement this procedure using the Vowpal Wabbit library for Python.[62] This simple procedure is remarkably effective; when trained on a randomly selected half of the customs data and then implemented on the reamining data for validation, the classifications are correct 95% of the time. Given this high success rate without regularization, 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.[63] That is, a classifier that solves

$$\min_{\mathbf{w}_g, a_g} \frac{1}{N_C} \sum_{i=1}^{N_C} \max\{0, 1 - y_{gi} \cdot (\mathbf{w}_g \cdot \mathbf{x}_i + a_g)\} \tag{OA.22}$$

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

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

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

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

TABLE 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

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 2 and table 3 is repeated using these alternative classifiers in figure OA.1 panels C and D and in table OA.11. 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 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.

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, возрасть, количество
16067065	44FZ	Скоросшиватель, пластиковый	3926100000	принадлежность, канцелярский, закладка, пластиковый, книга
18267299	44FZ	Скоросшиватель, пластиковый, Brauberg	3926100000	набор, канцелярский, настольный, индивидуальный, пластмассовый, упаковка, розничный, продажа

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.3.3 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.20). Second, the similarity between the sentence and the average sentence with the sentence's assigned classification in the *customs* data used to train the classifier.

To identify outlier sentences, we take the tokenized sentences that have been labeled as good g, $\mathbf{t}_g = \{\mathbf{t}_i : \hat{g}_C(\mathbf{x}_i) = g\}$ and transform them into vectors of indicators for the tokens $\mathbf{v}_{gi}$.[64] For each good, we then calculate the mean sentence vector in the customs data as $\mathbf{v}_g^C = \sum_{\mathbf{v}_{gi}, x_i \in \mathbf{X}^C} \mathbf{v}_{gi} / |\mathbf{t}_g|$. Then, to identify outlier sentences in the procurement data, we calculate each sentence's normalized cosine similarity with the good's mean vector,

$$\theta_{gi} = \frac{\bar{s}_g - s(\mathbf{v}_{gi}, \mathbf{v}_g)}{\bar{s}_g} \tag{OA.23}$$

where $s(\mathbf{v}_{gi}, \mathbf{v}_g) \equiv \cos(\mathbf{v}_{gi}, \mathbf{v}_g) = \frac{\mathbf{v}_{gi}\mathbf{v}_g}{\|\mathbf{v}_{gi}\|\|\mathbf{v}_g\|} = \frac{\sum_{k=1}^{K_g} t_{gik} t_{gk}}{\sqrt{\sum_{k=1}^{K_g} t_{gik}^2}\sqrt{\sum_{k=1}^{K_g} t_{gk}^2}}$ is the cosine similarity of the sentence vector $\mathbf{v}_{gi}$ with its good mean $\mathbf{v}_g$,[65] 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

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

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

similarity θ_{gi} was above the median. Figure OA.1 panels A and B and Table OA.12 show the robustness of our results to using the 45th or 55th percentile as thresholds.

OA.3.3 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

$$\min_{\mathbf{c}} \sum_{i=1}^{N} \left\| f\left(\mathbf{c}, \mathbf{t}_i\right) - \mathbf{t}_i \right\|^2 \tag{OA.24}$$

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

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

$$\eta\left(i\right) = \frac{b\left(i\right) - a\left(i\right)}{\max\left\{b\left(i\right), a\left(i\right)\right\}} \tag{OA.25}$$

where $a\left(i\right)$ is the average distance between sentence i and the other sentences in the same cluster, and $b\left(i\right)$ 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.4 Additional Figures and Tables

FIGURE OA.1: ROBUSTNESS OF EVENT STUDIES TO ALTERNATIVE TEXT CLASSIFIERS, CLASSIFICATION ACCURACY THRESHOLDS, AND OUTLIER TRIMMING

PANEL A: CLASSIFIER ACCURACY THRESHOLD 45TH PERCENTILE

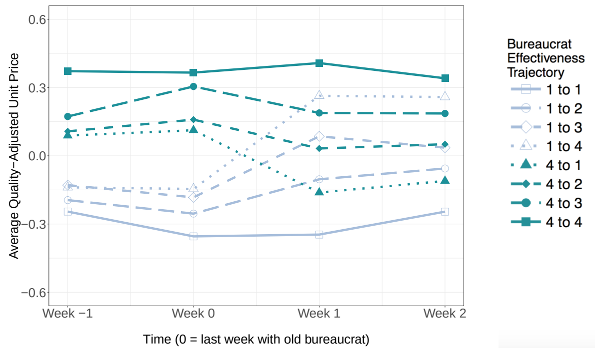

PANEL B: CLASSIFIER ACCURACY THRESHOLD 55TH PERCENTILE

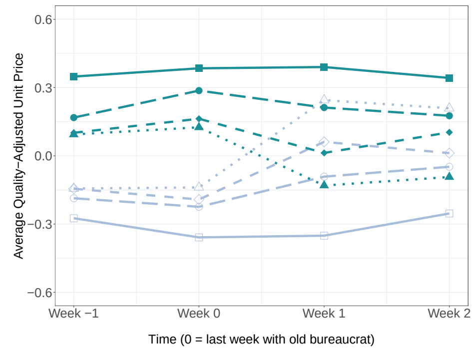

PANEL C: SUPPORT VECTOR MACHINE CLASSIFIER

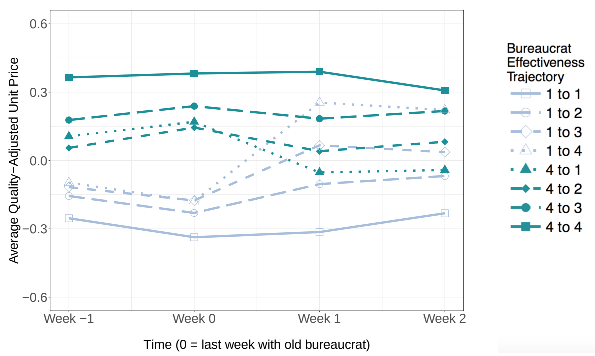

PANEL D: HIERARCHICAL MODEL CLASSIFIER

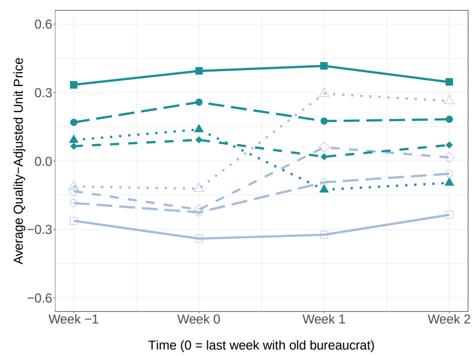

PANEL E: DROPPING TOP AND BOTTOM 2.5% OUTLIERS

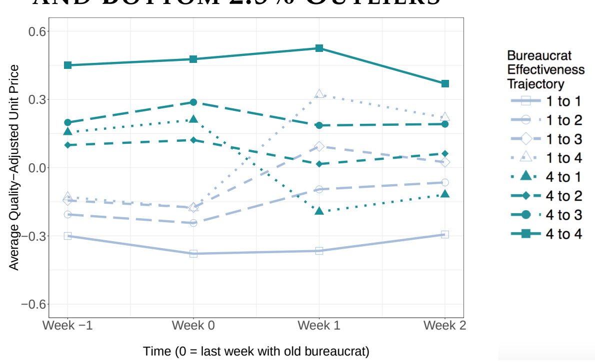

Each panel in the figure is analogous to Figure 2 (see notes to that figure for details of construction), with the following changes. In Panel A, rather than requiring the bureaucrat-organization pair to work together in two separate weeks, we require the pair to work together on two separate days. In 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.2: EVENT STUDY OF PROCUREMENT PRICES AROUND TIMES BUREAUCRATS AND ORGANIZATIONS SWITCH GOODS

PANEL A: BUREAUCRATS SWITCHING GOODS **PANEL B: ORGANIZATIONS SWITCHING GOODS**

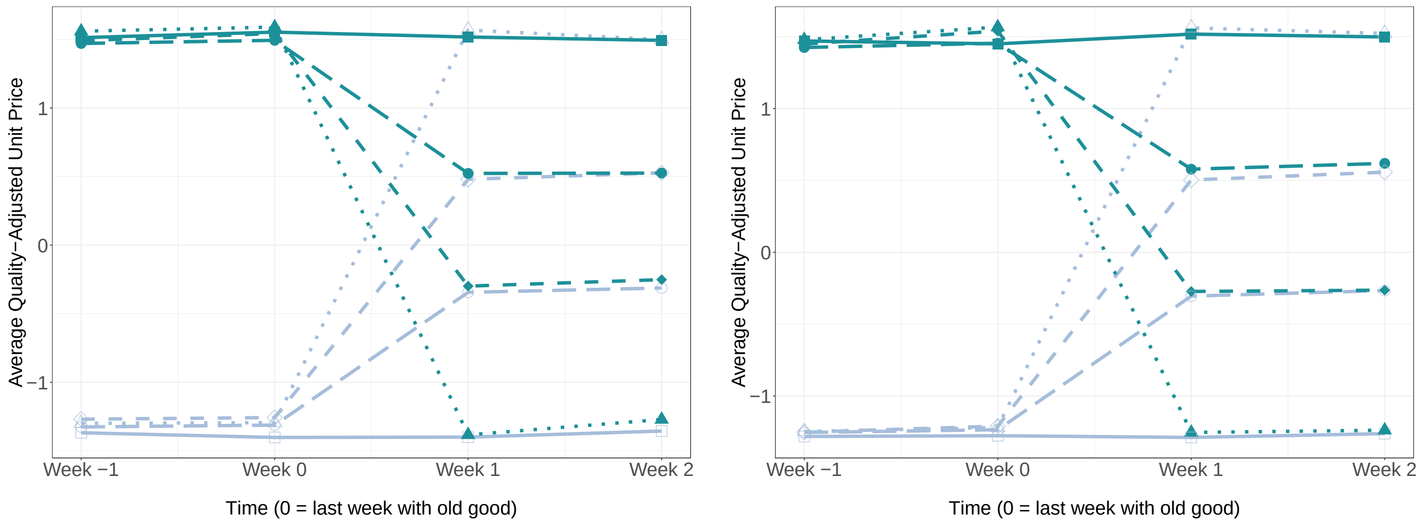

Each panel in the figure is analogous to Figure 2 that studies price changes around the time that organizations switch the bureaucrat making their purchases (see notes to that figure for details of construction). Panel A shows price changes around the time that bureaucrats switch the good they are purchasing. Panel B shows price changes around the time that organizations switch the good they are purchasing.

FIGURE OA.3: CORRELATION MATRIX FOR VARIABLES PREDICTING EFFECTIVE BUREAUCRATS AND ORGANIZATIONS

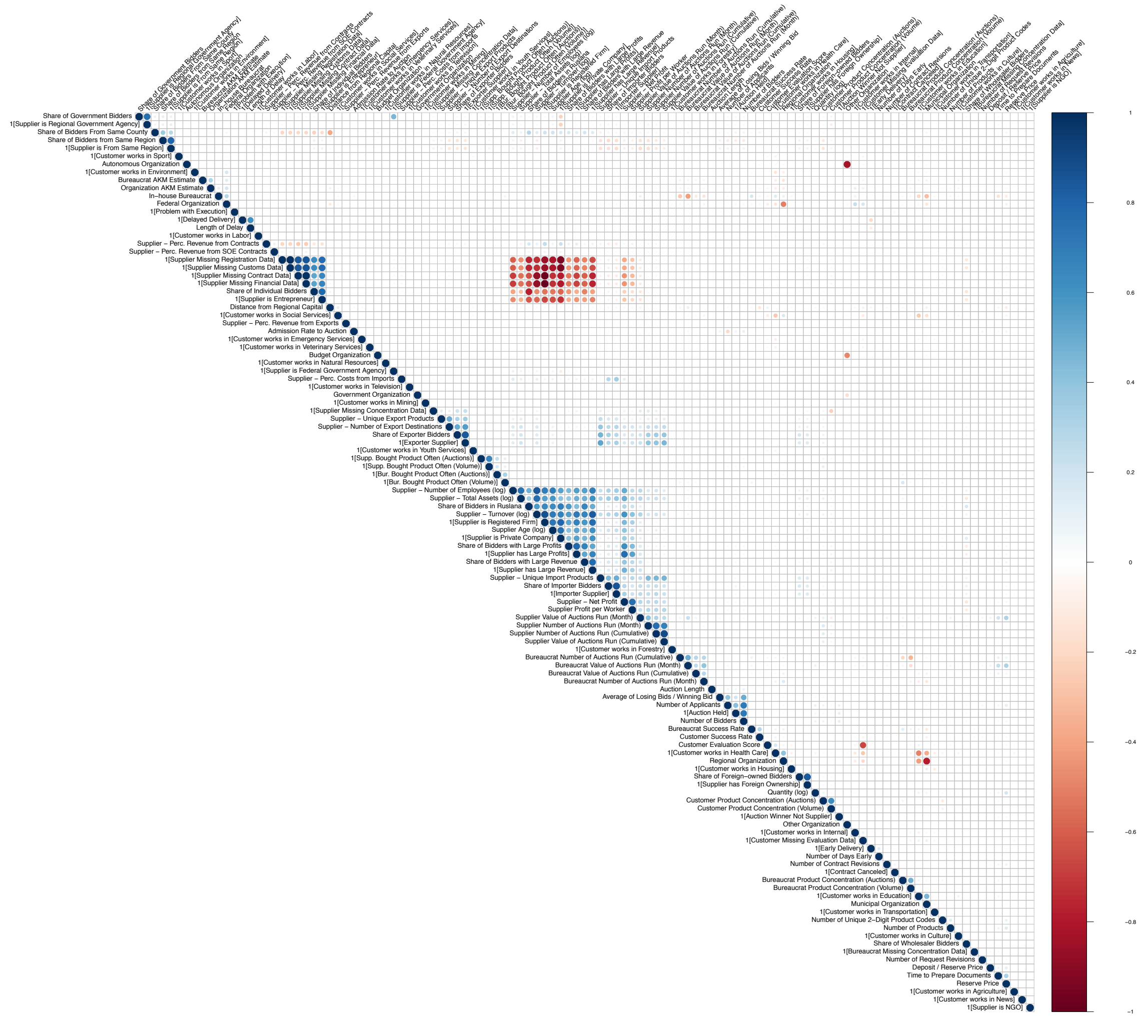

This figure shows correlations between each of the 113 potential explanatory variables for bureaucrat and organization effectiveness. Squares in blue denote positive correlations while negative correlations are shown in red. Larger circles and darker colors denote larger correlation coefficients.

Figure OA.4: Correlates of Bureaucrat Effectiveness: Lambda Returning Cross-Validation Error Within One Standard Error of the Minimum

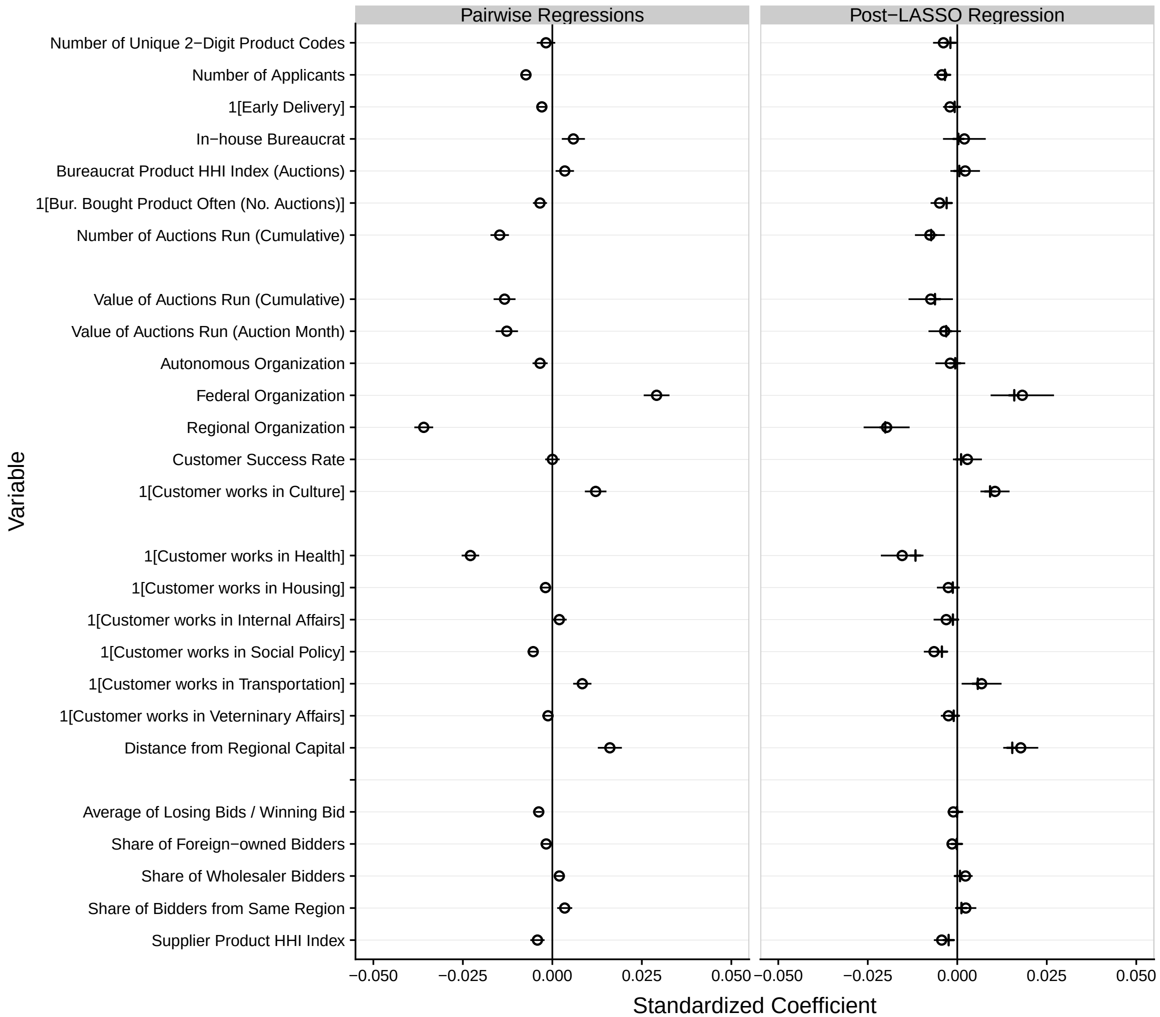

The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The circles in right column show the coefficients from a multivariate regression of the estimated bureaucrat effects on all the correlates that are selected by an elastic net regularization procedure, where the values of the regularization penalty lambda λ are chosen to return a cross-validation error within one standard deviation of the minimum. The elastic net mixing parameter used is 0.3. The coefficients from the elastic net regularization procedure are shown as crosses.

Figure OA.5: Correlates of Organization Effectiveness: Lambda Returning Cross-Validation Error Within One Standard Error of the Minimum

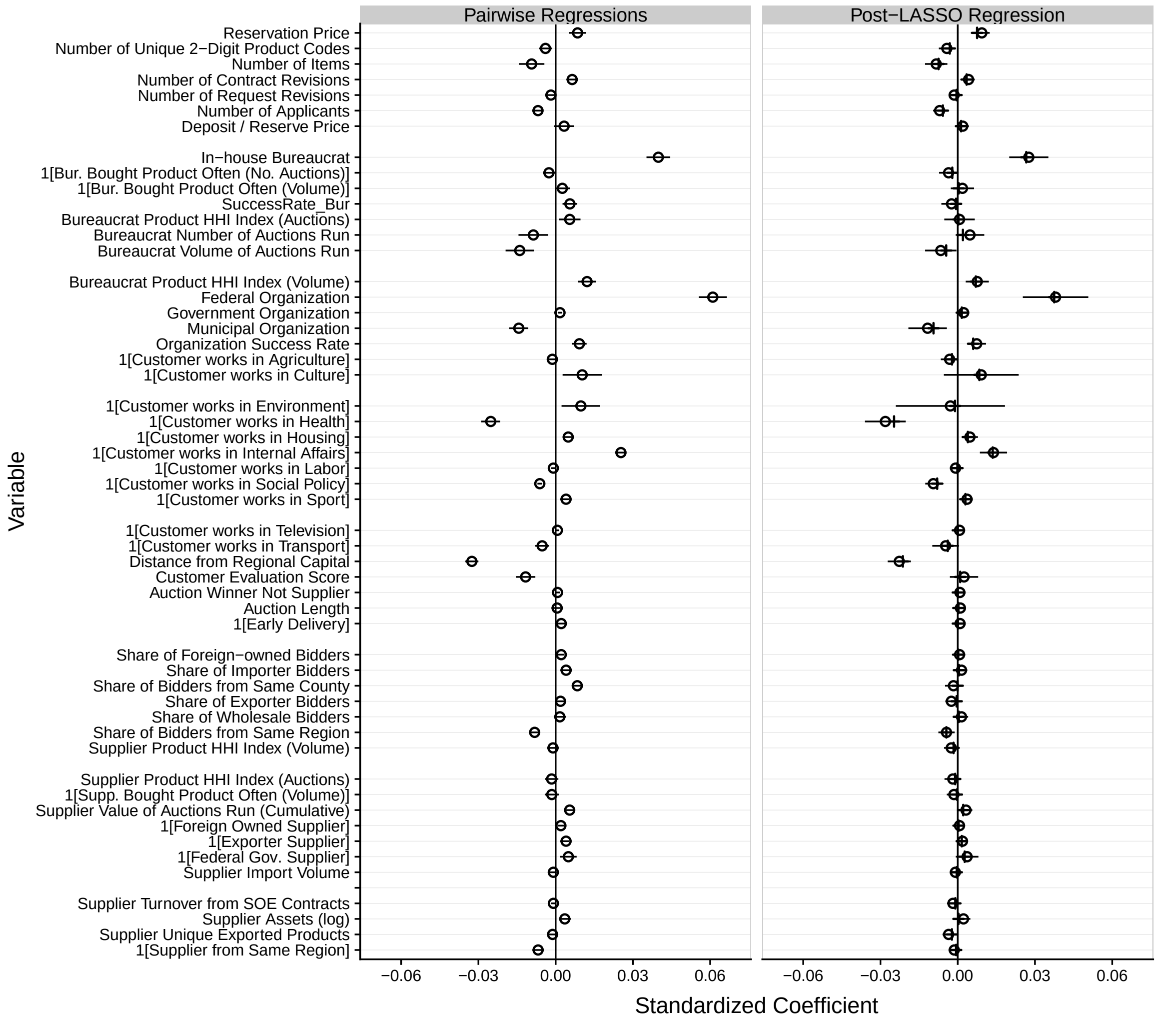

The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The circles in right column show the coefficients from a multivariate regression of the estimated bureaucrat effects on all the correlates that are selected by an elastic net regularization procedure, where the values of the regularization penalty lambda λ are chosen to return a cross-validation error within one standard deviation of the minimum. The elastic net mixing parameter used is 0.3. The coefficients from the elastic net regularization procedure are shown as crosses.

Figure OA.6: Correlates of Bureaucrat Effectiveness: Lambda Returning Cross-Validation Error Within One Standard Error of the Minimum Increasing Mixing Parameter to 0.7

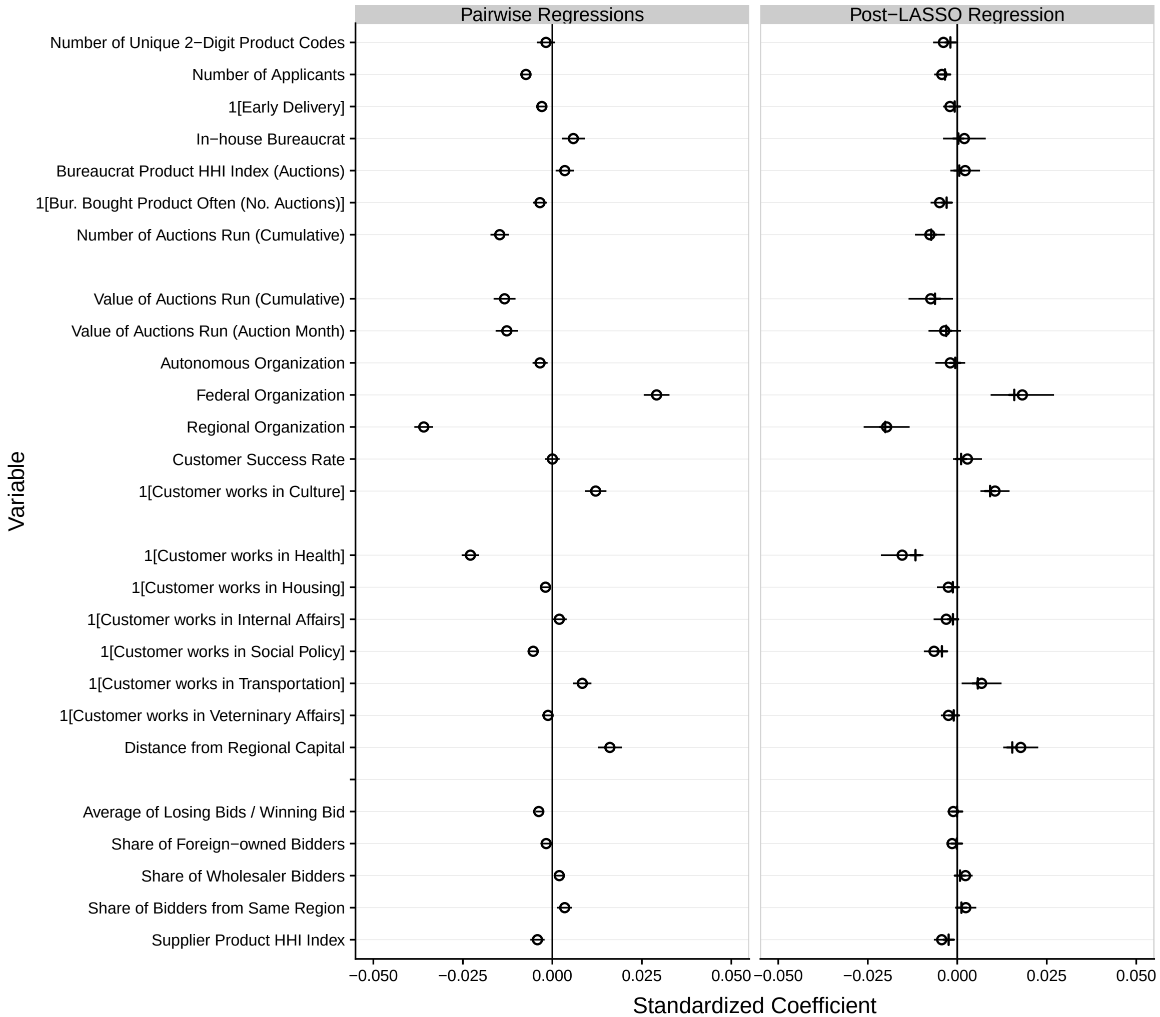

The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The circles in right column show the coefficients from a multivariate regression of the estimated bureaucrat effects on all the correlates that are selected by an elastic net regularization procedure, where the values of the regularization penalty lambda λ are chosen to return a cross-validation error within one standard deviation of the minimum. The elastic net mixing parameter used is 0.7. The coefficients from the elastic net regularization procedure are shown as crosses.

Figure OA.7: Correlates of Organization Effectiveness: Lambda Returning Cross-Validation Error Within One Standard Error of the Minimum Increasing Mixing Parameter to 0.7

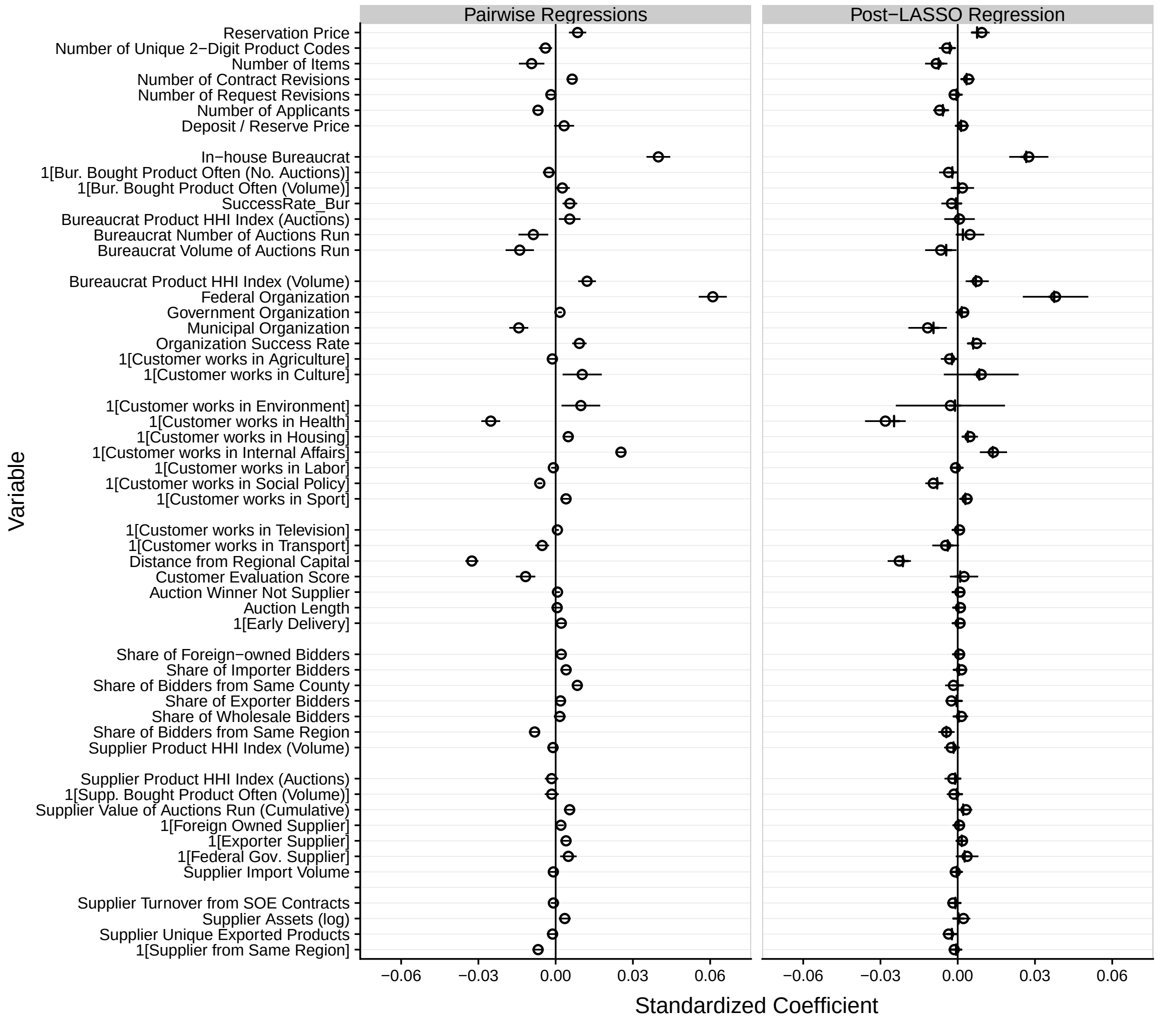

The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The circles in right column show the coefficients from a multivariate regression of the estimated bureaucrat effects on all the correlates that are selected by an elastic net regularization procedure, where the regularization parameter lambda is chosen to return a cross-validation error within one standard deviation of the minimum. The elastic net mixing parameter used is 0.7. The coefficients from the elastic net regularization procedure are shown as crosses.

Figure OA.8: Correlates of Bureaucrat Effectiveness: Elastic Net Regularization Coefficients Across Different Mixing Parameters

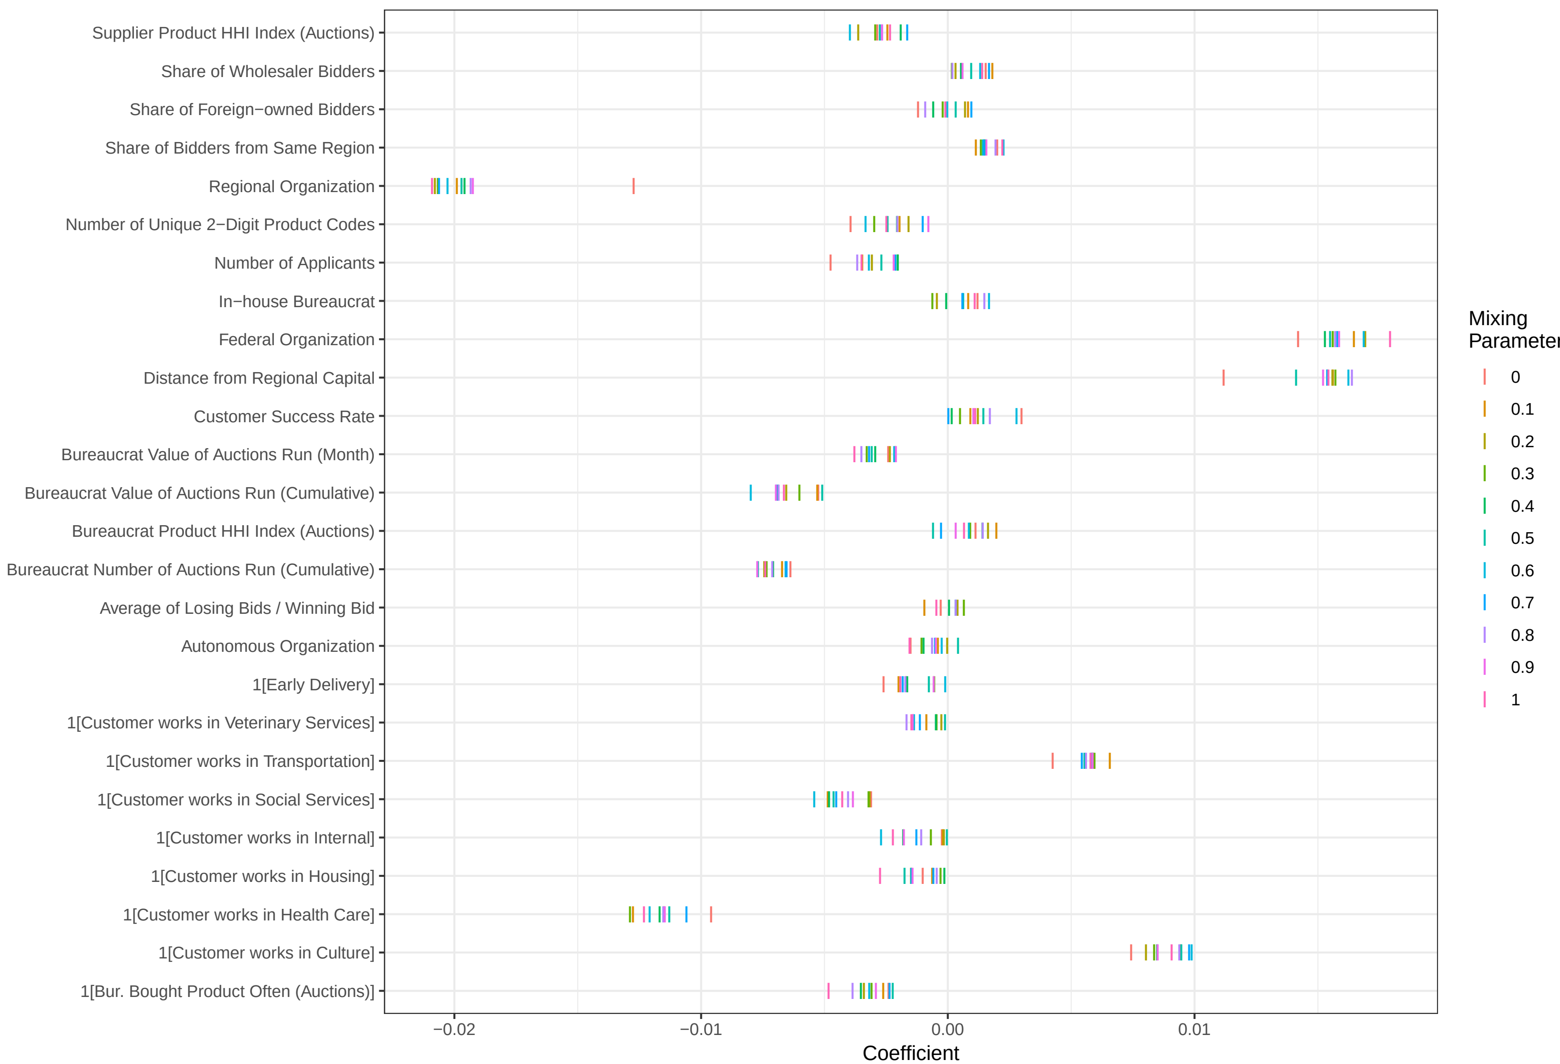

The figure shows the coefficients from the elastic net regularization procedure on the estimated bureaucrat effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. The variables shown are from the base model shown in Online Appendix Table OA.4 where the values of the regularization penalty lambda λ are chosen to return a cross-validation error within one standard deviation of the minimum.

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

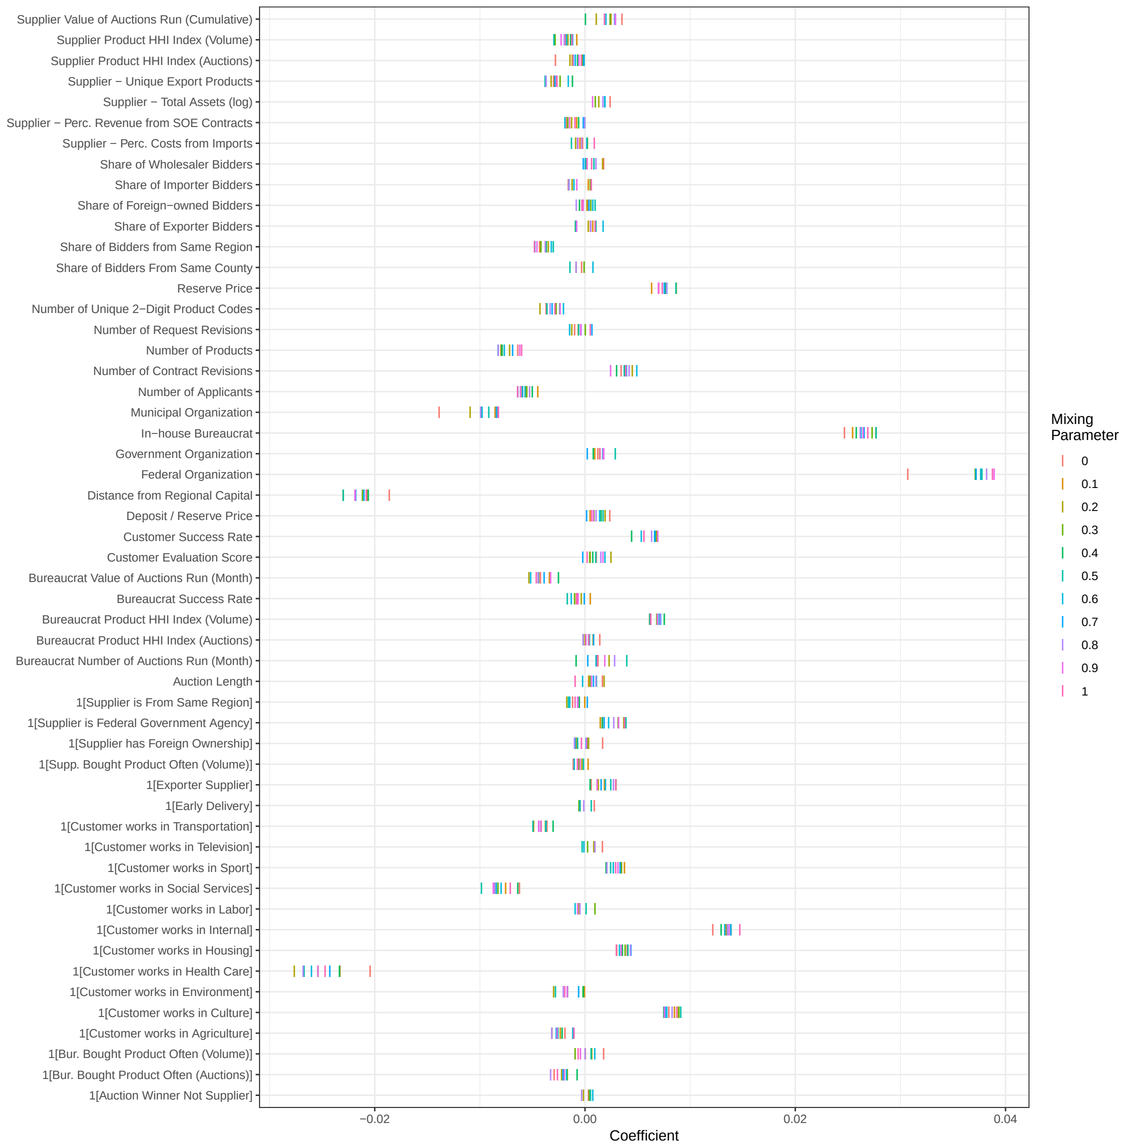

The figure shows the coefficients from the elastic net regularization procedure on the estimated bureaucrat effects across different values of the mixing parameters. Each coefficient is represented by a small vertical line corresponding by color to mixing parameters. The variables shown are from the base model shown in Online Appendix Table OA.4 where the values of the regularization penalty lambda λ are chosen to return a cross-validation error within one standard deviation of the minimum.

TABLE OA.3: 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 unpreferenced auctions in connected sets that fulfill three restrictions: singleton bureaucrat-organization, bureaucrat-good, and organization-good pairs are removed; each procurer (bureaucrats and organizations) implements a minimum of five purchases; and connected sets have at least three bureaucrats and organizations. Largest Connected Set is the largest connected set from the Analysis Sample (as measured by the number of organizations). Organizations working in Education include schools, universities, pre-schools, and youth organizations. Organizations working in Internal Affairs include police, emergency services, local administration, taxes, and transportation. Organizations working in Agriculture or the Environment include environmental protection funds, agricultural departments and nature promotion agencies. The Other category includes funds, monitoring agencies, and land cadasters, among many others. All sums are measured in billions of US dollars at an exchange rate of 43 rubles to 1 US dollar.

TABLE 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: Products Covered by Preference Laws, by Year

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

Table OA.6: 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 (10). 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 (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column 3 shows estimates from randomly splitting the sample in half, stratifying by bureaucrat-organization pair and calculating the covariance across the two noisy estimates 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 4.2.

TABLE OA.7: 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

The table implements the variance decomposition in equation (10) using the estimates from equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Column (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.8: ELASTIC NET NONZERO COEFFICIENTS ARE ROBUST ACROSS MIXING PARAMETERS (BUREAUCRAT EFFECTIVENESS)

Variable	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Admission Rate to Auction	0.000										
Bureaucrat Product HHI Index (Auctions)	0.001	0.001	0.001	0.001	0.001	0.000	0.001	0.000	0.001	0.000	0.001
Supplier Product HHI Index (Auctions)	-0.002	-0.003	-0.003	-0.002	-0.002	-0.002	-0.003	-0.002	-0.003	-0.002	-0.003
1[Auction Winner Not Supplier]	0.001						0.000				0.000
Auction Length	-0.001						0.000				0.000
1[Problem with Execution]	0.000										
1[Auction Held]	0.000										
1[Bur. Bought Product Often (Auctions)]	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.004	-0.003	-0.003	-0.002	-0.004
1[Supp. Bought Product Often (Auctions)]	0.001										
1[Bur. Bought Product Often (Volume)]	-0.001										
1[Supp. Bought Product Often (Volume)]	-0.001										
Bureaucrat Number of Auctions Run (Cumulative)	-0.006	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007	-0.007
Bureaucrat Value of Auctions Run (Cumulative)	-0.006	-0.006	-0.006	-0.006	-0.006	-0.006	-0.007	-0.006	-0.006	-0.006	-0.007
Bureaucrat Number of Auctions Run (Month)	-0.002										
Bureaucrat Value of Auctions Run (Month)	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003
1[Customer works in Agriculture]	0.000										
1[Customer works in Culture]	0.008	0.009	0.009	0.009	0.009	0.009	0.010	0.009	0.009	0.009	0.010
Distance from Regional Capital	0.012	0.015	0.015	0.015	0.015	0.015	0.016	0.015	0.016	0.015	0.016
1[Customer works in Education]	0.002										
1[Customer works in Emergency Services]	-0.001						0.000		0.000		0.000
1[Customer works in Environment]	0.000						0.000				0.000
1[Customer works in Forestry]	0.000										
1[Customer works in Health Care]	-0.010	-0.012	-0.012	-0.012	-0.012	-0.011	-0.013	-0.011	-0.012	-0.011	-0.013
1[Customer works in Housing]	-0.002	-0.001	-0.001	-0.001	-0.001	-0.001	-0.002	-0.001	-0.001	-0.001	-0.002
1[Customer works in Internal]	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.002	-0.001	-0.002	-0.001	-0.002
1[Customer works in Labor]	0.001										
1[Customer works in Mining]	0.000										
1[Customer works in Natural Resources]	0.000										
1[Customer works in News]	0.000										
Customer Evaluation Score	0.002						0.000				0.000
1[Customer works in Social Services]	-0.003	-0.004	-0.004	-0.004	-0.004	-0.004	-0.005	-0.004	-0.005	-0.004	-0.005
1[Customer works in Sport]	0.000										
1[Customer works in Television]	0.000										
1[Customer works in Transportation]	0.005	0.006	0.006	0.006	0.006	0.006	0.006	0.006	0.006	0.005	0.006
1[Customer works in Veterinary Services]	-0.002	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.002
1[Customer works in Youth Services]	0.000										
1[Delayed Delivery]	0.000										
1[Early Delivery]	-0.002	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	0.000	-0.001
Deposit / Reserve Price	-0.001										
Length of Delay	0.000										
Number of Days Early	0.000										
In-house Bureaucrat	0.002	0.000	0.000	0.000	0.000	0.000	0.001	0.000	0.001		0.001
Average of Losing Bids / Winning Bid	-0.001	0.000	0.000	0.000	0.000		0.000		0.000		0.000
1[Bureaucrat Missing Concentration Data]	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003
1[Supplier Missing Concentration Data]	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
1[Customer Missing Evaluation Data]	0.003	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000
1[Supplier Missing Registration Data]	0.005	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004
1[Supplier Missing Contract Data]	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
1[Supplier Missing Customs Data]	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
1[Supplier Missing Financial Data]	-0.005	-0.005	-0.005	-0.005	-0.005	-0.005	-0.005	-0.005	-0.005	-0.005	-0.005
Number of Bidders	0.000										
Number of Unique 2-Digit Product Codes	-0.003	-0.002	-0.002	-0.002	-0.002	-0.002	-0.003	-0.002	-0.002	-0.001	-0.003
Number of Products	0.000										
Number of Contract Revisions	-0.001	0.000	0.000				0.000		0.000		0.000
Number of Request Revisions	0.000										
Autonomous Organization	-0.001	-0.001	-0.001	-0.001	-0.001	0.000	-0.001	0.000	-0.001	0.000	-0.001
Budget Organization	0.000										
Federal Organization	0.015	0.016	0.016	0.016	0.016	0.016	0.017	0.016	0.016	0.015	0.017
Government Organization	-0.001						0.000		0.000		0.000
Municipal Organization	0.005										
Other Organization	0.001										
Regional Organization	-0.013	-0.020	-0.020	-0.020	-0.020	-0.020	-0.020	-0.020	-0.020	-0.020	-0.020
Organization AKM Estimate	0.251	0.250	0.250	0.250	0.250	0.250	0.249	0.250	0.250	0.250	0.249
Time to Prepare Documents	0.000										
Quantity (log)	-0.005	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.005	-0.004
Reserve Price	0.000										
Share of Individual Bidders	0.000										
Share of Foreign-owned Bidders	-0.001	0.000	0.000	0.000	0.000	0.000	-0.001	0.000	0.000		-0.001
Share of Bidders in Ruslana	-0.001						0.000				0.000
Share of Exporter Bidders	0.000										
Share of Government Bidders	0.000										
Share of Bidders From Same County	-0.001										
Share of Importer Bidders	0.000										
Share of Bidders with Large Profits	0.001										
Share of Bidders with Large Revenue	-0.001										
Share of Wholesaler Bidders	0.002	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.000	0.001
Share of Bidders from Same Region	0.002	0.001	0.001	0.001	0.001	0.001	0.002	0.001	0.001	0.001	0.002
Bureaucrat Success Rate	0.000										
Customer Success Rate	0.002	0.001	0.001	0.001	0.001	0.001	0.002	0.001	0.001	0.001	0.002
Supplier Number of Auctions Run (Cumulative)	0.000										
Supplier Value of Auctions Run (Cumulative)	0.000										
Supplier Number of Auctions Run (Month)	-0.002						0.000				0.000
Supplier Value of Auctions Run (Month)	0.000										
1[Supplier has Large Profits]	0.001						0.000		0.000		0.001
1[Supplier has Large Revenue]	0.000										
1[Supplier has Foreign Ownership]	0.000										
1[Supplier is NGO]	0.001										0.000
Supplier - Number of Export Destinations	0.000										
Supplier - Number of Employees (log)	0.001										
1[Exporter Supplier]	0.001										
1[Supplier is Federal Government Agency]	0.001										
Supplier Age (log)	0.000										
1[Importer Supplier]	0.000										0.000
Supplier - Perc. Revenue from Exports	0.000										
Supplier - Perc. Costs from Imports	0.000										
Supplier - Perc. Revenue from Contracts	0.000										
Supplier - Perc. Revenue from SOE Contracts	0.000										
Supplier - Net Profit	0.000										
1[Supplier is Private Company]	0.000										
Supplier Profit per Worker	0.000										
1[Supplier is Regional Government Agency]	-0.001										
Supplier - Total Assets (log)	-0.001										
Supplier - Turnover (log)	0.000										
Supplier - Unique Export Products	0.000										
Supplier - Unique Import Products	0.001										0.000
1[Contract Canceled]	0.000										
Number of Applicants	-0.004	-0.003	-0.003	-0.003	-0.003	-0.003	-0.004	-0.003	-0.004	-0.003	-0.004
Bureaucrat Product HHI Index (Volume)	0.001										
Supplier Product HHI Index (Volume)	-0.001										
1[Supplier is Registered Firm]	0.000										
1[Supplier is Entrepreneur]	-0.001						0.000				0.000
1[Supplier is From Same Region]	0.001						0.000				0.000

This table shows the nonzero coefficients predicting bureaucrat effectiveness taken from the elastic net regularization procedures that had values of the regularization penalty λ that minimized the mean squared error in K-fold cross-validation. Each column shows the coefficients from different mixing parameters being used, which can range from 0 to 1 for the basic elastic net approach. Note that we require the regularization procedure to include nine indicators in the elastic net procedure due their importance in predicting the bureaucrat effectiveness estimates: quantity (log), the organization effectiveness estimate, and seven indicators of missingness in data on suppliers, bureaucrats, and customers.

TABLE OA.9: ELASTIC NET NONZERO COEFFICIENTS ARE ROBUST ACROSS MIXING PARAMETERS (OGRANIZATION EFFECTIVENESS)

Variable	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Admission Rate to Auction	0.000										
Bureaucrat Product HHI Index (Auctions)	0.001	0.000	0.000	0.000		0.000	0.000	0.000	0.000	0.000	0.000
Supplier Product HHI Index (Auctions)	-0.002	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
1[Auction Winner Not Supplier]	0.001		0.000	0.000		0.000	0.000		0.000		
Auction Length	0.001	0.000	0.001	0.000		0.001	0.000	0.000	0.001	0.000	0.000
1[Problem with Execution]	0.000										
1[Auction Held]	0.000										
1[Bur. Bought Product Often (Auctions)]	-0.003	-0.002	-0.003	-0.002	-0.001	-0.003	-0.002	-0.002	-0.003	-0.002	-0.002
1[Supp. Bought Product Often (Auctions)]	0.000										
1[Bur. Bought Product Often (Volume)]	0.001	0.000	0.001	0.000		0.001	0.001		0.001	0.000	0.000
1[Supp. Bought Product Often (Volume)]	-0.001	0.000	-0.001	0.000		-0.001	-0.001		-0.001	0.000	0.000
Bureaucrat Number of Auctions Run (Cumulative)	-0.001										
Bureaucrat Value of Auctions Run (Cumulative)	0.001		0.000			0.000			0.000		
Bureaucrat Number of Auctions Run (Month)	0.002	0.001	0.003	0.002	0.000	0.003	0.002	0.001	0.003	0.002	0.002
Bureaucrat Value of Auctions Run (Month)	-0.005	-0.004	-0.005	-0.004	-0.003	-0.005	-0.005	-0.004	-0.005	-0.004	-0.004
Bureaucrat AKM Estimate	0.322	0.320	0.319	0.320	0.320	0.319	0.319	0.320	0.319	0.320	0.320
1[Customer works in Agriculture]	-0.002	-0.002	-0.003	-0.002	-0.002	-0.003	-0.002	-0.002	-0.003	-0.002	-0.002
1[Customer works in Culture]	0.008	0.008	0.009	0.008	0.008	0.009	0.008	0.008	0.009	0.008	0.008
Distance from Regional Capital	-0.019	-0.021	-0.022	-0.021	-0.021	-0.022	-0.022	-0.021	-0.022	-0.021	-0.021
1[Customer works in Education]	0.000		-0.001			-0.001	0.000		-0.001		
1[Customer works in Emergency Services]	0.000		0.000			0.000	0.000		0.000		
1[Customer works in Environment]	-0.001	0.000	-0.002	-0.001	0.000	-0.002	-0.001	-0.001	-0.002	-0.001	-0.001
1[Customer works in Forestry]	0.000										
1[Customer works in Health Care]	-0.020	-0.023	-0.027	-0.025	-0.023	-0.027	-0.026	-0.024	-0.027	-0.025	-0.024
1[Customer works in Housing]	0.004	0.004	0.004	0.004	0.003	0.004	0.004	0.004	0.004	0.004	0.004
1[Customer works in Internal]	0.013	0.013	0.014	0.014	0.013	0.014	0.014	0.014	0.014	0.014	0.014
1[Customer works in Labor]	-0.001		0.000	0.000		0.000	0.000		0.000	0.000	
1[Customer works in Mining]	-0.001		0.000			0.000	0.000		0.000		
1[Customer works in Natural Resources]	0.000		0.000			0.000			0.000		
1[Customer works in News]	0.000										
Customer Evaluation Score	0.002	0.001	0.002	0.001	0.000	0.001	0.001	0.000	0.001	0.001	0.001
1[Customer works in Social Services]	-0.006	-0.007	-0.009	-0.008	-0.007	-0.009	-0.008	-0.008	-0.009	-0.008	-0.008
1[Customer works in Sport]	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
1[Customer works in Television]	0.001		0.000	0.000		0.000	0.000		0.000		
1[Customer works in Transportation]	-0.004	-0.004	-0.004	-0.004	-0.003	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
1[Customer works in Veterinary Services]	0.000		0.000			0.000			0.000		
1[Customer works in Youth Services]	0.000		0.000			0.000			0.000		
1[Delayed Delivery]	-0.001										
1[Early Delivery]	0.001		0.000	0.000		0.000	0.000		0.000		
Deposit / Reserve Price	0.002	0.001	0.002	0.001	0.001	0.002	0.001	0.001	0.002	0.001	0.001
Length of Delay	0.001		0.000								
Number of Days Early	0.000										
In-house Bureaucrat	0.024	0.026	0.027	0.027	0.026	0.027	0.027	0.026	0.027	0.027	0.027
Average of Losing Bids / Winning Bid	0.001		0.000			0.000	0.000		0.000		
1[Bureaucrat Missing Concentration Data]	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
1[Supplier Missing Concentration Data]	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
1[Customer Missing Evaluation Data]	0.010	0.008	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007
1[Supplier Missing Registration Data]	-0.009	-0.008	-0.009	-0.008	-0.008	-0.009	-0.009	-0.009	-0.009	-0.009	-0.009
1[Supplier Missing Contract Data]	0.002	0.000	0.001	0.000	0.000	0.001	0.000	0.000	0.001	0.000	0.000
1[Supplier Missing Customs Data]	0.007	0.008	0.008	0.008	0.008	0.008	0.008	0.008	0.008	0.008	0.008
1[Supplier Missing Financial Data]	0.000	0.002	0.001	0.001	0.002	0.001	0.001	0.002	0.001	0.001	0.002
Number of Bidders	0.000										
Number of Unique 2-Digit Product Codes	-0.003	-0.003	-0.004	-0.003	-0.002	-0.003	-0.003	-0.003	-0.003	-0.003	-0.003
Number of Products	-0.007	-0.007	-0.008	-0.007	-0.007	-0.008	-0.008	-0.007	-0.008	-0.007	-0.007
Number of Contract Revisions	0.004	0.003	0.004	0.003	0.003	0.004	0.004	0.003	0.004	0.003	0.003
Number of Request Revisions	-0.001	0.000	-0.001	0.000		-0.001	-0.001	0.000	-0.001	0.000	0.000
Autonomous Organization	0.000										
Budget Organization	-0.001		0.000			0.000			0.000		
Federal Organization	0.030	0.037	0.038	0.037	0.037	0.038	0.038	0.037	0.038	0.038	0.038
Government Organization	0.002	0.001	0.002	0.002	0.001	0.002	0.002	0.001	0.002	0.002	0.001
Municipal Organization	-0.013	-0.009	-0.010	-0.009	-0.008	-0.010	-0.010	-0.009	-0.010	-0.009	-0.009
Other Organization	-0.001										
Regional Organization	-0.006										
Time to Prepare Documents	-0.001		0.000			0.000	0.000		0.000		
Quantity (log)	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005	0.005
Reserve Price	0.008	0.007	0.008	0.008	0.007	0.008	0.008	0.007	0.008	0.007	0.007
Share of Individual Bidders	0.001										
Share of Foreign-owned Bidders	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Share of Bidders in Ruslana	0.001		0.000			0.000			0.000		
Share of Exporter Bidders	0.001	0.000	0.001	0.000	0.000	0.001	0.001	0.000	0.001	0.000	0.000
Share of Government Bidders	0.002						0.000				
Share of Bidders From Same County	0.000		-0.001	0.000		-0.001	0.000		-0.001		
Share of Importer Bidders	-0.002	0.000	-0.001	0.000		-0.001	-0.001		-0.001	0.000	0.000
Share of Bidders with Large Profits	0.000										
Share of Bidders with Large Revenue	0.001		0.000			0.000	0.000		0.000		
Share of Wholesaler Bidders	0.001	0.000	0.001	0.000		0.001	0.001	0.000	0.001	0.000	0.000
Share of Bidders from Same Region	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
Bureaucrat Success Rate	-0.001	0.000	-0.001	-0.001		-0.001	-0.001	0.000	-0.001	-0.001	0.000
Customer Success Rate	0.006	0.006	0.006	0.006	0.005	0.006	0.006	0.006	0.006	0.006	0.006
Supplier Number of Auctions Run (Cumulative)	-0.004		0.000								
Supplier Value of Auctions Run (Cumulative)	0.004	0.002	0.002	0.002	0.001	0.002	0.002	0.002	0.002	0.002	0.002
Supplier Number of Auctions Run (Month)	0.002		0.000			0.000	0.000		0.000		
Supplier Value of Auctions Run (Month)	-0.001										
1[Supplier has Large Profits]	-0.001										
1[Supplier has Large Revenue]	-0.001										
1[Supplier has Foreign Ownership]	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
1[Supplier is NGO]	0.000										
Supplier - Number of Export Destinations	0.000										
Supplier - Number of Employees (log)	-0.001										
1[Exporter Supplier]	0.001	0.001	0.002	0.002	0.001	0.002	0.002	0.001	0.002	0.002	0.002
1[Supplier is Federal Government Agency]	0.002	0.002	0.003	0.003	0.002	0.003	0.003	0.002	0.003	0.003	0.003
Supplier Age (log)	0.000										
1[Importer Supplier]	0.000										
Supplier - Perc. Revenue from Exports	0.000										
Supplier - Perc. Costs from Imports	-0.001	0.000	-0.001	0.000		-0.001	0.000	0.000	-0.001	0.000	0.000
Supplier - Perc. Revenue from Contracts	0.000										
Supplier - Perc. Revenue from SOE Contracts	-0.001	-0.001	-0.001	-0.001	0.000	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
Supplier - Net Profit	0.001										
1[Supplier is Private Company]	-0.002		-0.001			-0.001	0.000		-0.001		
Supplier Profit per Worker	0.000										
1[Supplier is Regional Government Agency]	-0.002										
Supplier - Total Assets (log)	0.002		0.001	0.000		0.001	0.001		0.001	0.000	
Supplier - Turnover (log)	0.001										
Supplier - Unique Export Products	-0.003	-0.002	-0.003	-0.002	-0.001	-0.003	-0.002	-0.002	-0.003	-0.002	-0.002
Supplier - Unique Import Products	0.001										
1[Contract Canceled]	0.000		0.000			0.000			0.000		
Number of Applicants	-0.006	-0.005	-0.006	-0.006	-0.005	-0.006	-0.006	-0.005	-0.006	-0.006	-0.006
Bureaucrat Product HHI Index (Volume)	0.006	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007
Supplier Product HHI Index (Volume)	-0.002	-0.001	-0.002	-0.002	-0.001	-0.002	-0.002	-0.001	-0.002	-0.002	-0.001
1[Supplier is Registered Firm]	0.000										
1[Supplier is Entrepreneur]	0.000										
1[Supplier is From Same Region]	-0.002	-0.001	-0.001	-0.001	0.000	-0.001	-0.001	0.000	-0.001	-0.001	0.000

This table shows the nonzero coefficients predicting organization effectiveness taken from the elastic net regularization procedures that had values of the regularization penalty λ that minimized the mean squared error in K-fold cross-validation. Each column shows the coefficients from different mixing parameters being used, which can range from 0 to 1 for the basic elastic net approach. Note that we require the regularization procedure to include nine indicators in the elastic net procedure due their importance in predicting the bureaucrat effectiveness estimates: quantity (log), the organization effectiveness estimate, and seven indicators of missingness in data on suppliers, bureaucrats, and customers.

TABLE OA.10: BID PREFERENCES ARE MORE EFFECTIVE WHEN IMPLEMENTED BY LESS EFFECTIVE BUREAUCRATS (USING RAW FIXED EFFECTS ESTIMATES OF EFFECTIVENESS)

	All Products		Pharmaceuticals		
	Log Price	No. Bidders	Log Price	No. Bidders	Domestic Winner
	(1)	(2)	(3)	(4)	(5)
log Standardized Quantity	-0.330^{***}	0.045^{***}	-0.027^{***}	0.009^{**}	0.003^{***}
	(0.014)	(0.003)	(0.003)	(0.003)	(0.001)
Bureaucrat FE * Preferenced * Policy Active	-0.209^{***}	0.064^{**}	-0.121^{*}	0.366^{**}	0.081^{***}
	(0.020)	(0.025)	(0.066)	(0.149)	(0.029)
Organization FE * Preferenced * Policy Active	-0.155^{***}	0.058^{***}	-0.004	0.276^{**}	0.019
	(0.018)	(0.021)	(0.063)	(0.114)	(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

*** p<0.01, ** p<0.05, * p<0.1 This table estimates the triple-differences equation (14): $p_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Preferenced}_{gt} \times \text{PolicyActive}_t + \gamma\text{Preferenced}_{gt} \times \hat{\alpha}_b + \zeta\text{Preferenced}_{gt} \times \hat{\psi}_j + \eta\text{PolicyActive}_t \times \hat{\alpha}_b + \theta\text{PolicyActive}_t \times \hat{\psi}_j + \pi\text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \nu\text{Preferenced}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt}$. 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 Preferenced = 1 if the type of good appears on the list of goods covered by the preferences policy for that year. In the Pharmaceuticals sample, Preferenced = 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.11: 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 (10) in different samples. The decomposition uses the fixed effect estimates from equation (9): $p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Column (2) replicates the findings in column (1) of table 3. 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.3 instead of logistic regression. Column (4) uses the hierarchical classifier described in Section OA.3 instead of logistic regression.

TABLE OA.12: 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 (10) in different samples. The decomposition uses the fixed effect estimates from equation (9): $p_i = \mathbf{X}_i \boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Column (5) replicates the findings in column (1) of table 3. As described in Section OA.3, column (4) deemed sentences to be correctly classified if their predicted value $\hat{g}_C(\mathbf{x}_i)$ 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.