Individuals and Organizations as Sources of State Effectiveness, and Consequences for Policy Design*

Michael Carlos Best
Stanford Institute for Economic Policy Research & CEPR

Jonas Hjort
Columbia University & BREAD & NBER

David Szakonyi
George Washington University & Higher School of Economics

December 8, 2016

For the most recent version, please visit http://bit.ly/2dlCTyo

Abstract

Why does state effectiveness vary so much both across and within countries? Can we attribute the variation to the individuals and organizations carrying out the tasks defined by government policies? And what are the implications for policy design? We first use a new text-based machine learning method to assign the goods purchased in 25 million public procurement auctions conducted in Russia from 2011 to 2016 to comparable categories. We show that the individual bureaucrats and organizations in charge of procurement together explain one-third of the within-category variation in prices achieved, and that effective procurers lower auction entry costs, which in turn lowers prices. We then analyze the implications of heterogeneity in effectiveness for the impact of a ubiquitous procurement policy: granting bid preferences to a specific group of bidders. Consistent with a simple endogenous entry auction model with variation in auctioneer effectiveness, we find that when the bureaucracy is effective, favoring firms supplying domestically produced goods lowers entry and increases prices, but when effectiveness is low, the effect is reversed. These results demonstrate that there are large returns to the state from improving bureaucratic effectiveness, but that appropriately designed policies can compensate for low effectiveness.

*mbest@stanford.edu, hjort@columbia.edu, dszakonyi@gwu.edu We thank John M. Abowd, Daron Acemoglu, Richard Blundell, Raj Chetty, Allan Collard-Wexler, Francesco Decarolis, Ben Faber, Matthew Gentzkow, Josh Gottlieb, Caroline Hoxby, Amit Khandelwal, Brad Larsen, Imran Rasul, Jimmy Roberts, Jonah Rockoff, Orie Shelef, Andrei Shleifer, and Danny Yagan for helpful discussions; Andrei Yakovlev for institutional guidance; Georgiy Syunyaev and Andrey Druzhaev for outstanding research assistance; and seminar participants at U.C. Berkeley, Central European University, University of Copenhagen, Columbia, Duke, the Econometric Society, McGill, NEUDC, Santa Clara, Stanford, University College London, and University of Toronto for comments. 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'. The plethora of remaining errors are ours alone.

1 Introduction

Many policies work well in some countries or regions and poorly in others. Value-added taxes generate the intended paper trails and tax compliance in most developed countries, but rarely do so in developing countries (Bird & Gendron, 2007; Collier, 2014). The NREGA employment guarantee scheme supports poor workers and helpes complete important infrastructure programs in some Indian states, but is largely unused in others (Gulzar & Pasquale, forthcoming). 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). The list of examples is long and covers nearly all areas of policymaking (Rodrik, 2009). At the same time, recent research has documented dramatic differences in the characteristics of the bureaucrats and organizations that implement states' policies both across and within countries.[1] To what extent does the effectiveness of the bureaucratic apparatus help explain the variation in public sector "output" achieved under a given policy regime? And what are the implications for policy design?

To shed light on these questions, we focus on a well-defined form of output produced throughout the public sector – prices paid for goods procured – and use administrative data covering the universe of procurement auctions in Russia from 2011 to 2016. We show that under a standard policy regime that treats all suppliers equally, a third of the variation in prices paid is attributable to the bureaucratic apparatus. However, when a policy regime favoring domestic suppliers is introduced, supplier entry and prices improve in auctions run by ineffective procurers, but worsen in auctions run by effective procurers. Our results demonstrate that there are large returns to increasing bureaucratic effectiveness, but that policies designed taking the existing level of bureaucratic effectiveness into account can act as a partial substitute for increasing individual and organizational effectiveness.

There are three parts to our empirical analysis. We start by developing a text-based machine learning method that assigns procurement purchases to homogeneous bins. This allows us to compare bureaucrats and organizations across the country performing the same task. We then exploit the fact that many bureaucrats (procurement officers) are observed working with multiple end-user organizations (for example ministries, schools or hospitals) and vice versa, providing us with thousands of quasi-experiments that can be used to estimate the causal effect of specific bureaucrats and organizations on prices paid. To do so, we combine the variance decomposition method introduced by Abowd *et al.* (1999, 2002) with split-sample and shrinkage tools to correct for sampling error. In the third part of the paper, instead of holding the policy environment constant and varying the procurers, we hold the bureaucrat and organization constant and vary whether a particular procurement policy applies. We study a policy that provides bid preferences to suppliers of domestically produced goods. Using the fact that the policy applies to an evolving set of goods and is "turned off" parts of each calendar year, we estimate its average impact on prices paid, and how the impact differs for bureaucrats and organizations of different levels of baseline effectiveness.

To guide our empirical analysis, we develop a stylized model of public procurement auctions. In the

[1] Finan *et al.* (forthcoming) present a great overview of the literature.

model, potential suppliers decide whether to participate in procurement auctions by trading off their expected profits against the participation costs imposed on them by the government's bureaucracy. As a result, ineffective bureaucracies that impose high participation costs attract fewer participants and pay higher prices for the items they procure.[2] Introducing bid preferences for domestically produced goods, which we assume to be more expensive on average has two effects. First, it encourages entry by suppliers of such goods, and second, it discourages entry by suppliers of foreign goods, such that the net effect depends on the baseline level of entry. We show that for ineffective bureaucracies, the expected net effect is higher entry and lower prices, while for effective bureaucracies, the expected net effect is the opposite.

We present four main empirical findings. First, we show that both individual bureaucrats and individual public sector organizations significantly affect prices paid in a standard policy regime that treats all suppliers equally. Together, the individuals and organizations of the bureaucratic apparatus account for a third of the variation in the prices the government pays for its inputs. Of this effect, bureaucrats and organizations each account for roughly equal shares of the variation. A battery of tests gives no indication that the additivity of bureaucrat and organization effects and "exogenous mobility" assumptions needed to interpret our variance decomposition results causally are violated (as also tends to hold in the labor economics literature on workers and firms in the private sector[3]). The variance decomposition exercise thus informs us of the degree to which state effectiveness, in weak institutional contexts such as Russia, can be enhanced by attracting more individuals at the high end of the performance range observed public sector-wide, or by lifting organization-wide characteristics such as management or "organizational culture" towards the high end of the range.

Second, having estimated the effectiveness of each of the bureaucrats and organizations in our sample, we correlate these reduced-form performance measures with a rich set of indicators on how successful and less successful procurement auctions play out to determine what effective procurers do differently. We find that bureaucrats that perform well make fewer amendments to auction requests, demand lower security deposits, and end up attracting more entrants, consistent with our conceptual framework in which bureaucratic effectiveness lowers entry costs. Similarly, organizations that lower entry costs perform better.[4]

Third, our difference-in-differences analysis of Russia's "buy local" policy shows that, on average, the bid preferences achieve the goal of channeling demand to the intended beneficiaries – suppliers of domestically-made goods – at no cost to the government in that average prices paid are unaffected. This contrasts with the results of studies of similar preference policies using more structural empirical methods in higher state effectiveness contexts. For example, a 5 percent bid preference for small businesses in Californian road construction procurement is estimated to increase average costs by between 1 percent and 4 percent (Marion, 2007; Krasnokutskaya & Seim, 2011).

Fourth, interacting the "buy local" policy with our estimates of bureaucrats' and organizations' effec-

[2]See Samuelson (1985); Levin & Smith (1994); Bulow & Klemperer (1996); Menezes & Monteiro (2000); Gentry & Li (2014); Li & Zhang (2015).

[3]See e.g. Mendes *et al.* (2010); Card *et al.* (2013a,b, 2015); Goldschmidt & Schmieder (2015); Shelef & Nugyen-Chyung (2015); Alvarez *et al.* (2016).

[4]This evidence complements an innovative paper by Lacetera *et al.* (2016) studying auctioneer effects in a different setting, by shedding light on the channels through which successful auctioneers in the public sector are able to achieve lower prices.

tiveness reveals that the average treatment effects obscure considerably heterogeneity across the range of policy implementer effectiveness. The prices achieved by ineffective bureaucrats and organizations *decrease* by about 15 percent when preferences apply, while the prices achieved by effective procurers increase by a similar magnitude when preferences apply. Because ineffective bureaucrats and organizations have low baseline entry rates, for them, increased entry by suppliers of domestic goods outweighs the tilting of the playing field against importers. The opposite is true for effective procurers, consistent with our conceptual framework.

This paper contributes to the literatures on state effectiveness, individual workers and firms in the labor market, and methods for "transporting" estimated treatment effects across settings. The particular state activity we consider is public procurement, which makes up roughly 8 percent of worldwide GDP and 10 percent of Russia's non-resource GDP (Schapper *et al.* , 2009). Procurement is one of the few state activities where "output" produced – prices obtained – is well-defined, measurable, and comparable across the entire public sector. The arrival of comprehensive procurement data has thus enabled researchers to begin to answer some of the key open questions on state effectiveness (see e.g. Bandiera *et al.* , 2009; Ferraz *et al.* , 2015; Lewis-Faupel *et al.* , 2015). We make a methodological contribution to this and other literatures where classification of goods is a challenge: we show how text analysis can be used to ensure that within-category quality differences are minimal while maintaining generality by not restricting the sample to very specific types of goods.[5] Substantively, we demonstrate the extent to which state effectiveness depends on the individuals and organizations that make up the bureaucratic apparatus. This evidence adds to and extends recent studies on the more localized benefits of attracting "better" public sector workers (see e.g. Dal Bo *et al.* , 2013; Ashraf *et al.* , 2014; Hanna & Wang, 2015; Bertrand *et al.* , 2016; Callen *et al.* , 2016; Deserranno, 2016) or improving management or organizational culture (see e.g. Bandiera *et al.* , 2011; Rasul & Rogger, forthcoming; Bloom *et al.* , 2015a,b).

Having demonstrated that state effectiveness is partly embodied in the bureaucratic apparatus, we document that there are important consequences for policy design. The type of procurement policy we focus on – favoring a specific group of firms – is very common in both rich and poor countries (see Athey *et al.* , 2013). Our results add to those from studies of the effects of asymmetric auction rules on participation and prices (see for example Brannman & Froeb, 2000; Flambard & Perrigne, 2006; Marion, 2007; Krasnokutskaya & Seim, 2011; Bhattacharya *et al.* , 2014). However, using our effectiveness measures, we are able to go beyond average effects to explore how optimal policy depends on the individuals and organizations in charge of procurement. We show a specific way in which the policy rules governing public procurement should arguably differ when the entities implementing policy are more versus less effective. These results build on recent theoretical (Laffont, 2005; Estache & Wren-Lewis, 2009; Besley & Persson, 2009) and empirical work (Greenstone & Jack, 2015; Best *et al.* , 2015; Burgess *et al.* , 2012; Duflo

[5] The difficulty of categorizing goods accurately so as to ensure like for like comparisons has long dogged several literatures. Existing work tends to take one of three approaches: Pricing (observable) quality attributes through hedonic regressions; comparing only within pre-existing good categories such as the codes assigned to traded products by customs authorities; and/or restricting attention to especially homogeneous goods such as cement or block ice. In instead using text analysis to classify categories to compare, we follow an innovative study by Hoberg & Phillips (2016). Their method has similarities to ours, but differs in that they classify *firm* similarity based on text listing the various goods firms produce (we instead classify *good* similarity based on text describing each purchase), and in that they pre-specify the desired number of clusters/categories.

et al., 2013, 2014; Jia, 2014) starting to unpack how policies that work well in rich countries may not work well in lower income countries, and how precisely policy should be tailored to context.

We also contribute to the literature that estimates effects attributable to individual workers and firms in the labor market. We follow the seminal work of Abowd *et al.* (1999, 2002) (hereafter "AKM") on private sector labor markets showing how worker and firm fixed effects can be separately identified within sets connected by worker mobility.[6] This paper is to our knowledge the first to estimate how the individuals and organizations that make up the public sector matter.[7] In addition, our application differs in that, rather than wages, we study direct measures of performance (or "productivity").[8] While private sector applications have combined workers performing different tasks,[9] we are able to focus on bureaucrats and organizations who perform the same task. We also show how to combine the AKM method with tools to correct for sampling error deriving from having finite samples on each individual and each organization (Scott, 1948; Lancaster, 2000). To do so, we adapt split-sample methods (Finkelstein *et al.*, 2016; Silver, 2016) and methods directly estimating the signal and noise components of the estimates' variances and combining them to form minimum-mean-squared-error predictors, akin to Kane & Staiger (2008); Chetty *et al.* (2014); Chetty & Hendren (2015).[10] This allows us to estimate how much of the variation in output across the public sector is explained by the bureaucratic apparatus.[11]

Finally, we contribute to a recent methodological literature studying how an average treatment effect (ATE) estimated in one setting can be used to predict the effects of the relevant policy or program in another (Vivalt, 2016; Dehejia *et al.*, 2016; Gechter, 2016; Rokkanen, 2016). In decomposing the ATE into conditional treatment effects that are specific to bureaucrats and organizations of a given level of effectiveness, we follow the literature on heterogeneous treatment effects (see e.g. Heckman & Smith, 1997; Angrist, 2004; Deaton, 2010; Heckman, 2010) – although few previous studies consider treatment effects that condition on an unobserved (and therefore estimated) characteristic such as effectiveness. Our findings extend those of recent studies comparing program effects across branches of private firms or private-versus-public status of the implementing agency (see Bold *et al.*, 2013; Allcott, 2015; Vivalt, 2015; Blader *et al.*, 2016) by documenting how the impact of public policies can depend on the effectiveness of the particular individual or organization in the public sector who is in charge of implementation.

The rest of the paper is organized as follows. Section 2 presents an endogenous entry auction model with variation in auctioneer ability that guides our analysis. Background on the Russian public procure-

[6] 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 3. See also Bertrand & Schoar (2003) and the literature that followed on CEO effects.

[7] Jones & Olken (2005) and Yao & Zhang (2015) show how national and sub-national political leaders matter for economic growth.

[8] Wages do not necessarily reflect productivity (Eeckhout & Kircher, 2011; Card *et al.*, 2015), but are important objects in and of themselves.

[9] Carneiro *et al.* (2012) and Cardoso *et al.* (2016) show the potential importance of accounting for differences in tasks.

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

[11] In a related innovative paper, Bertrand *et al.* (2016) study how the incentives of elite bureaucrats in Indian states (due to variation in cohort size/competition for promotion) matter for bureaucratic performance and aggregate outcomes. Duflo *et al.* (2013); Khan *et al.* (2016); Callen *et al.* (2016); Deserranno (2016) also present important experimental evidence on how performance incentives affect bureaucrats' or public sector workers' performance. The reason that the methodology we use has not been feasible in previous work is partly that measuring output and thus productivity in the public sector is notoriously difficult, and public sector–wide datasets with sufficient power from countries where state effectiveness is an issue are rare.

ment system and information on the data we use is in sections 3 and 4. In Section 5, we estimate the effectiveness of individual bureaucrats and organizations and their contribution to public sector output. In Section 6 we analyze the impact of the "buy local" policy and its interaction with procurer effectiveness. Section 7 concludes.

2 Conceptual Framework

In this section we present a stylized model of public procurement in which two potential suppliers choose whether to try to get to sell an item to the government. The government uses an auction to award the contract and determine the price at which it buys the item. Suppliers must pay an entry cost to enter the auction; these entry costs serve as our reduced-form device for modeling state effectiveness. In Sub-section 2.1 we trace out how the level of state effectiveness affects supplier participation and prices achieved in procurement. Then, in Sub-section 2.2, we show how introducing bidding preferences for specific types of suppliers can have opposite effects depending on whether state effectiveness is high or low.

2.1 A simple model of procurement auctions with endogenous entry

Consider a government wishing to purchase an item from one of two potential suppliers. To make the purchase, the government uses a second-price descending auction with a publicly announced reservation price normalized to 1. In order to participate in the auction, bidders must pay a participation cost of c. This c represents the direct costs of preparing the technical and other documents required to participate, the liquidity costs of paying the deposit for participation, and the cost of attending the online auction. c may also include side payments to the procurer.[12]

In the first stage of the procurement process, the two potential suppliers, F and L, observe the announcement of the item to be procured and the participation cost c, and each supplier privately learns her cost of fulfilling the contract, v_i, $i = F, L$. The suppliers' fulfillment costs are independently distributed, but bidder F is, on average, more efficient than bidder L. To capture this as simply as possible, we assume that both bidders' fulfillment costs are uniformly distributed with CDFs $G_F(v_F) = \mathcal{U}[0,1]$ and $G_L(v_L) = \mathcal{U}[\mu, 1]$, where $0 < \mu < 1$.[13] Upon learning their fulfillment cost, the suppliers simultaneously decide whether or not to pay the entry cost and enter the auction.

In the second stage of the procurement process, if only one supplier chose to enter the auction, she is awarded the contract at the reservation price of 1. If neither supplier chose to enter, the procurer randomly picks a supplier and awards her the contract at a price of 1.[14] Finally, if both suppliers enter,

[12] Bandiera *et al.* (2009) find that 80 percent of waste in Italian procurement is due to low bureaucratic ability rather than corruption. In this paper we are agnostic about whether some procurers display higher entry costs than others because they are corrupt or because they are less effective. In our framework, the two sources of entry costs have the same impact on the equilibrium outcomes we focus on, namely participation and prices.

[13] The positions of the upper and lower bounds of the distribution are innocuous. Uniformity is a simplifying assumption that, while unrealistic, allows us to derive simple, closed-form expressions.

[14] A more realistic assumption would be that if no supliers enter the procurer has to re-run the auction at some cost, which would make the model dynamic. The assumption we make simplifies the exposition by making the model static. The qualita-

they take part in a second-price descending auction.

Both suppliers choose their entry and bidding strategies to maximize expected profits. Since bidder valuations are independent, it is a dominant strategy for bidders to bid their fulfillment cost in the auction. Denoting the bidding strategy of supplier i with fulfillment cost x by $b_i(x)$, we have $b_F(x) = b_L(x) = x$.[15] As a result, the winner is the bidder with the lowest fulfillment cost; she receives the contract at the other bidder's fulfillment cost. At the entry stage, we posit that the equilibrium involves supplier F entering if her fulfillment cost is below a threshold value $\overline{d}_F$, and bidder L entering if her fulfillment cost is below a threshold $\overline{d}_L$.[16]

We outline the equilibrium here, relegating a detailed characterization of the equilibrium and the proofs of propositions to Appendix A. Working backwards from the second stage, we write supplier i's expected profits if she enters with fulfillment cost v and suppliers enter according to $\overline{d}_F,\ \overline{d}_L$ as

$$U_i\left(v;\overline{d}_F,\overline{d}_L\right) = m_i\left(v;\overline{d}_F,\overline{d}_L\right) - q_i\left(v;\overline{d}_F,\overline{d}_L\right)v \tag{1}$$

where $m_i\left(v;\overline{d}_F,\overline{d}_L\right)$ is the expected payment supplier i receives if she enters with fulfillment cost v, and $q_i\left(v;\overline{d}_F,\overline{d}_L\right)$ is the probability that supplier i receives the contract if she enters when her fulfillment cost is v. The probabilities of winning are given by

$$q_i\left(v;\overline{d}_F,\overline{d}_L\right) = \Pr\left(b_i(v) < b_j(v_j)\,|v_g \leq \overline{d}_j\right)\Pr\left(v_j \leq \overline{d}_j\right) + \Pr\left(v_j > \overline{d}_j\right)\ \ i,j \in \{F,L\},\ i \neq j \tag{2}$$

Since the bidding strategies are chosen optimally, we can use the integral-form envelope theorem (Milgrom & Segal, 2002; Milgrom, 2004) to rewrite expected net profits and expected payments as[17]

$$\begin{aligned} U_i\left(v;\overline{d}_F,\overline{d}_L\right) &= \int_v^1 q_i\left(x;\overline{d}_F,\overline{d}_L\right)dx \\ m_i\left(v;\overline{d}_F,\overline{d}_L\right) &= \int_v^1 q_i\left(x;\overline{d}_F,\overline{d}_L\right)dx + q_i\left(v;\overline{d}_F,\overline{d}_L\right)v \end{aligned}$$

The entry thresholds are given by the suppliers who are indifferent between entering and paying the entry cost, and staying out and receiving the contract with probability 1/2 if the other supplier also stays out. That is, the entry thresholds satisfy

$$U_F\left(\overline{d}_F;\overline{d}_F,\overline{d}_L\right) - c = \frac{1}{2}(1-\overline{d}_F)\frac{1-\overline{d}_L}{1-\mu} \tag{3}$$

$$U_L\left(\overline{d}_L;\overline{d}_F,\overline{d}_L\right) - c = \frac{1}{2}(1-\overline{d}_F)(1-\overline{d}_L) \tag{4}$$

tive results are unlikely to be affected by this simplification.

[15] See, for example, Milgrom (2004) or Krishna (2010).

[16] This is the equilibrium that the auction literature with endogenous entry has focussed on (see, for example, Samuelson (1985), Krasnokutskaya & Seim (2011), Roberts & Sweeting (2015), Gentry & Li (2014)), though other equilibria may exist.

[17] Strictly, $U_i\left(v;\overline{d}_F,\overline{d}_L\right) = U_i\left(1;\overline{d}_F,\overline{d}_L\right) + \int_v^1 q_i\left(x;\overline{d}_F,\overline{d}_L\right)dx$. However, since a supplier with fulfillment cost of 1 never makes a profit, $U_i\left(1;\overline{d}_F,\overline{d}_L\right) = 0$.

In this equilibrium, the expected number of entrants is

$$\mathsf{E}[n] = G_F(\overline{d}_F) + G_L(\overline{d}_L) \tag{5}$$

and the expected price the government will pay for the item is

$$\mathsf{E}[p] = \mathsf{E}_{v_F}\left[m_F\left(v_F;\overline{d}_F,\overline{d}_L\right)\right] + \mathsf{E}_{v_L}\left[m_L\left(v_L;\overline{d}_F,\overline{d}_L\right)\right] + \left[1 - G_F\left(\overline{d}_F\right)\right]\left[1 - G_L\left(\overline{d}_L\right)\right] \tag{6}$$

combining expected payments to the entrants with the payment in the case of no entrants.

The following proposition shows how the entry costs that procurers impose on potential suppliers relate to the number of entrants and the prices the government pays.

Proposition 1. *Procurers who impose higher entry costs on potential suppliers (i) attract fewer entrants, and (ii) pay higher prices.*

$$\frac{d\mathsf{E}[n]}{dc} < 0 \quad \& \quad \frac{d\mathsf{E}[p]}{dc} > 0 \tag{7}$$

Proof. See Appendix A.2. □

2.2 Introducing bidding preferences for domestic goods

In the previous Sub-section, while the suppliers were asymmetric, the government treated them symmetrically. In this Sub-section we introduce bidding preferences favoring local products. Specifically, if bidder F bids b_F and wins, she only receives γb_F, where $\gamma \leq 1$, while if bidder L wins, she receives her full bid. In this setting, it is optimal for bidder F to shade her bids so that what is received when she wins is equal to her true fulfillment cost v_F. As a result, her optimal bid function is $b_F(x) = x/\gamma$. Bidder L's optimal strategy is the same as in a standard second-price auction–to bid her true value $b_L(x) = x$. Apart from this, the procurement process is as before.

In this case, the probability of winning is

$$\begin{aligned} q_F\left(x;\overline{d}_F,\overline{d}_L\right) &= \Pr\left(b_F(x) < b_L(v_L)\,|v_L \leq \overline{d}_L\right)\Pr\left(v_L \leq \overline{d}_L\right) + 1 \times \Pr\left(v_L > \overline{d}_L\right) \\ &= \Pr\left(v_L > \frac{x}{\gamma}|v_L \leq \overline{d}_L\right)\frac{\overline{d}_L - \mu}{1-\mu} + \frac{1-\overline{d}_L}{1-\mu} \end{aligned} \tag{8}$$

$$\begin{aligned} q_L\left(x;\overline{d}_F,\overline{d}_L\right) &= \Pr\left(b_L(x) < b_F(v_F)\,|v_F\overline{d}_F\right)\Pr\left(v_F \leq \overline{d}_F\right) + 1 \times \Pr\left(v_F > \overline{d}_F\right) \\ &= \Pr\left(\frac{v_F}{\gamma} > x|v_F \leq \overline{d}_F\right)\overline{d}_F + \left(1 - \overline{d}_F\right) \end{aligned} \tag{9}$$

but otherwise all the steps in characterizing the equilibrium are as before.[18]

The following proposition summarizes the impact of introducing bidding preferences favoring local products, emphasizing how the effects are different depending on the entry costs procurers impose on sellers.

[18] Appendix A.3 contains the details.

Proposition 2. *Bidding preferences favoring local manufacturers have opposite effects for buyers who impose high and low entry costs. For buyers who impose high entry costs, preferences lead them to attract more bidders and pay lower prices, while for bidders who impose low entry costs, preferences lead them to attract fewer bidders and pay higher prices. Price changes and changes in participation rates are monotonically decreasing in baseline prices and participation rates, respectively.*

Formally, (i) for every $\gamma \in \left(\overline{\gamma}_p, 1\right)$*, there exists a* $\tilde{c}_p(\gamma) \in [0, \overline{c}]$ *such that* $E[p|c,\gamma] - E[p|c,\gamma=1] < 0$ *for all* $c > \tilde{c}_p(\gamma)$ *and* $E[p|c,\gamma] - E[p|c,\gamma=1] > 0$ *for all* $c < \tilde{c}_p(\gamma)$*, where* $\overline{\gamma}_p = \arg\min_\gamma [p|\overline{c},\gamma] < 1$.

Similarly, for every $\gamma \in (\overline{\gamma}_n, 1)$ *there exists a* $\tilde{c}_n(\gamma) \in [0, \overline{c}]$ *such that* $E[n|c,\gamma] - E[n|c,\gamma=1] > 0$ *for all* $c > \tilde{c}_n(\gamma)$ *and* $E[n|c,\gamma] - E[n|c,\gamma=1] < 0$ *for all* $c < \tilde{c}_n(\gamma)$*, where* $\overline{\gamma}_n = \arg\max_\gamma E[n|\overline{c},\gamma] < 1$.

Moreover, (ii)

$$\frac{\partial E[p|c,\gamma] - E[p|c,\gamma=1]}{\partial c} < 0 \tag{10}$$

$$\frac{\partial E[n|c,\gamma] - E[n|c,\gamma=1]}{\partial c} > 0 \tag{11}$$

Proof. See Appendix A.4. □

Intuitively, without preferences ($\gamma = 1$), higher cost procurers depress entry and hence raise prices. They do so particularly for local bidders, since they tend to have higher fulfilment costs and hence lower expected profits from entering the auction.[19] Then, when preferences are introduced, this lowers expected profits for foreign suppliers and so discourages their entry. On the other hand, the preferences increase expected profits for local suppliers by giving them a better chance of winning, and so encourage their entry. This latter effect is strongest for high cost procurers, who were suppressing entry by local bidders the most in the absence of preferences. As a result, for high cost procurers the net effect is to increase participation and lower prices. Conversely, for low cost procurers, who weren't suppressing entry by local bidders as much in the absence of preferences, the net effect is to decrease participation and increase prices.

3 Public Procurement Auctions in Russia

3.1 A decentralized system with centralized rules

In 1991, following the collapse of the Soviet Union, and alongside the creation of market institutions, the Russian government created the institutional capacity to perform public procurement. As with most other state institutions, the system created was, and remains, extremely decentralized. Each government agency has the legal authority to make its own purchases and there are no centralized purchases (such as framework contracts).

While the legal authority to make purchases is decentralized, the legal framework governing procurement is centralized. Competitive bidding for all purchases above USD 35,000 became mandatory

[19] Formally, we show in appendix A.3 that the entry thresholds of the foreign and local bidder satisfy $\overline{d}_F - \gamma\overline{d}_L = \sqrt{2\gamma c\mu}$. Hence, the gap between the foreign and local bidders' entry thresholds is increasing in c.

in 1997 (Yakovlev *et al.* , 2010),[20] and in 2005, the procurement rules and regulations governing tender processes at all levels of government were harmonized.[21] The main aims of the legislation were to encourage greater competition, save on government expenditures, increase transparency, and reduce corruption (Krylova & Settles, 2011). For example, new provisions assigned criminal and administrative liability for individuals and legal entities violating anti-monopoly legislation, all enforced by the Federal Antimonopoly Service and Arbitration Courts.[22] In addition, a key innovation of the law was the creation of a centralized official procurement website (http://zakupki.gov.ru/), launched on January 1, 2011, which provides comprehensive information to the public and suppliers about all federal, regional, and municipal level purchases, and which is our main data source.

3.2 Procurement through auctions

During our data period, procurement represents an average of 10 percent of Russia's non-resource GDP, and a large share of the government's budget. Four main types of procurement mechanisms are available: electronic (open) auctions (53.5 percent of procurement), open tenders (19.8 percent), open requests for quotations (2.3 percent), and single source procurement (21.3 percent).[23] Since bureaucrats and organizations may affect procurement outcomes in very different ways under different purchase mechanisms, we restrict our attention to electronic auctions, which represent the bulk of public procurement. Government officials view auctions as the most potent way to reduce the scope for bureaucrats or organizations to collude with sellers, so our estimates of how individual bureaucrats and organizations affect outcomes should be viewed as lower bounds.

Since July 10, 2010, all auctions are conducted through one of five designated web sites.[24] All announcements, protocols, results, and contracts from the auctions on these five sites are also housed on the central nationwide procurement website (http://zakupki.gov.ru/). Figure 1 traces the steps involved in a procurement process; we now go through these.

Each procurement starts with an auction announcement. The announcement contains technical details of the product(s) to be purchased, a maximum allowable price, the required security deposit (between 0.5 and 5 percent of the maximum price), other participation requirements and the date of the electronic auction. Our data, described in greater detail below, contains 5,054,498 announcements. In order to participate in an auction and compete for a contract, suppliers must first obtain accreditation. This

[20]However, the absence of comprehensive legislation at this time contributed to a multitude of problems, including the inability to effectively monitor purchases across the country and gather statistical data on spending (McHenry & Pryamonosov, 2010).

[21]This was done by Federal Law No. 94-FZ On the Placement of Orders for the Procurement of Goods, Work and Services for State and Municipal Needs, which entered into force on January 1, 2006.

[22]Convicted violations could incur significant fines, disqualification, and imprisonment of up to five years for perpetrators.

[23]These four account for roughly 97 percent of all procurement during the time period. Appendix table D.3 shows usage of these methods over time. Other methods used much more rarely include closed auctions, two-stage tenders, and closed two-stage tenders. In interesting contemporaneous work to our paper, Andreyanov *et al.* (2016) use the timing of bid submissions in open requests for quotations to study corruption and collusion in procurement conducted under this mechanism.

[24]The five platforms are run by the Republic of Tatarstan (http://etp.zakazrf.ru/), the Moscow city government (www.roseltorg.ru), Sberbank (www.sberbank-ast.ru), RTS Index Agency (www.rts-tender.ru) and MMVB-Information Technologies (www.etp-micex.ru). Government agencies posting tenders can choose which platform they will use.

requires that suppliers are not in a state of bankruptcy, are not currently being punished by administrative law, do not have substantial unpaid taxes, and are not listed in the registry of suppliers who have committed violations of procurement rules during the last two years. Suppliers must also submit their security deposit. Finally, suppliers must prepare a formal application, consisting of two parts. The first part describes the good or service that they are offering to fulfill the procurement order; it can include product codes, sketches, and photographs to explain how the terms of the order will be met. The second part contains information on the supplier itself (name, address, etc.), licenses, accreditation, and past records of any other contracts fulfilled. Importantly, until the auction is concluded only the electronic trading platform has access to the second part of the application.

A procuring commission designated to oversee the auction receives and evaluates the first part of the application before the auction is held; no information on the identity of the supplier is revealed at this time. Applications to participate in auctions are denied if the supplier cannot pay the security deposit, their accreditation expires in three months or less, or their proposal is deemed not to comply with the requirements of the auction. In the event that only one supplier is approved to participate in the auction, the auction is declared "not held", the procuring commission receives the second part of the supplier's application, and a contract is drawn up with that supplier at the initial (maximum) price. This is a relatively common occurrence; in 1,344,825 cases, or 27 percent of the purchases, there is only one eligible participant. If there are no approved applicants, either because no suppliers apply or because all applicants are rejected, the purchase is cancelled. This occurs in 13 percent of auctions.

If more than one supplier is approved, the auction is held. All eligible suppliers are given a new and unique "participant number" to track their activities during the auction and protect their identity. At the specified time, all participants log in to the online platform and participate in a descending second-price auction. Qualifying bids must lower the current winning bid by discrete increments of between 0.5 and 5 percent of the initial (maximum) price. Information on the amount of a bid, the time entered, and the participant number is immediately made all 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 procuring commission receives and reviews the second part of the applications. These contain the identifying information for the auction participants, but do not allow for suppliers to be linked to the specific bids they submitted during the auction. During this second stage, the procuring commission checks the applications to make sure the suppliers' accreditations, names, tax ID numbers, registration, founding documents, and documents confirming participation in the tender 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.[25]

While the auction mechanism we study contains many built-in safeguards against corruption and collusion (such as anonymizing all documents and bidder identities during the auction) we cannot rule out that collusion between participants – or between participants and the bureaucrat and organization in charge of the auction – takes place. Note, however, that we are fairly confident that the goods purchased

[25] Declining to sign a contract after winning an electronic auction carries strict penalties for a supplier, including a three year ban from participating in future procurement processes.

are delivered.[26] That is, to the extent that corruption takes place in these auctions, it likely affects the final price paid and which firm wins the contract by manipulating how many and which firms participate in the auction. As such, the effects of collusion on the number and type of participants and the final price will be captured by our effectiveness estimates if individual bureaucrats and organizations are differentially corrupt (or differentially able to stop bidders from collusion). We thus believe that our estimates capture what a government should care about in the short term: the price paid for goods of a given type/quality that are actually delivered. From a longer term perspective, governments should also care about the allocative question of whether the "right" firms win contracts, but this is beyond the scope of our paper. If corruption channels demand to potentially less productive firms, the associated welfare consequences would not be captured by our estimates.

3.3 The role of bureaucrats and organizations in procurement

Public procurement purchases are made on behalf of a public sector entity that we will refer to as the *organization*. This organization requests that an item be procured, accepts delivery of the purchased item, uses the item, and pays for it. The organization may for example be a school, hospital or ministry at municipal, regional or federal level. As described above, the legal authority to make procurement purchases is decentralized to the level of these individual organizations. In order to make a purchase, the organization must pair with a procurement officer – we refer to these individuals as *bureaucrats* – to help organize and conduct the auction. Bureaucrats can either be "in-house" employees of the organization, or be employees of an external public agency whose bureaucrats conduct procurement auctions on behalf of multiple organizations. Such agencies can be organized by a given authority (for example an education or health ministry/department), at federal, regional, or municipal level, and its bureaucrats either required to conduct a specified subset (defined by the type of good to be purchased and/or reservation price of the contract) of the purchases of the organizations that fall under the authority's oversight.[27] Importantly, organizations are not able to choose whether to perform procurement in-house or through a procurement agency. Each federal authority, region or municipality sets rules dictating the use of either an in-house or external bureaucrat, depending on the size of the contract and the nature of the items being purchased that must be followed.

During the Soviet Union period, the civil service was centralized, and specialized bureaucrats were trained within a network of Higher Party Schools that was designed to supplement higher education in universities (Huskey, 2004). Afte the Russian Federation declared independence in 1991, the network of Higher Party schools collapsed under both fiscal and political pressure, leaving academies to fend for themselves in a new market for higher education. As a result, the education and labor market

[26] Less than 1 percent of the auctions in our sample suffered from "bad execution" (the supplier not carrying out his or her duties adequately, suppliers going bankrupt or disappearing, documents not being properly entered after the signing, a court order canceled the contract because of a dispute, etc.).

[27] Part of the motivation for allowing the creation of public agencies with bureaucrats who can handle purchases for multiple organizations 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 highly decentralized nature of procurement means that such joint purchases are rare, both because each participating organization is required to specify its own reservation price and technical documentation, and also because organizations need to initiate and coordinate joint purchases themselves.

for procurement bureaucrats is extremely decentralized.[28] Individuals interested in working in public procurement seek out educational and employment opportunities on an independent basis. In some instances, employers might request that candidates have completed professional training or obtained a certificate from an accredited institute that validates their knowledge of the electronic tender platforms. But interviews with experts and a review of recent procurement officer job vacancies posted to open online job boards revealed that the primary requirements are simply a legal education and knowledge of the existing laws (94-FZ, 44-FZ, and related normative acts) governing state tenders.[29] Moreover, in all cases we are aware of, the procurement bureaucrats are paid a flat salary, with no pay for performance element.

Since 2014, the division of labor between the organization and a potential external bureaucrat has been specified by law. The organization must submit all technical documentation, and choose and justify a reservation price. After this documentation is posted online, the organization and bureaucrat together designate a "procuring commission" consisting of at least five members to oversee the auction process, including assigning a chairman, deputy chairman, and secretary. These individuals must not have any connections to potential suppliers of the procurement order. The bureaucrat is on the committee, except in special circumstances. The organization also signs the contract once the winning bidder has been chosen. The external bureaucrat, with the help of the committee, is in charge of first stage review of applications, the auction itself, and second stage review of applications.[30] As far as we are aware, the same or a similar division of labor between the bureaucrat and his/her superiors in the organization applies when in-house bureaucrats are used, and also applies in purchased conducted before 2014. There is thus wide scope for both the bureaucrat and organization in charge to affect how the procurement process is conducted, and hence on the final outcomes.

3.4 Preferences for domestically-made goods

As part of reforms passed in 2005, the Russian government established a system to provide for special treatment of – "preferences" for – some types of firms when they participate in electronic auctions and open tenders. Three types of preferences were created: preferences for small and medium enterprises, preferences for organizations working with the disabled, and preferences for suppliers that make their goods and services in Russia. In this paper, we focus on the preferences for producers of Russian goods, who received a 15 percent preference from 2010-2015.

The preferences regime worked as follows. Each year from 2011 to 2014 a list of goods for which preferences for local producers applied was drawn up. The government order defining this list was passed

[28] The Russian government has not adopted a single approach to educating bureaucrats nor does it operate a centralized civil service administration to recruit, train, and assign public servants to postings around the country (Barabashev & Straussman, 2007). Instead, officials have encouraged market-based, institutional pluralism whereby regional academies, nongovernmental education institutes, and private companies compete in offering educational services to budding bureaucrats. Examples of such institutes offering trainings in the procurement sector include ArtAleks `http://artaleks.ru/`, the Granit Center `http://www.granit.ru/`, and the Higher School of Economics `https://igz.hse.ru/`.

[29] We include screenshots of recent online job ads in the Appendix.

[30] The one exception to this are "Kazennyie organizations", which can delegate all functions of the process to a centralized bureaucrat.

by the Ministry of Economic Development in May or June and remained in effect for the remaining part of the calendar year (through December 31st), after which the system of preferences ceased to operate until a new list had been created for the following year. The 2014 list was extended through December 31st, 2015; the first time a list had been in effect for more than a year. As such, preferences were never applied to procurements conducted the first period of each year from 2010-2014. The goods for which the preferences regime applied were identified by a unique product code, either from the OKDP (2011-2013) or from the OKPD (2014-2015), two official national classification systems used in Russia. Organizations filing procurement requests for any goods on this annual list were required to publicly inform participants that preferences for producers of Russian goods would be applied.[31] Goods where preferences were in place spanned numerous categories, including various types of food products, textile and furs, clocks, medical equipment, and automobiles. Importantly, the preferences regime did not ban foreign companies from participating in auctions and tenders. The country of origin of a good was defined as that country where the good was completely produced, or where it underwent significant reprocessing.

For the preferences to be applied, at least one application offering a foreign-made good and at least one application offering a Russian-made good had to have been submitted during the first stage of the auction process. If the firm that submitted the winning bid in the electronic auction had offered a foreign good in its application, then the contract it was offered to sign would be for 85 percent of its final bid. Therefore, the advantage given to local producers was that a winning firm supplying a foreign good would receive 15 percent less than their final winning bid. As always, declining to sign a contract after winning an electronic auction carried strict penalties for a supplier.

4 Data

4.1 Auction data

Our data on auction requests and final contracts comes from the Unified Register of Federal and Municipal Contracts located at http://zakupki.gov.ru/. This centralized site houses information on all public procurement in Russia from 2011 to the present-day. We collected data on the universe of electronic auction requests, review protocols, auction protocols, and contracts from January 1, 2011 until December 31, 2015. In all, we have information on 5,054,498 requests, though not all of these tenders are successfully concluded with a contract signed at the end of the process. Figure 1 maps our data onto the specific procurement procedures described in Section 3.

A great deal of previous work has faced the challenge of assigning products to categories so as to ensure that quality differences within categories are minimal. Broadly, three potential approaches have been pursued. First, much literature in housing and industrial organization thinks of products as bundles of underlying utility-generating attributes, collects data directly on the attributes assumed to be relevant, and uses hedonic regression analysis to estimate consumers' demand for attributes and/or sup-

[31] Significant penalties were incurred for those caught not applying preferences to an auction where the good being procured was on the list for that calendar year. Individual bureaucrats could be fined 15,000 rubles (∼USD 500), while organizations could be fined 50,000 rubles (∼ USD 1600).

pliers' costs of producing attributes (Griliches, 1971; Rosen, 1974; Epple, 1987). The second approach, prevalent in the trade literature, partitions the set of products into subsets within which products are assumed to be relatively homogeneous, using product classifications recorded by customs or other authorities (such as the Harmonized System codes) (Rauch, 1999; Schott, 2004; Broda & Weinstein, 2006), or more recently, international barcodes used by retailers (see, for example, Bronnenberg *et al.* (2015) and Atkin *et al.* (2015)). A third approach has been to restrict attention to products that are by nature especially homogeneous, such as cement and block-ice (see e.g. Syverson, 2004; Hortacsu & Syverson, 2007; Foster *et al.* , 2009). However, an issue with existing approaches to partitioning goods into comparable categories is that success is typically achieved at the cost of losing generality.[32] We are also in the common situation of text fields being the most detailed information available on the goods transacted in our data. We thus use text analysis methods from the machine learning literature to construct groups of homogeneous products. These methods are not yet widely used in economics, but we show how they can be used to construct consistent good classifications from good descriptions.[33]

Our method consists of four 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 10-digit 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. Fourth, we assign each cluster in the third step a more aggregate 6-digit HS product code in order to match to existing measures of product homogeneity for use in robustness exercises (Rauch, 1999; Khandelwal, 2010). Details are in Appendix C.

4.2 Pharmaceutical data

We also collect detailed data on procurement requests for pharmaceuticals, a sector where additional information on market price and goods' country of origin can be inferred using the brand name of the drug procured. Preferences applied to all pharmaceuticals goods each year from 2011-2015. The Russian government regulates the pharmaceutical market to ensure that certain drugs are available to prevent and treat certain illnesses, compelling manufacturers of these medicines to register in a List of Vital and Essential Medicinal Drugs (LVEMD). This list includes information on each drug's International Nonpro-

[32] Generality suffers both from restricting attention to very specific types of goods, and, in a methodological sense, from the fact that the existing methods that are more successful (at constructing homogeneous bins of goods) are often applicable only when unusual types of data are available.

[33] In using text analysis to classify categories to compare, we follow Hoberg & Phillips (2016). Their goal differs in that they classify *firm* similarity based on text listing the various goods firms produce, whereas we classify *good* similarity based on text describing each purchase. Their method also differs in that they pre-specify the desired number of clusters/categories. See also Gentzkow & Shapiro (2010); Hansen *et al.* (2014).

prietary Name (INN) (a globally recognized term to denote the chemical substance of the medicine)[34]; 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 the LVEMD using each drug's international brand name (international trademark name), active ingredient (international nonproprietary name - INN), dosage (mg, g, mkg), active units (IU), concentration (mg/ml, mkg/ml), volume (ml), and units (tablets, packages). This matching allows us to identify the exact drug and producer for each contract item.

Table 1 presents summary statistics on pharmaceutical auctions that were held from 2011 to 2015. Column 1 presents information on the sample of medicinal products that we were able to match to entries in the LVEMD (referred to below as the LVEMD sample). Column 2 includes all auctions in the pharmaceutical sector that had similar product codes to those in the LVEMD sample, but which we were unable to find an LVEMD entry for. This missingness could either be because the medicine being procured was not classified by the Russian government as "essential" or because sufficient information on dosage and quantity was not available in contract data. Table 1 indicates that the two samples are broadly similar. Though the number of medicinal auctions in our LVEMD sample is roughly one-third of overall number in relevant medicine categories, the LVEMD sample covers the vast majority of Russian regions and includes a large number of procuring organizations. Similarly, the proportions of procuring organizations in the LVEMD sample working at various administrative levels (federal, regional, or municipal) or across sectors (e.g. education, health, and internal affairs) are similar to those of the full sample (rows 6-11).

4.3 Supplier firm data

We also collected data on all suppliers that participate in any stage of the procurement process. The primary dataset on suppliers is *Ruslana*, which is collected by the Bureau Van Dijk (BVD) agency. Ruslana contains over four million Russian, Ukrainian, and Kazakh companies, covering a vast majority of registered companies that file financial information. All companies 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 private agencies such as BVD for dissemination to end-clients.

5 Individuals and Organizations as Sources of State Effectiveness

We begin our empirical analysis by focusing on purchases made under a policy regime that treats all suppliers equally. This allows us to focus on the role that the bureaucratic apparatus plays in shaping public sector output, while holding the policy environment fixed. As shown in Figure [histogram of prices], prices vary dramatically across purchases within good categories. In this section, we ask how much of this variation can be attributed to individual bureaucrats and organizations.

[34] 'Essential Medicines and Health Products.' World Health Organization. "http://www.who.int/medicines/services/inn/en/". Accessed October 8, 2015.

5.1 Empirical model

We model the final price paid in a procurement purchase as follows. Each item i destined for organization j is procured by a bureaucrat indexed by $b(i,j)$. The log price paid is

$$p_i = \mathbf{X}_i \boldsymbol{\beta} + \tilde{\alpha}_{b(i,j)} + \tilde{\psi}_j + \varepsilon_i \tag{12}$$

where $\mathbf{X}_i$ is a vector of item-level controls, including log quantity, good fixed effects[35], month fixed effects and interactions between 2-digit HS product categories, years, regions, and lot size[36]; $\tilde{\alpha}_{b(i,j)}$ is the bureaucrat effect, $\tilde{\psi}_j$ is the organization effect; and ε_{ia} is a residual. If bureaucrats are important drivers of prices achieved, then we expect $\mathsf{Var}\left(\tilde{\alpha}_{b(i,j)}\right) > 0$, and similarly, for organizations that drive prices, $\mathsf{Var}\left(\tilde{\psi}_j\right) > 0$.

5.2 Identification and estimation

Separate identification of the bureaucrat and organization effects is made possible by the fact that some bureaucrats make purchases with and for multiple organizations, and some organizations use multiple bureaucrats to make purchases. However, it is not possible to identify the effects of all the bureaucrats and organizations in the sample. Organizations are linked to each other by bureaucrats who make purchases for multiple organizations, allowing us to partition the N_b bureaucrats and N_j organizations into N_s mutually exclusive connected sets, each of which contains all the bureaucrats and organizations that can be linked by chains of bureaucrat "mobility". As shown by Abowd *et al.* (2002), 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 the $\tilde{\alpha}_{b(i,j)}$'s and $\tilde{\psi}_j$'s. Within each connected set, the bureaucrat and organization effects are identified only relative to each other and so must be normalized to enable interpretation. This also implies that comparisons across connected sets can only be made relative to the normalizations made in each connected set. That is, we will be able to identify $\mathsf{E}\left[\mathsf{Var}\left(\alpha_{b(i,j)}\right)|s\right]$ but not the unconditional variance, and similarly for the bureaucrat effects.

Faced with this issue, previous work has tended to restrict attention to the largest connected set (which in administrative datasets on the private sector tends to contain over 90 percent of the workers and firms), normalizing an arbitrary firm effect to 0, and estimating unconditional variances.[37] However, due to the decentralized nature of the Russian procurement system and lower worker mobility in the public sector, our data contains 32,127 connected sets. Nevertheless, the largest connected set contains 36,733 of the 116,436 organizations in the full sample. To assuage potential concerns about the representativeness of the largest connected set, we conduct our analysis in two samples. First we remove any bureaucrat-organization pair that only ever occurs together (as in this case it is impossible to distinguish bureaucrats from organizations, and similarly for bureaucrat-good pairs and organization-good

[35] Hereinafter we refer to the good categories constructed using the method described in Sub-section 4.1 as "goods".

[36] That is, we interact a more aggregate good-type indicator than the one we use to assign items to good categories with these other controls.

[37] An exception is Card *et al.* (2015) who study the largest male and female connected sets in Portuguese data, and who normalize the average effects of a subset of firms in each connected set to 0.

pairs. We also require that all bureaucrats and organizations purchase at least five items. We label the resulting sample – which contains all connected sets that fulfill these restrictions – the *analysis* sample. In a second approach we restrict attention to the largest connected set in the analysis sample. Table 2 compares the analysis and largest connected set samples to the full sample. All three are broadly similar in terms of the mean numbers of applicant and bidders, the sizes of the requests, as well as item-level characteristics, such as quantity, price, and price per unit.

In the analysis sample that contains multiple connected sets, a natural normalization of the bureaucrat and organization effects is to normalize each to have mean zero within each connected set and augment the model to include an intercept, $\gamma_{s(b,j)}$, specific to each connected set.[38] We rewrite the model in equation (12) as

$$p_i = \mathbf{X}_i\boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i \tag{13}$$

We use (13) to decompose the variation in prices into its constituent parts as follows

$$\begin{aligned}\mathsf{Var}(p_i) = \mathsf{Var}\left(\alpha_{b(i,j)}\right) + \mathsf{Var}\left(\psi_j\right) + \mathsf{Var}\left(\gamma_{s(b,j)}\right) + 2\mathsf{Cov}\left(\alpha_{b(i,j)}, \psi_j\right) + \mathsf{Var}\left(\mathbf{X}_i\boldsymbol{\beta}\right) \\ + 2\mathsf{Cov}\left(\alpha_{b(i,j)} + \psi_j, \gamma_{s(b,j)} + \mathbf{X}_i\boldsymbol{\beta}\right) + 2\mathsf{Cov}\left(\gamma_{s(b,j)}, \mathbf{X}_i\boldsymbol{\beta}\right) + \mathsf{Var}\left(\varepsilon_i\right)\end{aligned} \tag{14}$$

As shown in Appendix B, the effects in this augmented model are related to the underlying bureaucrat and organization effects through the equations

$$\alpha_b = \tilde{\alpha}_b - \overline{\alpha}_{s(b)} \tag{15}$$

$$\psi_j = \tilde{\psi}_j - \overline{\psi}_{s(j)} \tag{16}$$

$$\gamma_{s(b,j)} = \overline{\alpha}_{s(b,j)} + \overline{\psi}_{s(b,j)} \tag{17}$$

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 the connected set containing organization j. This allows us to relate the variances of our estimated bureaucrat and organization effects to their variances within and between connected sets using the law of total variance:

$$\begin{aligned}\mathsf{Var}\left(\tilde{\alpha}_b\right) &\equiv \mathsf{E}\left[\mathsf{Var}\left(\tilde{\alpha}_b|s(b)\right)\right] + \mathsf{Var}\left(\mathsf{E}\left[\tilde{\alpha}_b|s(b)\right]\right) \\ &= \mathsf{Var}\left(\alpha_b\right) + \mathsf{Var}\left(\mathsf{E}\left[\tilde{\alpha}_b|s(b)\right]\right) \geq \mathsf{Var}\left(\alpha_b\right) \qquad (18) \\ \mathsf{Var}\left(\tilde{\psi}_j\right) &\equiv \mathsf{E}\left[\mathsf{Var}\left(\tilde{\psi}_j|s(j)\right)\right] + \mathsf{Var}\left(\mathsf{E}\left[\tilde{\psi}_j|s(j)\right]\right) \\ &= \mathsf{Var}\left(\psi_j\right) + \mathsf{Var}\left(\mathsf{E}\left[\tilde{\psi}_j|s(j)\right]\right) \geq \mathsf{Var}\left(\psi_j\right) \qquad (19) \\ \mathsf{Var}\left(\tilde{\alpha}_b + \tilde{\psi}_j\right) &\equiv \mathsf{E}\left[\mathsf{Var}\left(\tilde{\alpha}_b + \tilde{\psi}_j|s(b,j)\right)\right] + \mathsf{Var}\left(\mathsf{E}\left[\tilde{\alpha}_b + \tilde{\psi}_j|s(b,j)\right]\right) \\ &= \mathsf{Var}\left(\alpha_b + \psi_j\right) + \mathsf{Var}\left(\gamma_{s(b,j)}\right) \qquad (20)\end{aligned}$$

Equations (18)–(20) show that consistent estimates of the variances of the bureaucrat and organization

[38]Note that this is actually more normalizations than are strictly necessary. We only require one linear restriction per connected set, while this normalization imposes two restrictions.

effects in (13) provide lower bounds on the variances of the true bureaucrat and organization effects in (12), respectively, and that we can construct the variance of the total effect of bureaucrats and organizations using our estimated bureaucrat and organization effects and our estimated connected set intercepts.

Our variance decomposition method uses movements of organizations between bureaucrats and between goods, and movements of bureaucrats between goods to identify how specific bureaucrats and organizations affect prices. Our identification therefore relies on these movements being orthogonal to the error term in equation (13). To illustrate the possible sources of endogenous mobility, we follow Card *et al.* (2013b) and write the error term as consisting of five random effects:

$$\varepsilon_i = \eta_{b(i,j)j} + \theta_{b(i,j)g} + \kappa_{jg} + \zeta_j + \xi_{b(i,j)} + \nu_i \tag{21}$$

where g indexes the good being purchased. $\eta_{b(i,j)j}$ is a bureaucrat-organization match-specific effect, and similarly $\theta_{b(i,j)g}$ and κ_{jg} are match effects for bureaucrat-good and organization-good matches. $\xi_{b(i,j)}$ and ν_i are unit-root drift terms for bureaucrats and organizations respectively. ν_i is a transitory error term.

$\eta_{b(i,j)j}$ represents price discounts (premia) that organization j achieves (suffers) when working with bureaucrat b relative to $\alpha_{b(i,j)} + \psi_j$. Such match effects could arise if organizations work especially well (or poorly) with specific bureaucrats. Similarly, it is possible that organizations and/or bureaucrats are especially good (or bad) at procuring specific types of goods, which would be captured by κ_{jg} and $\theta_{b(i,j)g}$ respectively. The unit root components reflect potential drift in the general ability of an organization or bureaucrat over time. Such drift could for example reflect the organization/bureaucrat learning how to achieve low prices, or potential bidders learning about the desirability of participating in auctions managed by a particular organization/bureaucrat.[39] The transitory term captures any remaining components of the error term.

Stacking the N items, we can write the model in matrix form as

$$\mathbf{p} = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\alpha} + \mathbf{J}\boldsymbol{\psi} + \mathbf{S}\boldsymbol{\gamma} + \boldsymbol{\varepsilon} \tag{22}$$

where B is the $N \times N_b$ design matrix indicating the bureaucrat conducting each purchase $[b_1, b_2, \ldots, b_{N_b}]$; J is the $N \times N_j$ design matrix indicating the organization purchasing each item $\left[j_1, j_2, \ldots, j_{N_j}\right]$; and S is the design matrix of connected set dummies $[s_1, s_2, \ldots, \ldots, s_{N_s}]$. $\mathbf{X}$ contains the good category fixed effects so that we can write $\mathbf{X}\boldsymbol{\beta} = \mathbf{G}\boldsymbol{\delta} + \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}$, where $\mathbf{G}$ is the $N \times N_g$ design matrix indicating the good category to which the item being purchased belongs.

Estimating (22) by OLS[40] will then identify the effects $\boldsymbol{\alpha}, \boldsymbol{\psi}$, and $\boldsymbol{\gamma}$ under the following assumptions:

$$\mathsf{E}\left[b_b'\boldsymbol{\varepsilon}\right] = 0 \;\forall b; \quad \mathsf{E}\left[j_j'\boldsymbol{\varepsilon}\right] = 0 \;\forall j; \quad \mathsf{E}\left[g_g'\boldsymbol{\varepsilon}\right] = 0 \;\forall i; \quad \mathsf{E}\left[\tilde{\mathbf{X}}'\boldsymbol{\varepsilon}\right] = 0 \tag{23}$$

[39] We assume that each of the three match effects has mean zero, and that the ζ and ξ components have mean zero but contain a unit root. General time trends in the data will be captured by the month effects in $\mathbf{X}$.

[40] As the dimensions of the matrices of fixed effects involved are large, rather than inverting a high-dimensional matrix, we follow the AKM literature and solve the OLS normal equations directly using the `lfe` package for `R` written by Gaure (2015).

which together implies that $\mathsf{E}\left[s_s'\varepsilon\right] = 0\ \forall s$. These orthogonality conditions allow for rich patterns of sorting of bureaucrats, organizations, and goods. For example, bureaucrats can move to the higher performing organizations over time, or effective bureaucrats can move systematically to high (or low) performing organizations, without violating (23). Similarly, especially effective bureaucrats and organizations can specialize in the purchase of certain goods. What (23) does rule out is systematic sorting based on unmodelled match effects between bureaucrats and organizations, bureaucrats and goods, or organizations and goods. Such forms of "endogenous mobility" are a priori unlikely in the institutional context of Russian public procurement (see Section 3). Nevertheless, we follow the existing literature, especially Card *et al.* (2013b), and explore the possibility as follows.

First, bias can arise if organizations choose bureaucrats to work with based on a match-specific price effect (as modeled by Mortensen & Pissarides (1994) and the related search literature in the context of private sector labor markets). Under the assumptions in (21), an organization that switches from using bureaucrat 1 to bureaucrat 2 can expect that the prices it pays will change by

$$\mathsf{E}\left[p|b=1\right] - \mathsf{E}\left[p|b=2\right] = \alpha_1 - \alpha_2 + \mathsf{E}\left[\varepsilon_i|b=1\right] - \mathsf{E}\left[\varepsilon_i|b=1\right] \tag{24}$$

$$\begin{aligned} &= \alpha_1 - \alpha_2 + \mathsf{E}\left[\eta_{1j}\right] - \mathsf{E}\left[\eta_{2j}\right] + \mathsf{E}\left[\theta_{1g}\right] - \mathsf{E}\left[\theta_{2g}\right] \\ &\quad + \mathsf{E}\left[\zeta_j|b=1\right] - \mathsf{E}\left[\zeta_j|b=2\right] + \mathsf{E}\left[\nu_i|b=1\right] - \mathsf{E}\left[\nu_i|b=2\right] \end{aligned} \tag{25}$$

If organization-bureaucrat match effects influence organizations' choice of bureaucrats, $\mathsf{E}\left[\eta_{1j}\right] \neq \mathsf{E}\left[\eta_{2j}\right]$. To test for this possibility, we construct an event study analysis tracking organizations that replace the bureaucrat they work with. We define an "employment spell" as a sequence of at least three purchases an organization-bureaucrat pair conduct together with less than 400 days between purchases. Wherever possible, we then match each employment spell with the earliest future spell involving the same organization but a different bureaucrat. This change of bureaucrats then constistutes an event. We classify the two bureaucrats involved in the event by assigning the average price they achieve in purchases they make with/for *other* organizations to the relevant quartile of the distribution of all bureaucrats' average prices during the quarter that the spell ends (for the earlier spell) or starts (for the later spell).

Figure 2 presents the results. The horizontal axis displays event time, i.e. the average prices achieved on occasions when the organization-bureaucrat pair made a purchase, with event time = 0 indexing the last day on which the "old" pair made a purchase, and event time = 1 indexing the first day on which the "new" pair made a purchase. On the vertical axis we display average prices paid, residualizing out month and good fixed effects.

Prices paid change sharply when the organization switches to a "lower price" or "higher price" bureaucrat, suggesting that bureaucrats do indeed affect prices paid. Further, there do not appear to be systematic price changes associated with switching between bureaucrats in the same quartile, and the price changes associated with switching from a high to a lower quartile bureaucrat and vice versa appear symmetric. For example, organizations switching from a bureaucrat in the first quartile of prices to a bureaurat in the fourth quartile experience an equal but opposite price change to organizations switching from the fourth to the first quartile. These last two points together are compelling evidence against the existence of strong sorting on match effects. If there was sorting on match effects, we would expect all switchers to experience price drops, and we would expect those moving from the first to the fourth quartile to experience a smaller price increase than organizations moving in the opposite direction.

Another concern is that organizations that become better (or worse) at procurement over time may systematically switch to a different type of bureaucrat, or vice versa for bureaucrats. In that case $\mathsf{E}\left[\zeta_j|b=1\right] \neq \mathsf{E}\left[\zeta_j|b=2\right]$ in the example in (24). This could occur if for example organizations with deteriorating performance have their procurement officer reassigned by the central government so that $\mathsf{E}\left[\zeta_j|b=2\right] > \mathsf{E}\left[\zeta_j|b=1\right]$. However, we do not see any systematic time trend in the trajectories of switchers in Figure 2, suggesting that there is no strong correlation between drift and switching.

It is also possible that fluctuations in the idiosyncratic error term ν_i are correlated with organizations switching bureaucrats, if for example an unexpectedly high price leads the organization to replace the bureaucrat. This would lead us to overstate the difference in the bureaucrat effects since $\mathsf{E}\left[\nu_i|b=1\right] > \mathsf{E}\left[\nu_i|b=2\right]$. However, Figure 2 shows no systematic "Ashenfelter dips" just before a bureaucrat switches, suggesting that the transitory error ν_i is not correlated with switching.

Bias could also arise if bureaucrats or organizations specialize in goods for which they are better at achieving low prices. In the example in (24), it could be that bureaucrat 1 is more specialized in the goods the organization typically purchases than bureaucrat 2 is, in which case we would underestimate the difference in the bureaucrat effects since then $\mathsf{E}\left[\theta_{1g}\right] < \mathsf{E}\left[\theta_{2g}\right]$. To test for this possibility, we construct event study figures for organizations switching between goods and bureaucrats switching between goods by following a procedure analogous to that for Figure 2. The results are presented in figures 3 and 4. These figures show the same patterns as in Figure 2. In addition to alleviating any concerns due to the possibility of biases arising from unmodeled match effects between organizations and goods or bureaucrats and goods, figures 3 and 4 thus help rule out strong correlation between drift in the organization and/or bureaucrat effects or the transitory error and organizations/bureaucrats switching goods. Taken together, figures 2–4 provide strong support in favor of our identifying assumptions and thus our interpretation of the estimates as individual and organizational sources of prices paid by the government.

A separate set of estimation issues arise from finite sample biases. As is well known from the panel data literature, consistency of a single set of estimated fixed effects requires that the number of observations on each group, rather than simply the total sample size, tends to infinity (Scott, 1948; Lancaster, 2000). In our case, this *incidental parameters* problem is expected to lead the estimated bureaucrat and organization fixed effects to be overdispersed, biasing us towards finding an impact of the procurers even if there is none. In the case of two sets of fixed effects, the problem may be compounded by *limited mobility bias*, i.e. that the estimated covariance between the two sets of fixed effects is negatively biased when the network of workers and firms (here: bureaucrats and organizations) features few movers (Andrews *et al.* , 2008).[41]

[41] We are estimating models with three sets of high-dimensional fixed effects (for organizations, bureaucrats, and goods). (The models also contain month dummies to control for common time trends, but there are few enough of these month effects that "month-connectedness" is not an issue). To our knowledge, identification results for models with more than two sets of

We address the possibility of sampling error biases in three ways. First, when calculating standard errors for our variance decomposition, we use a randomization inference approach rather than analytical standard errors so that we can take into account the patterns of correlation in the residuals. 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 has limitations[42], but makes randomization inference computationally feasible with our large datasets.

Our second method for dealing with sampling error is non-parametric approach, as in Finkelstein *et al.* (2016); Silver (2016). We take the observations for each bureaucrat-organization pair and randomly split them into two samples. We then estimate equation (13) 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. Each coefficient is estimated with error due to the incidental parameters problem, and limited mobility bias, e.g. $\hat{\alpha}_b^k = \alpha_b + \mu_{\alpha,b}^k$ etc. However, the errors in the two estimates should be uncorrelated ($\mathsf{Cov}\left(\mu_{\alpha,b}^1, \mu_{\alpha,b}^2\right) = 0$), so we can create split-sample estimates of the relevant variance terms as follows:

$$\widehat{\mathsf{Var}}^{SS}(\alpha_b) = \mathsf{Cov}\left(\hat{\alpha}_b^1, \hat{\alpha}_b^2\right) \qquad \widehat{\mathsf{Var}}^{SS}(\psi_j) = \mathsf{Cov}\left(\hat{\psi}_j^1, \hat{\psi}_j^2\right)$$
$$\widehat{\mathsf{Var}}^{SS}(\gamma_s) = \mathsf{Cov}\left(\hat{\gamma}_s^1, \hat{\gamma}_s^2\right) \qquad \widehat{\mathsf{Var}}^{SS}(\alpha_b + \psi_j) = \mathsf{Cov}\left(\hat{\alpha}_b^1 + \hat{\psi}_j^1, \hat{\alpha}_b^2 + \hat{\psi}_j^2\right)$$

Finally, we take a more parametric approach and estimate the variance components directly and use these to "shrink" our fixed effect estimates, akin to Kane & Staiger (2008); Chetty *et al.* (2014); Chetty & Hendren (2015). 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 for bureaucrats and organizations respectively. The variance of our estimated bureaucrat effects is $\mathsf{Var}(\hat{\alpha}) = \sigma_\alpha^2 + \sigma_\mu^2$ and the variance of our estimated organization effects is $\mathsf{Var}\left(\hat{\psi}\right) = \sigma_\psi^2 + \sigma_\omega^2$.

Our permutation method to calculate standard errors described above yields estimates of the variance of the sampling error for each bureaucrat and organization effect, s_b^2 and s_j^2. We thus estimate the signal variance of the bureaucrat effects as $\hat{\sigma}_\alpha^2 = \mathsf{Var}(\hat{\alpha}) - \mathsf{E}_b\left[s_b^2\right]$, where expectations are taken across bureaucrats and with weights $1/s_b^2$. Similarly, we estimate the signal variance of the organization effects as $\hat{\sigma}_\psi^2 = \mathsf{Var}\left(\hat{\psi}\right) - \mathsf{E}_j\left[s_j^2\right]$.

With these estimated variances in hand, we can form the linear predictor of the bureaucrat and organization effects that minimizes the mean-squared error of the predictions. Formally, we find

$$\lambda_b = \arg\min_{\tilde{\lambda}} \mathsf{E}\left[\alpha_b - \tilde{\lambda}\hat{\alpha}_b\right] = \frac{\sigma_\alpha^2}{\sigma_\alpha^2 + \sigma_{\mu_b}^2}$$

fixed effects are not yet available (Gaure, 2013), and providing such results is beyond the scope of this paper.

[42]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 **X**, but this drastically speeds up computation.

and our shrinkage estimators replace these terms with their sample analogues:

$$\hat{\alpha}_b^{Sh} = \frac{\hat{\sigma}_\alpha^2}{\hat{\sigma}_\alpha^2 + s_b^2}\hat{\alpha}_b \qquad \hat{\psi}_j^{Sh} = \frac{\hat{\sigma}_\psi^2}{\hat{\sigma}_\psi^2 + s_j^2}\hat{\psi}_j$$

5.3 Results

Table 3 implements the variance decomposition in equation (14) in the analysis sample. The first column shows estimates of the variances using the fixed effects from the estimation of equation (13), while the second column shows estimates from using our split-sample approach. The standard deviations of the bureaucrat and organization effects are large, at 1.588 and 1.607, respectively in the first column. However, they are negatively correlated so that their joint effect has a standard deviation of 1.122, or 1.273 if we add in the connected set effects to estimate the total effect of bureaucrats and organizations both within and across connected sets. Comparing these to the standard deviation of log prices, 3.198, we see that bureaucrats and organizations jointly explain over a third of the standard deviation of log prices. The split-sample estimates in the second column are remarkably similar; the standard deviation of the joint effect is 1.205, only slightly smaller than the corresponding 1.273 in column 1.

In Table 4 we repeat the variance decomposition using only the largest connected set. The results are again remarkably similar. The standard deviation of log prices is 3.206, of which 1.171 (using fixed effects), or 1.355 (using the split-sample method) can be attributed to bureaucrats and organizations. The standard errors estimated using the permutation method described above indicate that these results are strongly significant. In the final column we report the results of the shrinkage procedure, which gives slightly smaller estimated standard deviations than the split-sample approach, but only modestly so. The standard deviation of bureaucrat effects is 1.126, and that of the organization effects is 1.182.

5.4 Robustness: log-linearity

The model we have estimated assumes that the final price 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, strong match effects would lead the residuals to be large in the top bureaucrat and organization deciles, and small in the bottom deciles. Figure **??** shows a heat map of residuals for the analysis sample, while Figure **??** shows it for the largest connected set. In both cases we see no clear patterns in the residuals.

As a further test of our log linear model of prices, we reestimate equation (13) but include fixed effects for each bureaucrat-organization pair, allowing for arbitrary patterns of complementarity between bureaucrats and organizations. If there are indeed strong match effects between bureaucrats and organizations that we are omitting, then we expect the pair effect model to fit significantly better than our baseline model. Table 5 compares results for the two models. Strikingly, the pair effect model does not fit the data much better than our baseline model. The RMSE of the residuals goes down by 2.8 percent, the adjusted R^2 improves by 1.5 percent, and the pair effects have a much smaller variance than the main effects in tables 3 and 4. Overall, we do not find evidence supporting a rejection of our log-linearity assumption.

While the AKM model was developed to study firms and workers in the private sector, these results suggest that it can be fruitfully applied to performance data from the public sector. Our results show that individual bureaucrats and organizations significantly affect prices paid, contrary to the "mechanistic" view of the bureaucracy taken by much of the existing literature. Moreover, these effects are large, the bureaucrat and organization effects account for upwards of a third of the standard deviation in log prices.

5.5 Robustness: good homogeneity

A concern with our measure of state effectiveness is the possibility that differences in unit prices reflect differences in the quality of the goods being purchased, in addition to the differences in prices per quality-adjusted unit we attempt to restrict attention to. If this is the case, procurers who pay more on average for the goods they purchase may simply be purchasing higher quality goods, and not necessarily be performing worse than those paying less. To address this concern, we conduct two robustness exercises.

First, we show that our results are robust to restricting the sample to goods that are more homogenous. We split the sample into quintiles of good homogeneity, using the measure of the scope for product differentiation in Khandelwal (2010). We then reestimate (13) on successive subsamples adding in less and less homogenous goods. Table 6 shows the results. As we move from left to right, we add in less and less homogeneous goods, until the final column replicates our results from the full sample. The results are virtually unchanged when we restrict to increasingly homogeneous samples of goods. This reassures us that our results are not driven by unaccounted for good heterogeneity.

In our second check of the good homogeneity assumption, we restrict the sample to a type goods that is by nature very homogeneous — medicines (Bronnenberg *et al.* , 2015). The share of the variance in prices expained by bureaucrats and organizations is similar to in our full sample. Moreover, using barcode-level good controls rather than our text-based method delivers very similar results.

5.6 Correlates of individual and organizational level state effectiveness

What do effective bureaucrats and organizations do differently? In this section we relate variation in observable measures of individual bureaucrat and organization behavior to variation in their estimated fixed effects, α_b and ψ_j. As discussed in Section 4, our data contains detailed information on how a given procurement process was conducted and its intermediate outcome, from the intial request document, through the auction itself, to the final contract signed with the supplier. Since we have many observables for each purchase, we use regularization techniques to select the variables that are most predictive of the bureaucrat and organization effects. We run a LASSO regression of the bureaucrat and organization effects on all procurement process-related variables in our data to select the most important observables

(**?**), and then run bivariate and multivariate regressions using only the variables selected by the LASSO procedure.[43]

Figures 7 and 8 show the results. The left panel of each figure shows regression coefficients from a series of bivariate regressions of the bureaucrat (in Figure 7) or organization (in Figure 8) effects on each of the observables alone. The right panel shows the coefficients from the multivariate regression of the bureacrat or organization effects on the post-LASSO selected variables. To facilitate comparison, all variables are standardized to have unit standard deviation so that the coefficients can be interpreted as the effect in standard deviations of the bureaucrat/organization effects of a one standard deviation change in the measure of procurement behavior.

Three key findings emerge from Figure 7. First, successful bureaucrats are those who achieve high supplier participation rates. The number of applicants and the number of bidders are strongly associated with better bureaucrat performance. Second, successful bureaucrats make the auctions accessible by setting relatively low guarantees (a bond that participants have to post in order to participate in the auction), and by not revising the request document often (for example to correct errors that might deter potential bidders). Third, the qualification stage is more important than the auction and contracting stages. Six of the 13 LASSO-selected variables relate to the qualification stage, and these have bigger coefficients in the multivariate regression.

Turning to the organizations, the findings from Figure 8 echo those of Figure 7 with some differences. For organizations, we observe geographic location, level of government, and type of government activity. These are important predictors of performance. Organizations that are further from their regional capital, and organizations at regional and municipal (as opposed to federal) level achieve lower prices. On measures of behavior in the procurement process, the heterogeneity across effective and ineffective organizations is otherwise similar to that for bureaucrats.

Overall, 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, consistent with the predictions of our model in Section 2. This suggests that efforts to improve bureaucratic effectiveness should be targeted at recruiting bureaucrats that are comparatively good at reducing the costs imposed on potential government suppliers (or, similarly, reassigning procurement responsibilities across organizations to those that reduce such costs), or helping existing bureaucrats and organizations perform better in this dimension.

6 Individual and Organizational Sources of Heterogeneous Policy Effects: the Case of Bid Preferences

6.1 Difference-in-differences results

In Section 5 we held constant the policy environment. We varied the bureaucrat and organization in charge of procurement, exploiting the thousands of quasi-experiments created by the "movement" of

[43] More precisely, we use the variables with nonzero coefficients in the LASSO regression with the regularization penalty λ that minimizes the mean squared error in K-fold cross-validation.

organizations across bureaucrats, and vice versa, to estimate the effect of individual procurers. In this section we instead hold constant the bureaucrat and organization in charge of procurement and vary whether a particular procurement policy applies. We do this for two reasons. First, "shocking" the policy environment and tracing out how the effects depend on the procurers in charge represents an additional way to test if the heterogeneity in effectiveness estimated in Section 5 is due to differences in how effective and less effective bureaucrats and organizations administer policies, as we argue. Second, the ultimate goal of this paper is to determine if there are policy design implications of micro level sources of state effectiveness. To do so, it is useful to exploit a policy change.

We first pursue a difference-in-differences strategy focused on the policy itself. The "buy local" policy is "turned on" each year in the spring, and then expires at the end of the year. Preferences apply to some goods and not others, and the list of goods varies year-on-year, albeit only moderately so. Moreover, for the policy to apply, there must be a minimum of one bidder in the auction offering a domestic (Russian-produced) good, and a minimum of one bidder offering a foreign-made good.

Procurers signal to potential bidders that the preference policy applies to the items in a given request by checking a box in the electronic submission form. The preferences information is then prominently displayed on the first page of the electronic auction request. Although significant penalties apply for procurers failing to apply the preference policy correctly, procurers sometimes fail to check the box. There are thus two scenarios to consider: (1) where procurers correctly indicate that preferences apply to a request, and (2) where procurers fail to indicate on a request that preferences apply.[44] We come back to (2) below.

In the case of (1), we estimate the average treatment effect (ATE) of the preferences policy, where the treatment is a dummy variable for whether the procurers explicitly applied the policy to a request, correctly or incorrectly. That equation takes the following form:

$$y_{igt} = \alpha + \beta \text{Preferenced}_g + \gamma \ln(\text{quantity}) + \mu_g + \lambda_t + \varepsilon_{igt} \tag{26}$$

where y_{igt} is outcome y (log Price) in item i for good g in month t, Preferenced_g is a dummy for whether the procurers indicate that preferences apply to the good, ln(quantity) is the logarithm of the standardized quantity of goods being procured, μ_c and λ_t are good and month fixed effects, and ε_{igt} is an error term. We two-way cluster standard errors by month and good.

Table 7 shows the results of estimating equation 26 using the full sample, while Table 8 shows those using the largest connected set. Column (1) in each table shows that the average effect on the (log) price, controlling for month and good fixed effects and standardized quantity, is a precisely estimated zero. This is despite the fact that preferences are thought to reduce the number of firms that choose to bid on the contract.

As tables 7 and 8 show, the preferences policy does achieve the goal of procuring more Russian-made goods. This finding is noteworthy; since shifting demand towards domestic producers comes at no direct cost to the government, the "buy local" policy might be seen as successful industrial policy

[44]In our data, we find that the third scenario of bureaucrats applying the preference policy to requests for goods not on decreed list for that year is exceedingly rare, accounting for 0.09% of all requests where the policy is applied.

from the government's perspective.[45] The finding also contrasts with the results of studies of similar preference policies in the U.S. using more structural empirical methods (Marion, 2007; Krasnokutskaya & Seim, 2011).

We hypothesize that the effectiveness of the bureaucrat and organization that run the auction may matter for how preferences affect prices. To test this, we pursue a triple-difference approach, interacting the policy treatment with the estimated bureaucrat and organization effects from Section 5. Columns (2), (3) and (4) of both Table 7 (the analysis sample) and Table 8 (the largest connected set) show the results using the AKM estimates. We alternately include the full interactions with bureaucrat effectiveness, organization effectiveness, and both effectiveness measures. We also include a 'connected set' fixed effect in order to control for intercepts specific to each connected set.

Several findings emerge from the tables. First, the results clearly show that the zero price effect on average combined a price *increase* among effective bureaucrats and organizations and a price *decrease* among ineffective procurers. The patterns for bureaucrat effectiveness are very similar to those for organization effectiveness. Although the effect on the triple interaction with organization effectiveness is less precisely estimated, the introduction of the "buy local" policy appears to result in convergence of "good" and "bad" bureaucrats' performance, as well as good and bad organizations' performance. Under the policy, effective bureaucrats are less able to secure lower prices for goods, while ineffective bureaucrats do better at decreasing the cost the government pays for the same goods.

In the case of scnario (2) from above, where there are concerns about compliance, we adopt an intent-to-treat (ITT) approach that turns on the treatment indicator if the goods procured under the request should have had preferences policy apply, irrespective of whether the procurers correctly checked the box or not. Each request can consist of multiple goods. To be conservativy, we create a dummy variable indicating intent-to-treat that takes a value of 1 if *any* of the goods in a request should have had preferences apply. The estimating equation is thus the same as in (26), except that the way we define Preferenced_g differs.

We present the results of the intent-to-treat approach in Table 13 (for the analysis sample) and Table 14 (for the largest connected set). The findings differ from the average treatment effects discussed above in one important respect. The preference policy has a strongly negative ITT effect on prices achieved. Requests that technically fall under the provisions of the preferences regime see prices for each item drop by approximately 18 percent. However, when the interactions with bureaucrat and organization effectiveness are introduced in columns (2), (3) and (4), the heterogeneity results are broadly similar to the ATEs. Under the preferences policy, ineffective bureaucrats and organizations see a price decrease, while effective procurers pay higher prices under when the policy applies.

Overall our findings in this section provide some of the first direct evidence that policies should be tailored to the effectiveness of those who will implement the policies. The performance of highly effective bureaucrats and organizations in Russia are hurt by distorting competition, similarly to what has been found for preferences policies in the United States (Marion, 2007; Krasnokutskaya & Seim, 2011). By contrast, ineffective bureaucrats and organizations perform better when importing suppliers

[45] Although note that pinning down the welfare consequences of channeling demand to potentially less productive firms is beyond the scope of this paper.

are handicapped. Through the lens of the model in Section 2, our results indicate that this is because the additional entry of suppliers of local goods induced by the policy more than offsets the additional entry costs that ineffective procurers impose on potential bidders, leading to an overall increase in entry and a decrease in prices.

7 Conclusion

In this paper we have presented evidence that both the individuals and the organizations tasked with implementing policy are significant sources of state effectiveness. Bureaucrats and public sector organizations together account for over one third of the variation in quality-adjusted unit prices paid by the Russian government. Consistent with our simple endogenous entry auction model, effective public procurers engage in practices that lower entry costs for potential suppliers. Such practices matter not only in a constant policy environment, but also for the impact of policy changes. Studying the impact of a "buy local" policy that applies bidding preferences to bids from suppliers of domestically manufactured goods, we find that the induced increase in entry by domestic suppliers outweighs the tilting of the playing field against importers for ineffective procurers, who have low baseline entry rates. The opposite is true for effective bureaucrats and organizations, as our conceptual framework predicts.

These findings have important implications for policy design and future research on state effectiveness. First, the degree of heterogeneity in effectiveness *within* a given state enterprise we document implies that there are large returns to the state of employing more bureaucrats at the high end of the observed performance range, and of improving organization-wide characteristics such as management or "organizational culture". An important question for future research is whether individual and organizational effectiveness can best be improved by hiring other types of individuals (or reshuffling tasks across organizations) or by training existing bureaucrats and organizations better. Another is how the impact of more versus less individuals and organizations compares to the impact of more versus less effectives *policies*. This latter question is key for a central government trying to allocate its budget and attention so as to maximize state effectiveness.

A second implication is methodological. Our findings imply that in order to "extrapolate" an average treatment effect of a public policy estimated in one setting to another setting, knowledge of and making use of the difference in local effectiveness across the two settings is essential.

Finally, our findings imply that policies that are suboptimal when state effectiveness is high may be second-best optimal when state effectiveness is low. Policies should thus be designed with the effectiveness of the individuals and organizations that will implement the policies in mind. This finding is especially important for policymaking for and in developing countries. In the past international organizations have, for example, often recommended the use of the same policies around the world. Of course, achieving "globally" optimal policy outcomes likely requires *both* maximizing the effectiveness of the bureaucratic apparatus and choosing policies that are designed for implementation by effective individuals and organizations.

References

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

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

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

ALVAREZ, JORGE, ENGBOM, NIKLAS, & MOSER, CHRISTIAN. 2016. *Firms and the Decline of Earnings Inequality in Brazil*. Working Paper.

ANDREWS, MARTYN J., SCHANK, T., & UPWARD, R. 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.

ANDREYANOV, PASHA, DAVIDSON, ALEC, & KOROVKIN, VASILY. 2016. *Corruption vs Collusion: Evidence from Russian Procurement Auctions*. mimeo: UCLA.

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

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

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

ATKIN, DAVID, FABER, BENJAMIN, & GONZALEZ-NAVARRO, MARCO. 2015. *Retail Globalization and Household Welfare: Evidence from Mexico*. forthcoming, Journal of Political Economy.

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.

BANDIERA, ORIANA, GUISO, LUIGI, PRAT, ANDREA, & SADUN, RAFFAELLA. 2011. *What Do CEOs Do?* CEPR Working Paper No. 8235.

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

BERTRAND, M., & SCHOAR, A. 2003. Managing with style: The effect of managers on firm policies. *The 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. 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 forthcoming*, **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.

BLADER, STEVE, GARTENBERG, CLAUDINE, & PRAT, ANDREA. 2016. *The Contingent Effect of Management Practices*. Working Paper.

BLOOM, NICHOLAS, LEMOS, RENATA, SADUN, RAFFAELLA, & REENEN, JOHN VAN. 2015a. *Does management matter in healthcare*. mimeo, Stanford.

BLOOM, NICHOLAS, LEMOS, RENATA, SADUN, RAFFAELLA, & REENEN, JOHN VAN. 2015b. Does management matter in schools. *Economic Journal*.

BOLD, TESSA, KIMENYI, MWANGI, MWABU, GERMANO, NG'ANG'A, ALICE, & SANDEFUR, JUSTIN. 2013. *Scaling Up What Works: Experimental Evidence on External Validity in Kenyan Education*. mimeo, Center for Global Development.

BRANNMAN, L., & FROEB, L. 2000. Mergers, Cartels, Set-Asides and Bidding Preferences in Asymmetric Second-price Auctions. *Review of Economics and Statistics*, 283–290.

BRODA, CHRISTIAN, & WEINSTEIN, DAVID E. 2006. Globalization and the Gains from Variety. *Quarterly Journal of Economics*, **121**, 541–585.

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

BULOW, JEREMY, & KLEMPERER, PAUL. 1996. Auctions Versus Negotiations. *American Economic Review*, **86**, 180–194.

BURGESS, R., HANSEN, M., OLKEN, B., POTAPOV, P., & SIEBER, S. 2012. The political economy of deforestation in the tropics. *The Quarterly Journal of Economics*, **127**, 1707–1754.

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

CARD, DAVID, DEVICIENTI, FRANCESCO, & MAIDA, AGATA. 2013a. Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data. *Review of Economic Studies*.

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

CARD, DAVID, CARDOSO, ANA RUTE, & KLINE, PATRICK. 2015. *Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women*. Forthcoming, 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.

CHETTY, RAJ, & HENDREN, NATHANIEL. 2015. *The Impacts of Neighborhoods on Intergenerational Mobility: Childhood Exposure Effects and County-Level Estimates*. Mimeo: Harvard University.

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

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

COLLIER, P. 2014. *Africa: New Opportunities, Old Impediments*. Wincott Lecture at the International Growth Centre.

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.

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. 2016. *Financial Incentives as Signals: Experimental Evidence from the Recruitment of Village Promoters in Uganda*. mimeo Northwestern University.

DUFLO, E., GREENSTONE, M., PANDE, R., & RYAN, N. 2013. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics*, **128**, 1499–1545.

DUFLO, ESTHER, GREENSTONE, MICHAEL, PANDE, ROHINI, & RYAN, NICHOLAS. 2014. *The Value of Regulatory Discretion: Estimates from Environmental Inspections in India*. NBER Working Paper.

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

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

ESTACHE, A., & WREN-LEWIS, L. 2009. Toward a theory of regulation for developing countries: Following Jean-Jaques Laffont's lead. *Journal of Economic Literature*, **47**(3), 729–770.

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

FINAN, F., OLKEN, B.A., & PANDE, R. forthcoming. The Personnel Economics of the State. *Handbook of Field Experiments*.

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

FLAMBARD, V., & PERRIGNE, I. 2006. Asymmetry in Procurement Auctions: Evidence from Snow Removal Contracts. *Economic Journal*, 1014–1036.

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.

GAURE, SIMEN. 2015. *lfe: Linear Group Fixed Effects*. R package version 2.2-1699.

GECHTER, MICHAEL. 2016. *Generalizing the Results from Social Experiments: Theory and Evidence from Mexico and India*. working paper, Penn State.

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.

GOLDSCHMIDT, DEBORAH, & SCHMIEDER, JOHANNES. 2015. *The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure*. NBER Working Paper No. 21366.

GREENSTONE, MICHAEL, & JACK, B. KELSEY. 2015. *Envirodevonomics: A Research Agenda for a Young Field*. forthcoming, Journal of Economic Literature.

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

GULZAR, S., & PASQUALE, B. forthcoming. Politicians, Bureaucrats, and Development: Evidence from India. *American Political Science Review*.

HANNA, REMA, & WANG, SHING-YI. 2015. *Dishonesty and Selection into Public Service: Evidence from India*. mimeo, Harvard.

HANSEN, STEPHEN, MCMAHON, MICHAEL, & PRAT, ANDREA. 2014. *Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach*. Mimeo: Columbia University.

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*.

HUSKEY, EUGENE. 2004. From higher party schools to academies of state service: the marketization of bureaucratic training in Russia. *Slavic Review*, **63**(2), 325–348.

JIA, R. 2014. *Pollution for promotion*. Working Paper.

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

KANE, T. J., & STAIGER, D. 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. *The Quarterly Journal of Economics*.

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

KRASNOKUTSKAYA, E., & SEIM, K. 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.

LEVIN, DAN, & SMITH, JAMES L. 1994. Equilibrium in Auctions with Entry. *The American Economic Review*, **84**(3), 585–599.

LEWIS-FAUPEL, SEAN, NEGGERS, YUSUF, OLKEN, BENJAMIN, & PANDE, ROHINI. 2015. *Can Electronic Procurement Improve Infrastructure Provision? Evidence from Public Works in India and Indonesia*. forthcoming American Economic Journal: Economic Policy.

LI, TONG, & ZHANG, BINGYU. 2015. Affiliation and Entry in First-Price Auctions with Heterogeneous Bidders: An Analysis of Merger Effects. *American Economic Journal: Microeconomics*, **7**(2), 188–214.

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

MCHENRY, WILLIAM K., & PRYAMONOSOV, DENIS. 2010. Emerging Electronic Procurement in Russia's Regional Governments. *Journal of Public Procurement*, **10**(2), 211–=246.

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.

MENEZES, FLAVIO M., & MONTEIRO, PAULO K. 2000. Auctions with endogenous participation. *Review of Economic Design*, **5**(1), 71–89.

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

MILGROM, PAUL, & SEGAL, ILYA. 2002. Envelope Theorems for Arbitrary Choice Sets. *Econometrica*, **70**, 583–601.

MORTENSEN, DALE T., & PISSARIDES, CHRISTOPHER A. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *The Review of Economic Studies*, **61**(3), 397–415.

RASUL, IMRAN, & ROGGER, DANIEL. forthcoming. 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.

ROBERTS, JAMES W., & SWEETING, ANDREW. 2015. *Bailouts and the Preservation of Competition*. National Bureau of Economic Research.

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

ROKKANEN, MIIKKA. 2016. *Exam Schools, Ability, and the Effects of Affirmative Action: Latent Factor Extrapolation in the Regression Discontinuity Design*. working paper, Columbia.

ROSEN, SHERWIN. 1974. Hedonic Prices and Implicit Markets: Product Dfiferentiation 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.

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

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.

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

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

VIVALT, EVA. 2015. Heterogeneous Treatment Effects in Impact Evaluations. *American Economic Review, Papers and Proceedings*, **105**, 467–470.

VIVALT, EVA. 2016. *How Much Can We Impact Evaluations Inform Policy Decisions?* Mimeo: Stanford.

YAKOVLEV, ANDREI, DEMIDOVA, OLGA, & BALAEVA, OLGA. 2010. Empirical Analysis Of Suppliers' Non-Performance Risks in Execution of Public Procurement Contracts in Russia. *Financi*, 6–20.

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

FIGURE 1: PROCESS FLOW-CHART

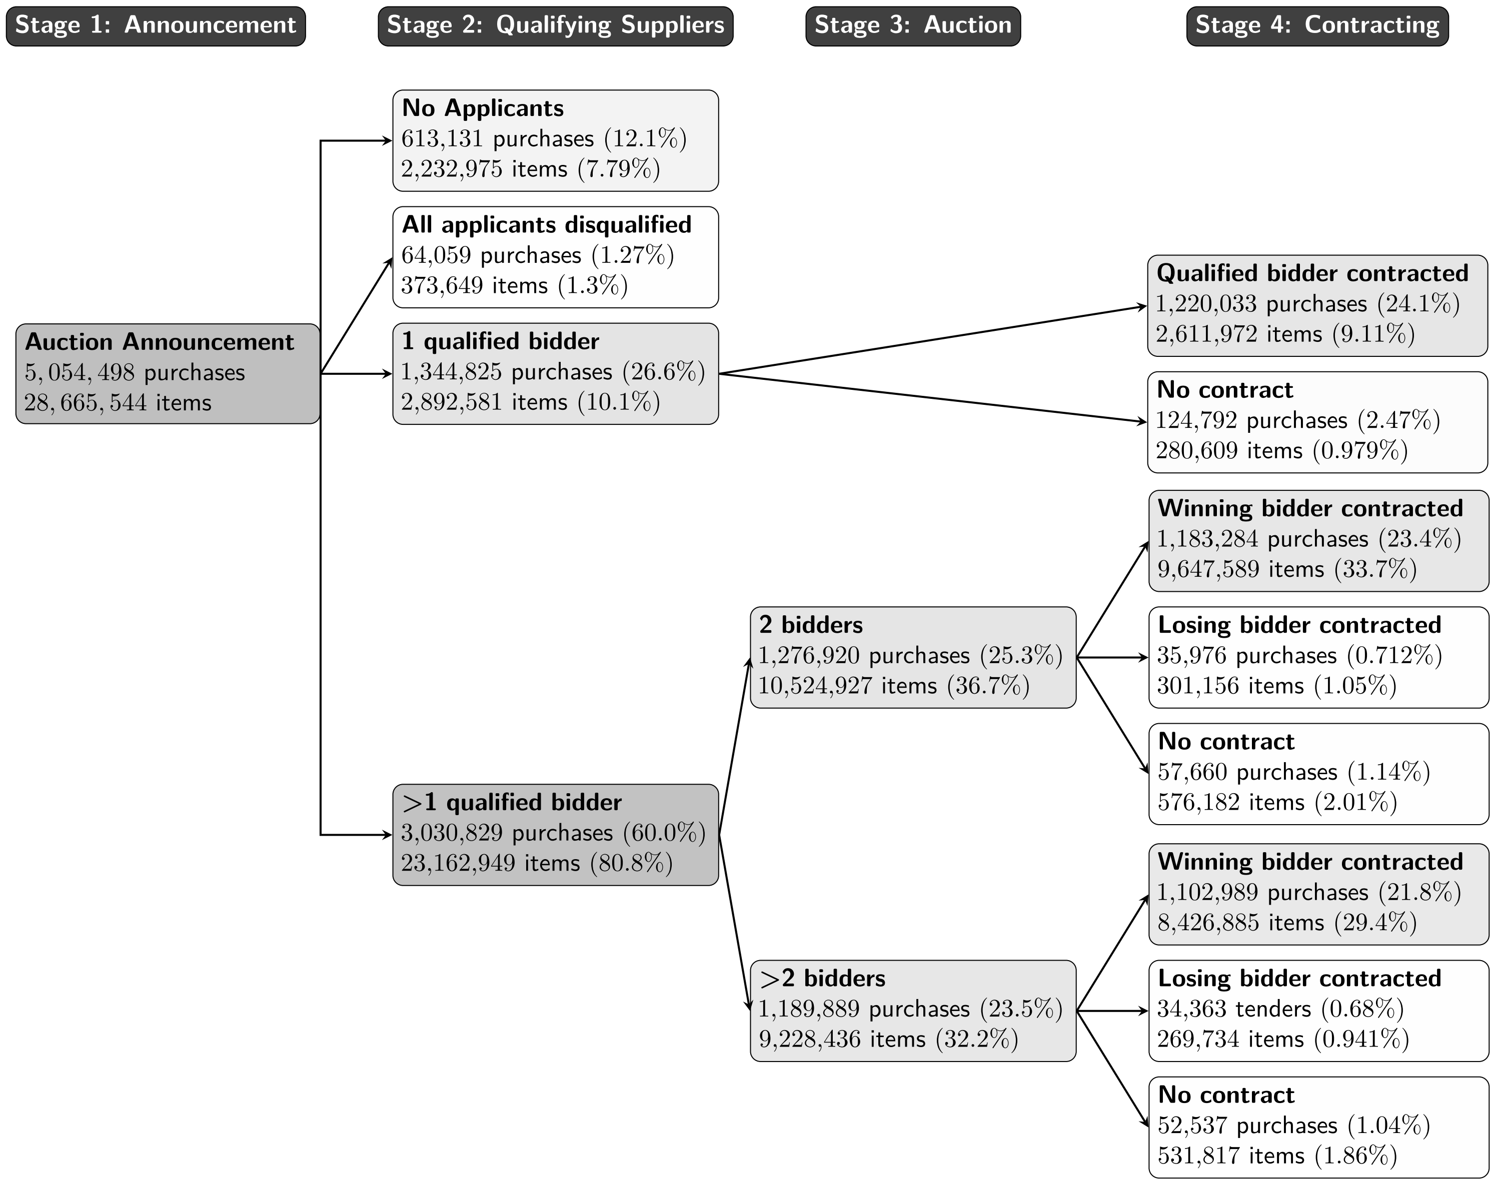

The figure summarizes the process public procurement purchases follow. Numbers are based on all purchases made under laws 94 and 44 in 2011-2015.

FIGURE 2: EVENT STUDY: ORGANIZATIONS SWITCHING BUREAUCRATS

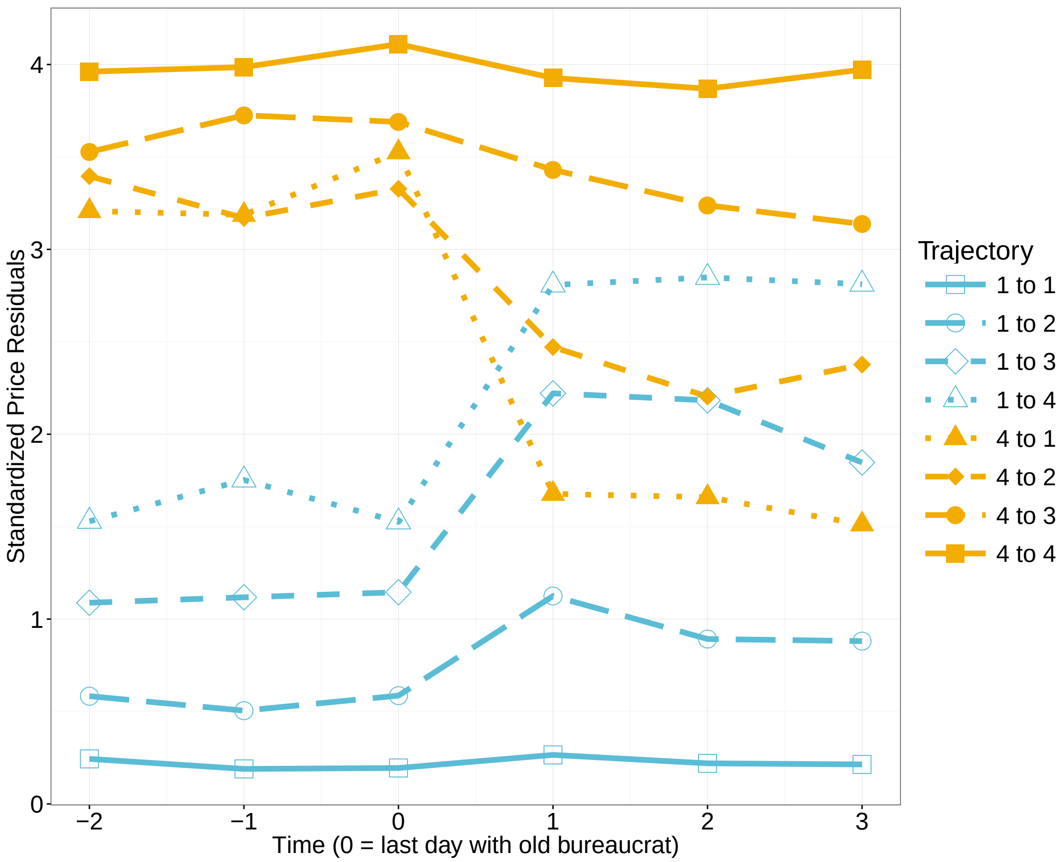

The figure shows...

FIGURE 3: EVENT STUDY: BUREAUCRATS SWITCHING GOODS

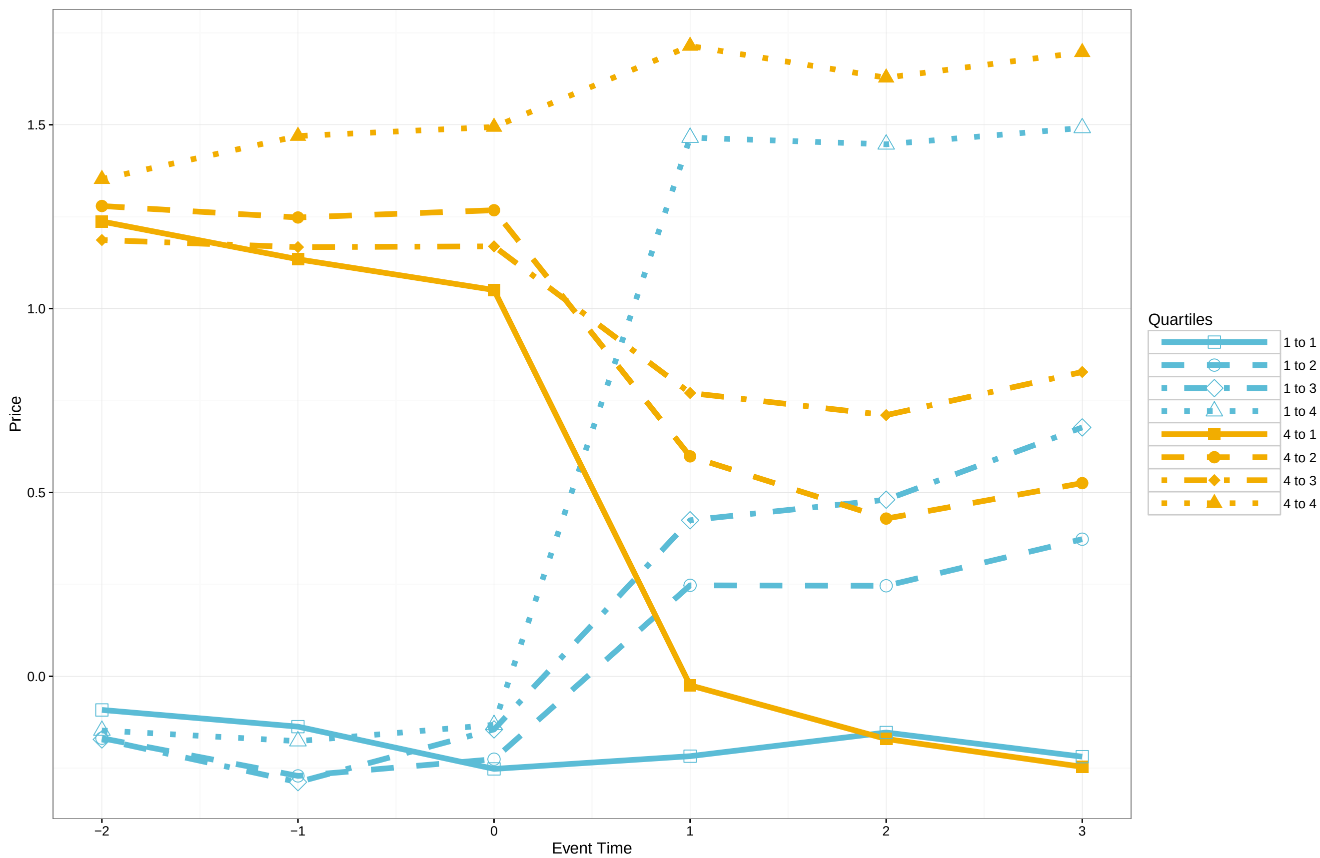

The figure shows...

FIGURE 4: EVENT STUDY: ORGANIZATIONS SWITCHING GOODS

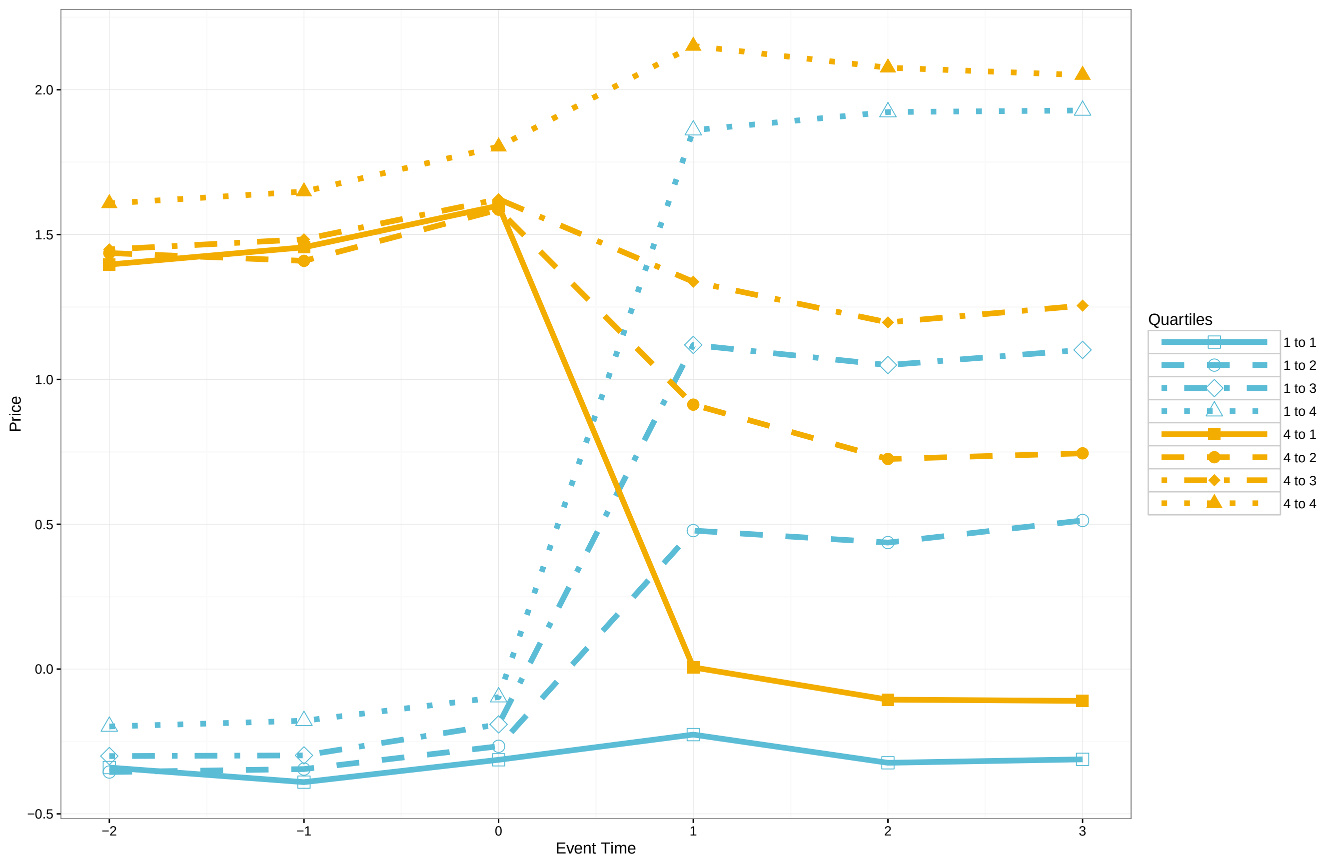

The figure shows...

FIGURE 5: NO SYSTEMATIC PATTERN IN RESIDUALS: ANALYSIS SAMPLE

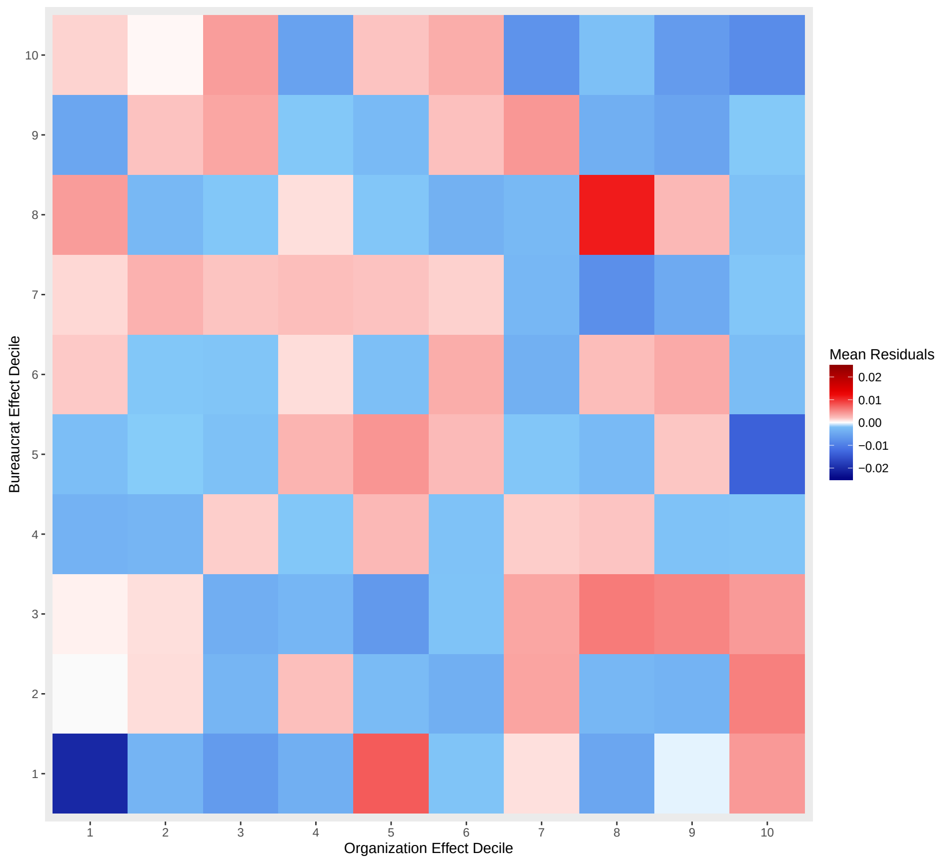

FIGURE 6: NO SYSTEMATIC PATTERN IN RESIDUALS: LARGEST CONNECTED SET

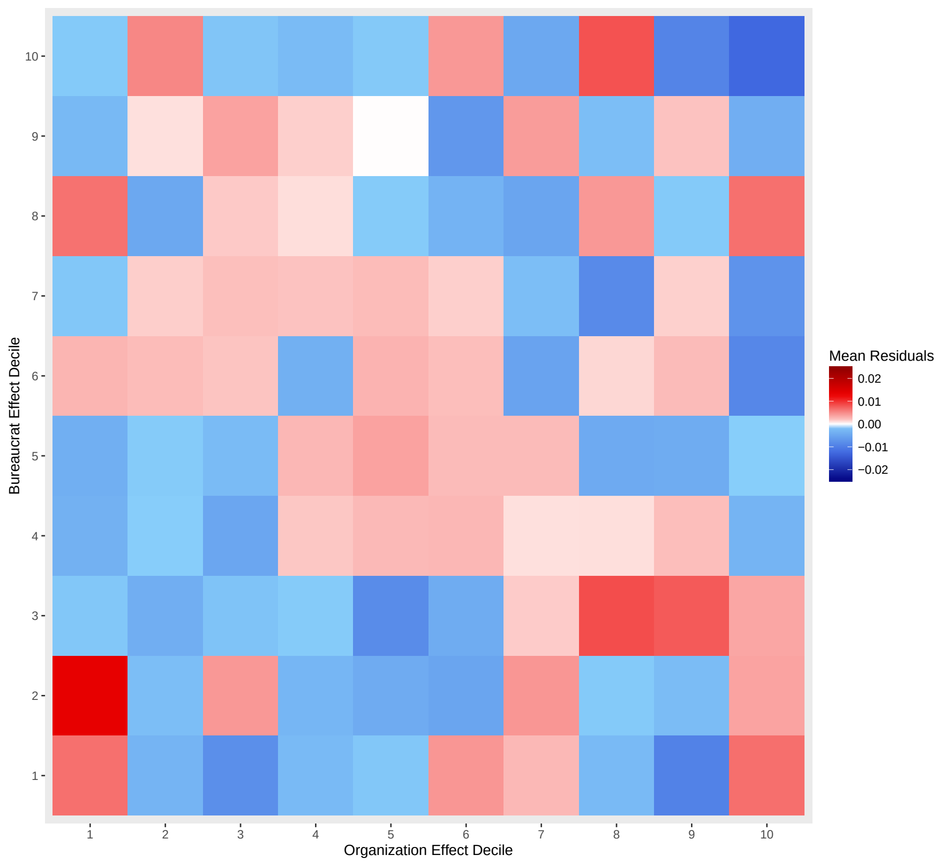

FIGURE 7: CORRELATES OF BUREAUCRATIC EFFECTIVENESS

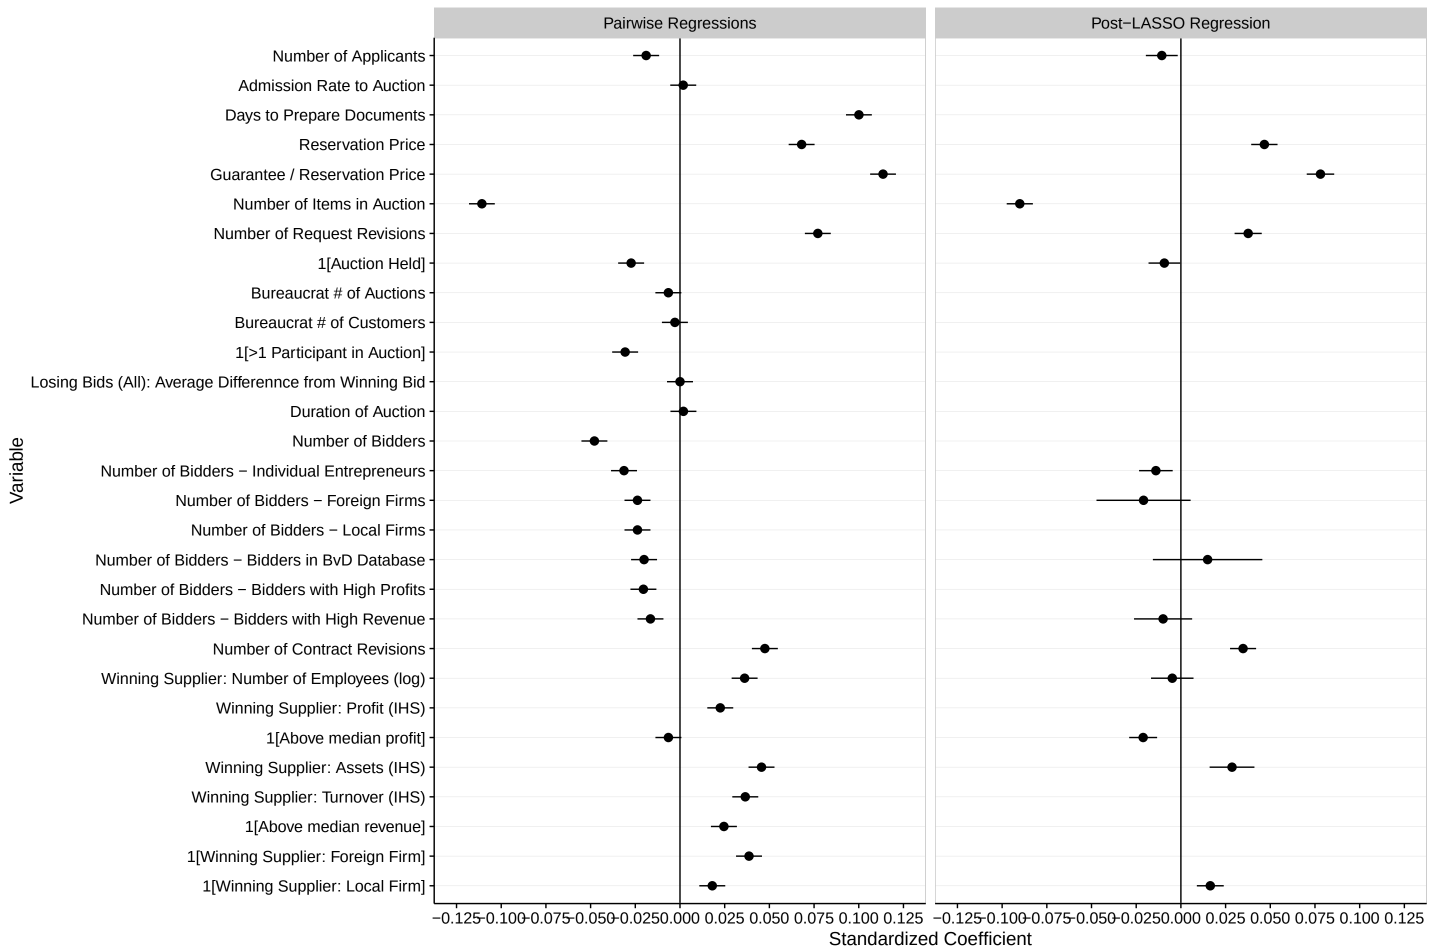

FIGURE 8: CORRELATES OF ORGANIZATION EFFECTIVENESS

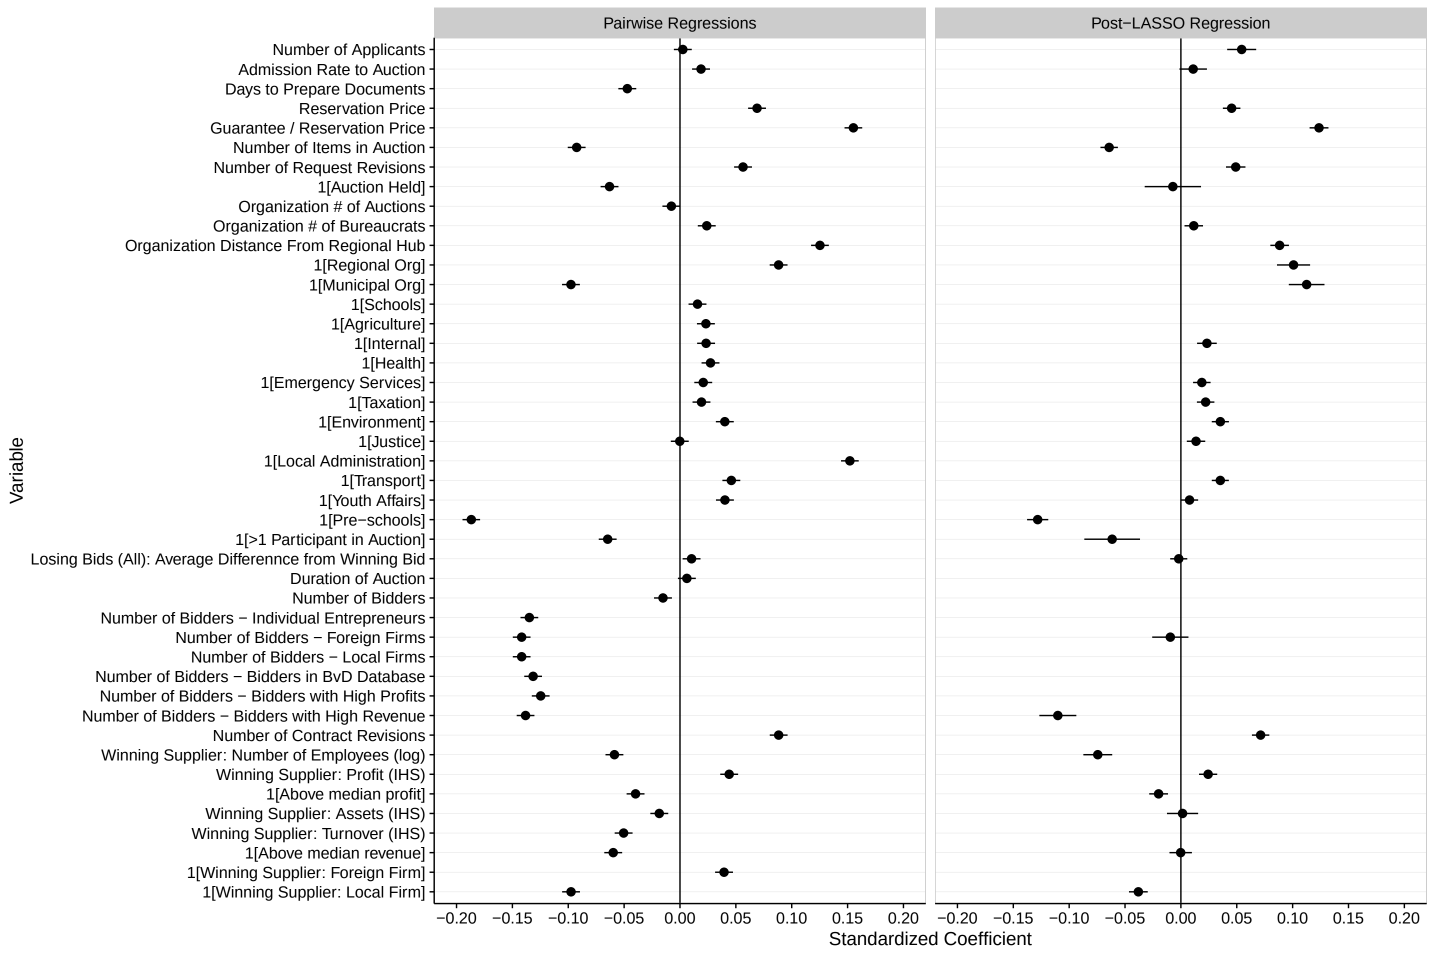

TABLE 1: SUMMARY STATISTICS: MEDICINES

	Medicines - LVEMD Sample	Medicines - Full Sample
(1) Auctions	123,679	276,033
(2) Organizations	2,677	4,695
(3) Regions	86	87
(4) Products	1,684	1,684
(5) Total Value (Bn Roubles)	138.18	218.17
(6) Federal	0.08	0.10
(7) Regional	0.80	0.78
(8) Municipal	0.12	0.12
(9) School	0.28	0.42
(10) Internal Affairs/Defense	0.01	0.01
(11) Health	0.71	0.57

TABLE 2: SUMMARY STATISTICS

	Full Sample	Analysis Sample	Largest Connected Set
(1) Number of Bureaucrats	135,632	41,470	15,791
(2) Number of Organizations	103,690	49,438	13,615
(3) Number of Connected Sets	29,532	647	1
(4) Number of Bureaucrats with >1 Org.	16,077	12,278	4,207
(5) Number of Organizations with >1 Bur.	63,468	41,919	12,228
(6) Number of Organization Types	13	13	13
(7) Number of Federal Organizations	13,976	2,028	587
(8) Number of Regional Organizations	27,004	16,255	5,523
(9) Number of Municipal Organizations	62,656	31,141	7,505
(10) Number of Goods	15,442	14,951	13,114
(11) Number of Regions	90	90	70
(12) Number of Requests	2,083,033	1,409,021	474,745
(13) Mean Number of Applicants	3.37	3.38	3.47
(14) Mean Number of Bidders	2.8	2.78	2.82
(15) Mean Reservation Price	950,856	1,017,075	1,443,743
(16) Quantity Mean	462	451	459
Median	21	25	20
SD	1,717	1,654	1,719
(17) Price Mean	11,685	10,042	20,248
Median	203	184	270
SD	60,984	55,654	88,953
(18) Unit Price Mean	10,270	8,771	18,308
Median	8.8	6.88	13
SD	81,963	88,813	131,159
(19) Number of Observations	18,052,927	17,705,338	11,304,918

TABLE 3: VARIANCE DECOMPOSITION RESULTS: ANALYSIS SAMPLE

	Fixed Effects	Split-Sample
s.d. of Bureaucrat Effects (across items)	1.588	1.490
s.d. of Organization Effects (across items)	1.607	1.490
s.d. of Connected Set Effects (across items)	0.601	0.445
s.d. of Bur + Org Effects Within CS	1.122	1.120
s.d. of Bur + Org Effects (Total)	1.273	1.205
s.d. of log P	3.198	3.198
Sample Size	17,705,354	17,705,354

The table implements...

TABLE 4: VARIANCE DECOMPOSITION RESULTS: LARGEST CONNECTED SET

	Fixed Effects	s.e.	Split-Sample	Min MSE
s.d. of Bureaucrat Effects (across items)	1.602	(0.152)	1.206	1.126
s.d. of Organization Effects (across items)	1.624	(0.151)	1.257	1.182
s.d. of Bur + Org Effects	1.171	(0.000411)	1.355	0.896
s.d. of log P	3.206		3.206	3.206
Sample Size	11,304,934		11,304,934	11,304,934

The table implements...

TABLE 5: VARIANCE DECOMPOSITION RESULTS: ADDING PAIR EFFECTS

	Analysis Sample	Largest Connected Set
RMSE of residuals	1.595	1.588
Adjusted R-squared	0.7456	0.749
RMSE of residuals (pair model)	1.551	1.533
Adjusted R-squared (pair model)	0.7566	0.7627
s.d of pair effect	0.3716	0.4141

The table implements...

TABLE 6: VARIANCE DECOMPOSITION RESULTS: ROBUSTNESS TO HETEROGENEOUS GOODS

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
s.d. of Bureaucrat Effects (across items)	1.367	1.571	1.365	1.205	1.588
s.d. of Organization Effects (across items)	1.369	1.566	1.345	1.299	1.607
s.d. of Bur + Org Effects Within CS	1.264	1.156	1.166	1.122	1.122
s.d. of Bur + Org Effects (Total)	1.490	1.323	1.336	1.275	1.273
s.d. of log P	3.425	3.329	3.300	3.251	3.198
s.d. of Bur Efs / s.d. of log P	0.399	0.472	0.414	0.371	0.496
s.d. of Org Efs / s.d. of log P	0.400	0.470	0.408	0.399	0.502
s.d. of Bur+Org Efs Within / s.d. of log P	0.369	0.347	0.353	0.345	0.351
s.d. of Bur+Org Efs Total / s.d. of log P	0.435	0.397	0.405	0.392	0.398
Sample Size	3,543,096	7,080,679	10,623,302	14,162,092	17,705,354

The table implements...

TABLE 7: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ATE, ANALYSIS SAMPLE

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.721^{***}	-0.700^{***}	-0.697^{***}	-0.660^{***}
	(0.026)	(0.025)	(0.025)	(0.022)
Preferenced	0.038	0.015	0.103^{*}	-0.206^{***}
	(0.057)	(0.051)	(0.054)	(0.035)
Bureaucrat FE		0.228^{***}		0.994^{***}
		(0.012)		(0.027)
Bureaucrat FE * Preferenced		-0.084^{***}		-0.308^{***}
		(0.010)		(0.020)
Organization FE			0.268^{***}	0.998^{***}
			(0.011)	(0.026)
Organization FE * Preferenced			0.003	-0.264^{***}
			(0.008)	(0.018)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	25,438,084	24,547,612	24,681,629	24,472,139
R^2	0.638	0.655	0.660	0.714

*** p<0.01, ** p<0.05, * p<0.1

TABLE 8: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ATE, LARGEST CONNECTED SET

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.721^{***}	-0.694^{***}	-0.689^{***}	-0.649^{***}
	(0.026)	(0.025)	(0.025)	(0.022)
Preferenced	0.038	0.010	0.112^{**}	-0.258^{***}
	(0.057)	(0.052)	(0.055)	(0.034)
Bureaucrat FE		0.224^{***}		1.000^{***}
		(0.014)		(0.027)
Bureaucrat FE * Preferenced		-0.085^{***}		-0.318^{***}
		(0.012)		(0.023)
Organization FE			0.277^{***}	1.002^{***}
			(0.013)	(0.025)
Organization FE * Preferenced			0.0003	-0.271^{***}
			(0.011)	(0.021)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	25,438,084	16,126,512	16,194,450	16,075,772
R^2	0.638	0.645	0.651	0.714

*** p<0.01, ** p<0.05, * p<0.1

TABLE 9: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ITT, ANALYSIS SAMPLE

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.718^{***}	-0.697^{***}	-0.694^{***}	-0.658^{***}
	(0.026)	(0.025)	(0.025)	(0.022)
Preferenced	-0.292^{***}	-0.259^{***}	-0.191^{***}	-0.403^{***}
	(0.063)	(0.057)	(0.065)	(0.044)
Bureaucrat FE		0.255^{***}		0.977^{***}
		(0.015)		(0.030)
Bureaucrat FE * Preferenced		-0.118^{***}		-0.346^{***}
		(0.011)		(0.022)
Organization FE			0.296^{***}	0.982^{***}
			(0.013)	(0.028)
Organization FE * Preferenced			-0.085^{***}	-0.322^{***}
			(0.011)	(0.023)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	25,438,084	24,284,618	24,514,444	24,165,243
R^2	0.639	0.656	0.660	0.708

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

TABLE 10: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ITT, LARGEST CONNECTED SET

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.718^{***}	-0.691^{***}	-0.684^{***}	-0.645^{***}
	(0.026)	(0.025)	(0.025)	(0.022)
Preferenced	-0.291^{***}	-0.274^{***}	-0.223^{***}	-0.440^{***}
	(0.063)	(0.065)	(0.074)	(0.050)
Bureaucrat FE		0.230^{***}		0.996^{***}
		(0.015)		(0.030)
Bureaucrat FE * Preferenced		-0.107^{***}		-0.381^{***}
		(0.012)		(0.025)
Organization FE			0.308^{***}	0.998^{***}
			(0.014)	(0.028)
Organization FE * Preferenced			-0.096^{***}	-0.347^{***}
			(0.012)	(0.027)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	25,438,084	15,887,143	16,011,515	15,809,065
R^2	0.639	0.644	0.651	0.707

*** p<0.01, ** p<0.05, * p<0.1

TABLE 11: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ITT, ANALYSIS SAMPLE, MEDICINES, BARCODE FE

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.049^{***}	-0.044^{***}	-0.044^{***}	-0.043^{***}
	(0.003)	(0.003)	(0.003)	(0.003)
Preferenced	-0.015	-0.018	-0.015	-0.016
	(0.014)	(0.015)	(0.015)	(0.014)
Bureaucrat FE		0.209^{***}		0.800^{***}
		(0.034)		(0.045)
Bureaucrat FE * Preferenced		-0.109^{***}		-0.321^{***}
		(0.031)		(0.056)
Organization FE			0.197^{***}	0.787^{***}
			(0.030)	(0.047)
Organization FE * Preferenced			-0.074^{***}	-0.305^{***}
			(0.023)	(0.061)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	544,347	487,852	487,852	487,852
R^2	0.935	0.938	0.938	0.940

*** p<0.01, ** p<0.05, * p<0.1

TABLE 12: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ITT, LARGEST CONNECTED SET, MEDICINES, BARCODE FE

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.049^{***}	-0.044^{***}	-0.048^{***}	-0.050^{***}
	(0.003)	(0.007)	(0.007)	(0.007)
Preferenced	-0.015	0.012	0.023	0.023
	(0.014)	(0.040)	(0.039)	(0.039)
Bureaucrat FE		0.377^{***}		0.700^{***}
		(0.111)		(0.147)
Bureaucrat FE * Preferenced		0.012		-0.102
		(0.116)		(0.133)
Organization FE			0.361^{***}	0.788^{***}
			(0.136)	(0.137)
Organization FE * Preferenced			-0.115	-0.315^{**}
			(0.135)	(0.127)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	544,347	36,060	36,060	36,060
R^2	0.935	0.942	0.942	0.943

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

TABLE 13: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ITT, ANALYSIS SAMPLE, MEDICINES, MACHINE LEARNING FE

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.326^{***}	-0.333^{***}	-0.344^{***}	-0.348^{***}
	(0.012)	(0.013)	(0.013)	(0.013)
Preferenced	-0.657^{***}	-0.669^{***}	-0.595^{***}	-0.572^{***}
	(0.034)	(0.035)	(0.034)	(0.031)
Bureaucrat FE		0.423^{***}		0.904^{***}
		(0.027)		(0.034)
Bureaucrat FE * Preferenced		-0.235^{***}		-0.384^{***}
		(0.020)		(0.024)
Organization FE			0.580^{***}	0.932^{***}
			(0.031)	(0.035)
Organization FE * Preferenced			-0.150^{***}	-0.291^{***}
			(0.021)	(0.026)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	544,347	487,852	487,852	487,852
R^2	0.450	0.473	0.487	0.514

*** p<0.01, ** p<0.05, * p<0.1

TABLE 14: AKM BUREAUCRAT AND ORGANIZATION EFFECTS AND PREFERENCE POLICY IMPACT: ITT, LARGEST CONNECTED SET, MEDICINES, MACHINE LEARNING FE

	Outcome: log Price			
	(1)	(2)	(3)	(4)
log Standardized Quantity	-0.326^{***}	-0.385^{***}	-0.433^{***}	-0.430^{***}
	(0.012)	(0.022)	(0.020)	(0.019)
Preferenced	-0.657^{***}	-0.535^{***}	-0.399^{***}	-0.369^{***}
	(0.034)	(0.076)	(0.071)	(0.069)
Bureaucrat FE		0.570^{***}		0.942^{***}
		(0.059)		(0.061)
Bureaucrat FE * Preferenced		-0.444^{***}		-0.424^{***}
		(0.059)		(0.055)
Organization FE			0.886^{***}	1.009^{***}
			(0.062)	(0.055)
Organization FE * Preferenced			-0.222^{***}	-0.266^{***}
			(0.047)	(0.044)
Month, Good FEs	Yes	Yes	Yes	Yes
Year x Product x Size x Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	No	Yes	Yes	Yes
Observations	544,347	36,060	36,060	36,060
R^2	0.450	0.490	0.549	0.567

*** p<0.01, ** p<0.05, * p<0.1

A Theoretical Appendix

A.1 Detailed Characterization of Equilibrium Without Bidding Preferences

As shown in section 2.1, the sellers' expected profits can be expressed in terms of their probabilities of winning. Using our assumptions about the distributions of seller fulfillment costs, the probabilities of winning are

$$\begin{aligned}
q_F\left(x;\overline{d}_F,\overline{d}_L\right) &= \mathsf{Pr}\left(b_F\left(x\right)<b_L\left(v_L\right)|v_L\le \overline{d}_L\right)\mathsf{Pr}\left(v_L\le\overline{d}_L\right)+1\times\mathsf{Pr}\left(v_L>\overline{d}_L\right)\\
&= \mathsf{Pr}\left(v_L>x|v_L\le\overline{d}_L\right)\frac{\overline{d}_L-\mu}{1-\mu}+\frac{1-\overline{d}_L}{1-\mu}\\
&= \begin{cases} 1 & \text{, if } x<\mu \\ \frac{1-x}{1-\mu} & \text{, if } x\in\left[\mu,\overline{d}_L\right) \\ \frac{1-\overline{d}_L}{1-\mu} & \text{, if } x\ge\overline{d}_L \end{cases} \qquad \text{(A.1)}\\
q_L\left(x;\overline{d}_F,\overline{d}_L\right) &= \mathsf{Pr}\left(b_L\left(x\right)<b_F\left(v_F\right)|v_F\le \overline{d}_F\right)\mathsf{Pr}\left(v_F\le\overline{d}_F\right)+1\times\mathsf{Pr}\left(v_F>\overline{d}_F\right)\\
&= \mathsf{Pr}\left(v_F>x|v_F\le\overline{d}_F\right)\overline{d}_F+\left(1-\overline{d}_F\right)\\
&= \begin{cases} 1-x & \text{, if } x\in\left[\mu,\overline{d}_F\right) \\ 1-\overline{d}_F & \text{, if } x\ge\overline{d}_F \end{cases} \qquad \text{(A.2)}
\end{aligned}$$

Integrating these probabilities we get the expected profits

$$\begin{aligned}
U_F\left(v;\overline{d}_F,\overline{d}_L\right) &= \int_v^1 q_F\left(x;\overline{d}_F,\overline{d}_L\right)dx\\
&= \begin{cases} \int_v^\mu 1\,dx+\int_\mu^{\overline{d}_L}\frac{1-x}{1-\mu}\,dx+\int_{\overline{d}_L}^1\frac{1-\overline{d}_L}{1-\mu}\,dx & \text{, if } v<\mu \\ \int_v^{\overline{d}_L}\frac{1-x}{1-\mu}\,dx+\int_{\overline{d}_L}^1\frac{1-\overline{d}_L}{1-\mu}\,dx & \text{, if } x\in\left[\mu,\overline{d}_L\right) \\ \int_v^1\frac{1-\overline{d}_L}{1-\mu}\,dx & \text{, if } x\ge\overline{d}_L \end{cases}\\
&= \begin{cases} \frac{2-2\overline{d}_L+\overline{d}_L^2-\mu^2}{2(1-\mu)}-v & \text{, if } v<\mu \\ \frac{2-2\overline{d}_L+\overline{d}_L^2}{2(1-\mu)}-\frac{2v-v^2}{2(1-\mu)} & \text{, if } v\in\left[\mu,\overline{d}_L\right) \\ \frac{(1-\overline{d}_L)(1-v)}{1-\mu} & \text{, if } v\ge\overline{d}_L \end{cases} \qquad \text{(A.3)}
\end{aligned}$$

And similarly for an entrant of type L with fulfillment cost v (where $\mu < \overline{d}_F$)

$$
\begin{aligned}
U_L\left(v; \overline{d}_F, \overline{d}_L\right) &= \int_v^1 q_L\left(x; \overline{d}_F, \overline{d}_L\right) dx \\
&= \int_v^{\overline{d}_F} (1-x)\, dx + \int_{\overline{d}_F}^1 \left(1-\overline{d}_F\right) dx \\
&= \begin{cases} 1 - v - \overline{d}_F + \frac{\overline{d}_F^2}{2} + \frac{1}{2}v^2 & \text{, if } v \in \left[\mu, \overline{d}_F\right) \\ \left(1-\overline{d}_F\right)(1-v) & \text{, if } v \geq \overline{d}_F \end{cases}
\end{aligned} \tag{A.4}
$$

To find the entry thresholds, we need to find the type-F supplier $\overline{d}_F$ and type-L supplier $\overline{d}_L$ who are indifferent between entering (in which case they receive $U_i\left(\overline{d}_i; \overline{d}_F, \overline{d}_L\right) - c$) and staying out of the second-stage auction (in which case they receive the contract at price 1 with probability $\frac{1}{2}\left[1 - F_j\left(\overline{d}_j\right)\right]$. That is, we need to solve the system of equations

$$
\begin{cases} U_F\left(\overline{d}_F; \overline{d}_F, \overline{d}_L\right) - c & = \frac{1}{2}(1-\overline{d}_F)\frac{1-\overline{d}_L}{1-\mu} \\ U_L\left(\overline{d}_L; \overline{d}_F, \overline{d}_L\right) - c & = \frac{1}{2}(1-\overline{d}_F)(1-\overline{d}_L) \end{cases} \tag{A.5}
$$

Since each of these equations has two cases, there are potentially two solutions, depending on whether $\overline{d}_F \lessgtr \overline{d}_L$. However, there is no solution when $\overline{d}_F < \overline{d}_L$. The solution with $\overline{d}_F > \overline{d}_L$ satisfies

$$
\begin{aligned}
&\begin{cases} \frac{(1-\overline{d}_L)(1-\overline{d}_F)}{1-\mu} & = c + \frac{1}{2}(1-\overline{d}_F)\frac{1-\overline{d}_L}{1-\mu} \\ 1 - (\overline{d}_L + \overline{d}_F) + \frac{1}{2}\left(\overline{d}_F^2 + \overline{d}_L^2\right) & = c + \frac{1}{2}(1-\overline{d}_F)(1-\overline{d}_L) \end{cases} \\
\Leftrightarrow &\begin{cases} 1 - \frac{2c(1-\mu)}{1-\overline{d}_L} & = \overline{d}_F \\ \frac{1}{2}(1-\overline{d}_F)(1-\overline{d}_L) + \frac{1}{2\gamma}\left(\overline{d}_F - \overline{d}_L\right)^2 & = c \end{cases} \\
\Leftrightarrow &\begin{cases} 1 - \frac{2c(1-\mu)}{1-\overline{d}_L} & = \overline{d}_F \\ \sqrt{2c\mu} + \overline{d}_L & = \overline{d}_F \end{cases}
\end{aligned} \tag{A.6}
$$

Solving, we see that

$$
\overline{d}_L = \frac{2 - \sqrt{2c\mu} - \sqrt{2c\,(2-\mu)}}{2} \tag{A.7}
$$

$$
\overline{d}_F = \frac{2 + \sqrt{2c\mu} - \sqrt{2c\,(2-\mu)}}{2} \tag{A.8}
$$

which characterize the entry strategies in this equilibrium. Given these, the expected number of entrants in the auction is

$$
\mathsf{E}[n] = F_A\left(\overline{d}_A\right) + F_B\left(\overline{d}_B\right) = \overline{d}_A + \frac{\overline{d}_B - \mu}{1-\mu} \tag{A.9}
$$

we can also calculate the expected payments to each bidder when their fulfillment cost is v

$$\begin{aligned}
m_F(v) &= U_F(v) + q_F(v)\,v \\
&= \begin{cases} \frac{2-2\overline{d}_L+\overline{d}_L^2-\mu^2}{2(1-\mu)} & \text{, if } v<\mu \\ \frac{2-2\overline{d}_L+\overline{d}_L^2}{2(1-\mu)} - \frac{v^2}{2(1-\mu)} & \text{, if } v\in\left[\mu,\overline{d}_L\right) \\ \frac{(1-\overline{d}_L)}{1-\mu} & \text{, if } v\geq \overline{d}_L \end{cases} && \text{(A.10)} \\
m_L(v) &= U_L(v) + q_L(v)\,v \\
&= 1-\overline{d}_F + \frac{\overline{d}_F^2}{2} - \frac{1}{2}v^2,\ v\leq \overline{d}_L < \overline{d}_F && \text{(A.11)}
\end{aligned}$$

The ex-ante expected profits of the two bidders are therefore

$$\begin{aligned}
\mathsf{E}_V\left[m_F(v)\right] &= \int_0^{\mu} \frac{2-2\overline{d}_L+\overline{d}_L^2-\mu^2}{2(1-\mu)}\,dv + \int_{\mu}^{\overline{d}_L} \frac{2-2\overline{d}_L+\overline{d}_L^2}{2(1-\mu)} - \frac{v^2}{2(1-\mu)}\,dv + \int_{\overline{d}_L}^{\overline{d}_F} \frac{(1-\overline{d}_L)}{1-\mu}\,dv \\
&= \frac{\overline{d}_L^3-\mu^3+3\overline{d}_F(1-\overline{d}_L)}{3(1-\mu)} && \text{(A.12)} \\
\mathsf{E}_V\left[m_L(v)\right] &= \int_{\mu}^{\overline{d}_L} \left(1-\overline{d}_F+\frac{\overline{d}_F^2}{2\gamma}-\frac{1}{2}v^2\right)\frac{1}{1-\mu}\,dv = \left[1-\overline{d}_F+\frac{\overline{d}_F^2}{2}\right]\frac{\overline{d}_L-\mu}{1-\mu} + \frac{(\mu^3-\overline{d}_L^3)}{6(1-\mu)} && \text{(A.13)}
\end{aligned}$$

Together, these imply that the price the auctioneer expects to pay is

$$\begin{aligned}
\mathsf{E}\left[p\right] &= \mathsf{E}_V\left[m_F(v)\right] + \mathsf{E}_V\left[m_L(v)\right] + \mathsf{Pr}\left(n=0\right) \\
&= 1-\left(1-\frac{\overline{d}_F}{2}\right)\overline{d}_F\frac{\overline{d}_L-\mu}{1-\mu} + \frac{\overline{d}_L^3-\mu^3}{6\left(1-\mu\right)} && \text{(A.14)}
\end{aligned}$$

A.2 Proof of Proposition 1

Proof. The proposition can be shown by simple differentiation. Starting with the expected number of entrants, differentiating (5), we see that

$$\frac{\partial \mathsf{E}\left[n\right]}{\partial c} = \frac{\partial \overline{d}_F}{\partial c} + \frac{1}{1-\mu}\frac{\partial \overline{d}_L}{\partial c} \qquad \text{(A.15)}$$

which depends on how the entry thresholds change with c. Differentiating the expressions for the entry thresholds (A.7) and (A.8),

$$\begin{aligned}
\frac{\partial \overline{d}_L}{\partial c} &= \frac{1}{2}\left[-\frac{\sqrt{2c\mu}}{2c} - \frac{\sqrt{2c(2-\mu)}}{2c}\right] \\
&= -\frac{1}{2c}\left(1-\overline{d}_L\right) < 0 \\
\frac{\partial \overline{d}_F}{\partial c} &= \frac{1}{2}\left[\frac{\sqrt{2c\mu}}{2c} - \frac{\sqrt{2c(2-\mu)}}{2c}\right] \\
&= -\frac{1}{2c}\left(1-\overline{d}_F\right) < 0
\end{aligned}$$

Plugging these into (A.15), we obtain

$$\frac{\partial \mathsf{E}[n]}{\partial c} = -\frac{1}{2c}\left[\left(1-\overline{d}_F\right) + \frac{1-\overline{d}_L}{1-\mu}\right] < 0 \tag{A.16}$$

showing the first part of the proposition. Following the same steps for the second part, the derivative of the expected price is

$$\mathsf{E}[p] = 1 - \left(1-\frac{\overline{d}_F}{2}\right)\overline{d}_F\frac{\overline{d}_L-\mu}{1-\mu} + \frac{\overline{d}_L^3-\mu^3}{6(1-\mu)}$$

and inserting the expressions for the thresholds' derivatives, we obtain

$$\begin{aligned}
\frac{\partial \mathsf{E}[p]}{\partial c} &= \frac{1}{2}\frac{\partial \overline{d}_F}{\partial c}\overline{d}_F\frac{\overline{d}_L-\mu}{1-\mu} - \left(1-\frac{\overline{d}_F}{2}\right)\left[\frac{\overline{d}_L-\mu}{1-\mu}\frac{\partial \overline{d}_F}{\partial c} + \frac{\overline{d}_F}{1-\mu}\frac{\partial \overline{d}_L}{\partial c}\right] + \frac{\overline{d}_L^2}{2(1-\mu)}\frac{\partial \overline{d}_L}{\partial c} \\
&= -\frac{\overline{d}_F}{2c}\frac{\overline{d}_L-\mu}{1-\mu}\left(1-\overline{d}_F\right) + \frac{\overline{d}_L-\mu}{1-\mu}\frac{1-\overline{d}_F}{2c} + \left(1-\frac{\overline{d}_F}{2}\right)\frac{\overline{d}_F}{1-\mu}\frac{1-\overline{d}_L}{2c} - \frac{\overline{d}_L^2}{2(1-\mu)}\frac{1-\overline{d}_L}{2c} \\
&= \frac{\overline{d}_L-\mu}{1-\mu}\frac{\left(1-\overline{d}_F\right)^2}{2c} + \frac{1-\overline{d}_L}{4c(1-\mu)}\left[\left(1-\overline{d}_L\right)^2 - \left(1-\overline{d}_F\right)^2 + 2\overline{d}_L\right] > 0
\end{aligned}$$

where the last inequality follows since $\overline{d}_F \geq \overline{d}_L$, completing the proof. □

A.3 Detailed Characterization of Equilibrium Without Bidding Preferences

As shown in section 2.1, the sellers' expected profits can be expressed in terms of their probabilities of winning. Using our assumptions about the distributions of seller fulfillment costs, the probabilities of winning are

$$
\begin{aligned}
q_F\left(x;\overline{d}_F,\overline{d}_L\right) &= \mathsf{Pr}\left(b_F\left(x\right)<b_L\left(v_L\right)|v_L\leq\overline{d}_L\right)\mathsf{Pr}\left(v_L\leq\overline{d}_L\right)+1\times\mathsf{Pr}\left(v_L>\overline{d}_L\right)\\
&= \mathsf{Pr}\left(v_L>\frac{x}{\gamma}|v_L\leq\overline{d}_L\right)\frac{\overline{d}_L-\mu}{1-\mu}+\frac{1-\overline{d}_L}{1-\mu}\\
&= \begin{cases} 1 & \text{, if } x<\gamma\mu \\ \frac{\gamma-x}{\gamma(1-\mu)} & \text{, if } x\in\left[\gamma\mu,\gamma\overline{d}_L\right) \\ \frac{1-\overline{d}_L}{1-\mu} & \text{, if } x\geq\gamma\overline{d}_L \end{cases} \qquad \text{(A.17)}\\
q_L\left(x;\overline{d}_F,\overline{d}_L\right) &= \mathsf{Pr}\left(b_L\left(x\right)<b_F\left(v_F\right)|v_F\overline{d}_F\right)\mathsf{Pr}\left(v_F\leq\overline{d}_F\right)+1\times\mathsf{Pr}\left(v_F>\overline{d}_F\right)\\
&= \mathsf{Pr}\left(\frac{v_F}{\gamma}>x|v_F\leq\overline{d}_F\right)\overline{d}_F+\left(1-\overline{d}_F\right)\\
&= \begin{cases} 1-x\gamma & \text{, if } x\in\left[\mu,\frac{\overline{d}_F}{\gamma}\right) \\ 1-\overline{d}_F & \text{, if } x\geq\frac{\overline{d}_F}{\gamma} \end{cases} \qquad \text{(A.18)}
\end{aligned}
$$

Integrating these probabilities we get the expected profits

$$
\begin{aligned}
U_F\left(v;\overline{d}_F,\overline{d}_L\right) &= \int_v^1 q_F\left(x;\overline{d}_F,\overline{d}_L\right)dx\\
&= \begin{cases} \int_v^{\gamma\mu}1\,dx+\int_{\gamma\mu}^{\gamma\overline{d}_L}\frac{\gamma-x}{\gamma(1-\mu)}\,dx+\int_{\gamma\overline{d}_L}^1\frac{1-\overline{d}_L}{1-\mu}\,dx & \text{, if } v<\gamma\mu \\ \int_v^{\gamma\overline{d}_L}\frac{\gamma-x}{\gamma(1-\mu)}\,dx+\int_{\gamma\overline{d}_L}^1\frac{1-\overline{d}_L}{1-\mu}\,dx & \text{, if } x\in\left[\gamma\mu,\gamma\overline{d}_L\right) \\ \int_v^1\frac{1-\overline{d}_L}{1-\mu}\,dx & \text{, if } x\geq\gamma\overline{d}_L \end{cases}\\
&= \begin{cases} \frac{2-2\overline{d}_L+\gamma\overline{d}_L^2-\mu^2\gamma}{2(1-\mu)}-v & \text{, if } v<\gamma\mu \\ \frac{2-2\overline{d}_L+\gamma\overline{d}_L^2}{2(1-\mu)}-\frac{2\gamma v-v^2}{2\gamma(1-\mu)} & \text{, if } v\in\left[\gamma\mu,\gamma\overline{d}_L\right) \\ \frac{(1-\overline{d}_L)(1-v)}{1-\mu} & \text{, if } v\geq\gamma\overline{d}_L \end{cases} \qquad \text{(A.19)}
\end{aligned}
$$

And similarly for an entrant of type L with fulfillment cost v (where $\mu<\frac{\overline{d}_F}{\gamma}$)

$$
\begin{aligned}
U_L\left(v;\overline{d}_F,\overline{d}_L\right) &= \int_v^1 q_L\left(x;\overline{d}_F,\overline{d}_L\right)dx\\
&= \int_v^{\overline{d}_F/\gamma}\left(1-x\gamma\right)dx+\int_{\overline{d}_F/\gamma}^1\left(1-\overline{d}_F\right)dx\\
&= \begin{cases} 1-v-\overline{d}_F+\frac{\overline{d}_F^2}{2\gamma}+\frac{\gamma}{2}v^2 & \text{, if } v\in\left[\mu,\frac{\overline{d}_F}{\gamma}\right) \\ \left(1-\overline{d}_F\right)\left(1-v\right) & \text{, if } v\geq\frac{\overline{d}_F}{\gamma} \end{cases} \qquad \text{(A.20)}
\end{aligned}
$$

To find the entry thresholds, we need to find the type-F supplier $\overline{d}_F$ and type-L supplier $\overline{d}_L$ who are indifferent between entering (in which case they receive $U_i\left(\overline{d}_i;\overline{d}_F,\overline{d}_L\right)-c$) and staying out of the second-stage auction (in which case they receive the contract at price 1 with probability $\frac{1}{2}\left[1-F_j\left(\overline{d}_j\right)\right]$.

That is, we need to solve the system of equations

$$\begin{cases} U_F\left(\bar{d}_F; \bar{d}_F, \bar{d}_L\right) - c & = \frac{1}{2}(1-\bar{d}_F)\frac{1-\bar{d}_L}{1-\mu} \\ U_L\left(\bar{d}_L; \bar{d}_F, \bar{d}_L\right) - c & = \frac{1}{2}(1-\bar{d}_F)(1-\bar{d}_L) \end{cases} \tag{A.21}$$

Since each of these equations has two cases, there are potentially two solutions, depending on whether $\bar{d}_F \lesseqgtr \gamma\bar{d}_L$. However, there is no solution when $\bar{d}_F < \gamma\bar{d}_L$. The solution with $\bar{d}_F > \gamma\bar{d}_L$ satisfies

$$\begin{aligned} &\begin{cases} \frac{(1-\bar{d}_L)(1-\bar{d}_F)}{1-\mu} & = c + \frac{1}{2}(1-\bar{d}_F)\frac{1-\bar{d}_L}{1-\mu} \\ 1 - (\bar{d}_L + \bar{d}_F) + \frac{1}{2\gamma}\left(\bar{d}_F^2 + \gamma^2\bar{d}_L^2\right) & = c + \frac{1}{2}(1-\bar{d}_F)(1-\bar{d}_L) \end{cases} \\ \Leftrightarrow &\begin{cases} 1 - \frac{2c(1-\mu)}{1-\bar{d}_L} & = \bar{d}_F \\ \frac{1}{2}(1-\bar{d}_F)(1-\bar{d}_L) + \frac{1}{2\gamma}\left(\bar{d}_F - \gamma\bar{d}_L\right)^2 & = c \end{cases} \\ \Leftrightarrow &\begin{cases} 1 - \frac{2c(1-\mu)}{1-\bar{d}_L} & = \bar{d}_F \\ \sqrt{2\gamma c\mu} + \gamma\bar{d}_L & = \bar{d}_F \end{cases} \end{aligned} \tag{A.22}$$

Solving, we see that

$$\bar{d}_L = \frac{1+\gamma-\sqrt{2\gamma c\mu} - \sqrt{\left[(1-\gamma)-\sqrt{2\gamma c\mu}\right]^2 + 4\gamma c\,(1-\mu)}}{2\gamma} \tag{A.23}$$

$$\bar{d}_F = \frac{1+\gamma+\sqrt{2\gamma c\mu} - \sqrt{\left[(1-\gamma)-\sqrt{2\gamma c\mu}\right]^2 + 4\gamma c\,(1-\mu)}}{2} \tag{A.24}$$

which characterize the entry strategies in this equilibrium. Given these, the expected number of entrants in the auction is

$$\mathsf{E}[n] = F_A\left(\bar{d}_A\right) + F_B\left(\bar{d}_B\right) = \bar{d}_A + \frac{\bar{d}_B - \mu}{1-\mu} \tag{A.25}$$

we can also calculate the expected payments to each bidder when their fulfillment cost is v

$$\begin{aligned} m_F(v) &= U_F(v) + q_F(v)\,v \\ &= \begin{cases} \frac{2-2\bar{d}_L+\gamma\bar{d}_L^2-\mu^2\gamma}{2(1-\mu)} & \text{, if } v < \gamma\mu \\ \frac{2-2\bar{d}_L+\gamma\bar{d}_L^2}{2(1-\mu)} - \frac{v^2}{2\gamma(1-\mu)} & \text{, if } v \in \left[\gamma\mu, \gamma\bar{d}_L\right) \\ \frac{(1-\bar{d}_L)}{1-\mu} & \text{, if } v \geq \gamma\bar{d}_L \end{cases} \end{aligned} \tag{A.26}$$

$$\begin{aligned} m_L(v) &= U_L(v) + q_L(v)\,v \\ &= 1 - \bar{d}_F + \frac{\bar{d}_F^2}{2\gamma} - \frac{\gamma}{2}v^2,\ v \leq \bar{d}_L < \bar{d}_F \end{aligned} \tag{A.27}$$

The ex-ante expected profits of the two bidders are therefore

$$
\begin{aligned}
\mathsf{E}_V\left[m_F(v)\right] &= \int_0^{\gamma\mu} \tfrac{2-2\overline{d}_L+\gamma\overline{d}_L^2-\mu^2\gamma}{2(1-\mu)}\,dv + \int_{\gamma\mu}^{\gamma\overline{d}_L} \tfrac{2-2\overline{d}_L+\gamma\overline{d}_L^2}{2(1-\mu)} - \tfrac{v^2}{2\gamma(1-\mu)}\,dv + \int_{\gamma\overline{d}_L}^{\overline{d}_F} \tfrac{(1-\overline{d}_L)}{1-\mu}\,dv \\
&= \tfrac{\gamma^2(\overline{d}_L^3-\mu^3)+3\overline{d}_F(1-\overline{d}_L)}{3(1-\mu)} \qquad \text{(A.28)}
\end{aligned}
$$

$$
\mathsf{E}_V\left[m_L(v)\right] = \int_\mu^{\overline{d}_L}\left(1-\overline{d}_F+\frac{\overline{d}_F^2}{2\gamma}-\frac{\gamma}{2}v^2\right)\frac{1}{1-\mu}\,dv = \left[1-\overline{d}_F+\frac{\overline{d}_F^2}{2\gamma}\right]\frac{\overline{d}_L-\mu}{1-\mu}+\frac{\gamma(\mu^3-\overline{d}_L^3)}{6(1-\mu)} \qquad \text{(A.29)}
$$

Together, these imply that the price the auctioneer expects to pay is

$$
\begin{aligned}
\mathsf{E}\left[p\right] &= \mathsf{E}_V\left[m_F(v)\right]+\mathsf{E}_V\left[m_L(v)\right]+\mathsf{Pr}\left(n=0\right) \\
&= 1-\left(1-\frac{\overline{d}_F}{2\gamma}\right)\overline{d}_F\frac{\overline{d}_L-\mu}{1-\mu}+\frac{\overline{d}_L^3-\mu^3}{6\left(1-\mu\right)}\gamma\left(2\gamma-1\right) \qquad \text{(A.30)}
\end{aligned}
$$

A.4 Proof of Proposition 2

Proof. To prove the proposition we proceed in three steps. First, we show that for any level of entry costs, there is a threshold γ below which introducing preferences at that rate causes prices to increase, and above which prices decrease. Second, we show that this threshold is decreasing in the entry costs procurers impose on suppliers. Third we argue that these first two steps imply the proposition. Our first step can be characterized in the following lemma.

Lemma 3. *For every $y \in [0,1]$ there exists a unique $\gamma^\star(y) \in [0,1]$ that satisfies*

$$
E\left[p|c=y,\gamma=\gamma^\star(y)\right] = E\left[p|c=y,\gamma=1\right], \qquad \text{(A.31)}
$$

Moreover, $E\left[p|c=y,\gamma=1\right] \le E\left[p|c=y,\gamma=\gamma^\star(y)\right],\ \forall\gamma\in\left[\gamma^\star(y),1\right]$ and $E\left[p|c=y,\gamma=1\right] \ge E\left[p|c=y,\gamma=\gamma^\star(y)\right],$ $[0,\gamma^\star(y)]$

Proof. Type up proof □

The second step is to show that higher-cost procurers have a lower $\gamma^\star$. The following lemma shows this

Lemma 4. *The price-equalizing γ is lower for procurers who impose larger entry costs on suppliers:*

$$
\frac{\partial\gamma^\star(y)}{\partial y} < 0 \qquad \text{(A.32)}
$$

Proof. Type up proof □

From these two lemmas the proof follows easily. To see part (i) consider a particular $x < \overline{\gamma}$. By lemma 3, $\mathsf{E}\left[p|c=y,\gamma=x\right] - \mathsf{E}\left[p|c=y,\gamma=1\right] < 0$ for all procurers whose entry costs y are such that $\gamma^\star(y) < x$. Conversely, $\mathsf{E}\left[p|c=y,\gamma=x\right] - \mathsf{E}\left[p|c=y,\gamma=1\right] > 0$ for all procurers whose entry costs y are such that $\gamma^\star(y) > x$. By lemma 4, $\gamma^\star(y) < x$ for all procurers with entry costs higher than $c^\star(x)$, and $\gamma^\star(y) > x$ for all procurers with entry costs above $c^{\star}(x)$, where $c^{\star}(x)$ is the unique cost level satisfying $\gamma^{\star}(c^{\star}(x)) = x$. Part (ii) follows immediately from the continuity of $\mathsf{E}[p]$ and $\mathsf{E}[n]$ in c and γ.

$\square$

B Identification of Bureaucrat and Organization Effects with Multiple Connected Sets

As shown in ACK 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.

ACK normalize the grand mean of the firm effects to be zero while the person effects have to be mean zero in each connected set. This is arbitrary, and 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. So we will have zero mean constraints on both the person and firm effects in each connected set. This also allows us to identify S connected set means $\gamma_s = \bar{\alpha}_s + \bar{\psi}_s$.

To see this explicitly, write the model as

$$\mathbf{p} = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\alpha} + \mathbf{F}\boldsymbol{\psi} \tag{B.33}$$

where **p** is the $N \times 1$ vector of item prices; **X** is an $N \times k$ matrix of control variables, **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; **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

$$\begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\alpha}}_{OLS} \\ \hat{\boldsymbol{\psi}}_{OLS} \end{bmatrix} = \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \end{bmatrix} \mathbf{p} \tag{B.34}$$

As ACK (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.

ACK 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, so that all the effects are to be interpreted as deviations from the grand mean. This makes comparisons across connected sets difficult, so we prefer to add $2N_s$ restrictions setting the mean of the bureaucrat and organization effects to 0 within each connected set, and adding N_s intercepts to capture the average bureaucrat and organization effect in each connected set.

Specifically, we augment the model to be

$$\mathbf{p} = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\tilde{\boldsymbol{\alpha}} + \mathbf{F}\tilde{\boldsymbol{\psi}} + \mathbf{S}\boldsymbol{\gamma} \tag{B.35}$$

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

$$\begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \begin{bmatrix} \mathbf{B} & \mathbf{F} & \mathbf{S} \end{bmatrix} \\ \begin{bmatrix} \mathbf{S_b} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S_o} & \mathbf{0} \end{bmatrix} \end{bmatrix} \begin{bmatrix} \hat{\alpha} \\ \hat{\psi} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{B}' \\ \mathbf{F}' \\ \mathbf{S}' \end{bmatrix} \mathbf{p} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \tag{B.36}$$

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 5 (Identification)**.** *If the true model is given by* (B.33)*, then $\hat{\boldsymbol{\alpha}}$, $\hat{\boldsymbol{\psi}}$, and $\hat{\boldsymbol{\gamma}}$, the estimators of $\tilde{\boldsymbol{\alpha}}$, $\tilde{\boldsymbol{\psi}}$ and $\boldsymbol{\gamma}$ in the augmented model* (B.35) *that solve the augmented normal equations* (B.36) *(i) are uniquely identified, and (ii) are related to the true bureaucrat and organization effects $\boldsymbol{\alpha}$ and $\boldsymbol{\psi}$ by*

$$\begin{bmatrix} \hat{\boldsymbol{\alpha}} \\ \hat{\boldsymbol{\psi}} \\ \hat{\gamma} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\alpha} - \mathbf{S_b}'\overline{\boldsymbol{\alpha}} \\ \boldsymbol{\psi} - \mathbf{S_o}'\overline{\boldsymbol{\psi}} \\ \overline{\boldsymbol{\alpha}} + \overline{\boldsymbol{\psi}} \end{bmatrix} \tag{B.37}$$

where $\overline{\boldsymbol{\alpha}}$ is the $N_s \times 1$ vector of connected-set bureaucrat effect means, and $\overline{\boldsymbol{\psi}}$ 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 (B.35) only has rank $N_b + N_o - N_s$. To see this, we note that $\mathbf{B}\mathbf{S_b}' = \mathbf{F}\mathbf{S_o}' = \mathbf{S}$ and so $2N_s$ columns of the $N \times (N_b + N_o + N_s)$ matrix $[\mathbf{B}\ \mathbf{F}\ \mathbf{S}]$ are collinear. However, the $2N_s$ restrictions $\mathbf{S_b}\hat{\boldsymbol{\alpha}} = 0$ and $\mathbf{S_o}\hat{\boldsymbol{\psi}} = 0$ are independent of the standard normal equations, so the first matrix in (B.36) has rank $N_b + N_o + N_s$ and hence the solution to (B.36) is unique.

To see the second part, it suffices to show that (B.37) solves (B.36). First, substitute the estimators out of (B.36) using (B.37) and substitute in the true model using (B.33) to rewrite (B.36) as

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

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 $\mathbf{S_o}\boldsymbol{\psi}$ is an $N_s \times 1$ vector in which each entry is the sum of the organization effects, shows that the two sides are equal, yielding the result $\square$

C Technical Appendix 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 four 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. Fourth, we assign each cluster in the third step a more aggregated 6-digit HS product code in order to match to external measures of product homogeneity (Rauch, 1999; Khandelwal, 2010) for use in robustness exercises.

C.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 section [LINK]. Each product description is parsed in the following way, using the Russian libraries for Python's Natural Language Toolkit[46]

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', and [A good example of the case of a word, but first look up which case NLK keeps]

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]. [47] Similarly, the product description "sodium bicarbonate - solution for

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

[47] 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].

infusion 5%,200ml" would result in the following tokens: [sodium, bicarbonate, solution, infusion, 5%, 200ml].[48]

C.2 Classification

In the second step of our procedure we train a classification algorithm to label each of the sentences in the customs data with one of the H_C labels in the set of labels in the customs dataset, $\mathcal{H}_C$. To prepare our input data, each of the N_C tokenized sentences $\mathbf{t}_i$ in the customs dataset is transformed into a vector of token indicators and indicators for each possible bi-gram (word-pair), denoted by $\mathbf{x}_i \in \mathcal{X}_C$.[49] Each sentence also has a corresponding good classification $g_i \in \mathcal{G}_C$, so we can represent our customs data as the pair $\{\mathbf{X}_C, \mathbf{g}_C\}$ and we seek to find a classifier $\hat{g}_C(\mathbf{x}) : \mathcal{X}_C \to \mathcal{H}_C$ that assigns every text vector $\mathbf{x}$ to a product code.

As is common in the literature, rather than solving this multiclass classification problem in a single step, we pursue a "one-versus-all" approach and reduce the problem of choosing among G possible good classifications to G_C binary choices between a single good and all other goods, and then combine them (RIFKIN 2004 CITE). 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:

$$\hat{g}_C(\mathbf{x}_i) = \arg\max_{g \in \mathcal{G}_C} p_g(\mathbf{x}_i) \tag{C.38}$$

Each binary classifier is a linear support vector machine, with a hinge loss function.[50] That is, it solves

$$\min_{\mathbf{w}_g, a_g} \frac{1}{N_C} \sum_{i=1}^{N_C} \max\{0, 1 - y_{gi} \cdot (\mathbf{w}_g \cdot \mathbf{x}_i + a_g)\} \tag{C.39}$$

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 (C.38). We implement this procedure using the Vowpal Wabbit library for Python.[51] 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

[48] 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].

[49] 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].

[50] 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)

[51] See http://hunch.net/ vw/.

TABLE C.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

a regularization penalty to improve out of sample fit.[INSERT NUMBER OF TIMES THE CROSSWALK GETS THE 6-DIGIT HS CODE RIGHT TOO]

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 [PAN & YANG CITE], 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 **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 C.1 (translated into English) and C.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.

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

TABLE C.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

described in section [LINK] below. To identify incorrectly labeled sentences, we identify sentences that have been classified as belonging to a certain good, but are very different from the average sentence with that classification. To do this, we take the tokenized sentences that have been labeled as good g, $\sqcup_g = \{\mathbf{t}_i : \hat{g}_C(\mathbf{x}_i) = g\}$ and transform them into vectors of indicators for the tokens $\mathbf{v}_{hi}$.[52] For each good, we then calculate the mean sentence vector as $\mathbf{v}_g = \sum_{\mathbf{v}_{gi}} \mathbf{v}_{gi} / |\mathbf{t}_g|$. Then, to identify outlier sentences, we calculate each sentence's normalized cosine similarity with the good's mean vector,

$$\theta_{gi} = \frac{\bar{s}_g - s(\mathbf{v}_{gi}, \mathbf{v}_g)}{\bar{s}_g} \tag{C.40}$$

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 $\mathbf{v}_g$,[53] 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. Sentences with a normalized cosine similarity above a threshold $\overline{\theta}$ are deemed to be misclassified. To choose the threshold $\overline{\theta}$, we use the customs data again. We apply the classification algorithm to the customs data, and identify correctly classified sentences ($\hat{g}_C(\mathbf{x}_i) = g_i$) and incorrectly classified sentences ($\hat{g}_C(\mathbf{x}_i) \neq g_i$). A good choice of

[52] 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

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

the threshold $\overline{\theta}$ will minimize the sum of type I and type II errors

$$V\left(\bar{\theta}\right)=\underbrace{\sum_{\hat{g}_{C}(\mathbf{x}_{i})\neq g_{i}} I\left\{\theta_{i}<\bar{\theta}\right\}}_{\text{Type I errors}}+\underbrace{\sum_{\hat{g}_{C}(\mathbf{x}_{i})=g_{i}} I\left\{\theta_{i}>\bar{\theta}\right\}}_{\text{Type II errors}} \tag{C.41}$$

In the customs data $V\left(\overline{\theta}\right)$ is roughly flat between 0.65 and 0.95, so we choose the midpoint $\bar{\theta}=0.8$ as our baseline threshold, and perform robustness to the choice of threshold by using $\bar{\theta}=0.65, 0.95$ instead.

C.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

$$\min_{\mathbf{c}} \sum_{i=1}^{N} \left\|f\left(\mathbf{c},\mathbf{t}_{i}\right)-\mathbf{t}_{i}\right\|^{2} \tag{C.42}$$

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

$$\eta\left(i\right)=\frac{b\left(i\right)-a\left(i\right)}{\max\left\{b\left(i\right),a\left(i\right)\right\}} \tag{C.43}$$

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. Picking $k=6,500$ minimizes the average silhouette coefficient in the data. For robustness, we also use $k=3,000$ and $k=10,000$.

[add a few examples]

C.4 Assigning Clusters to 6-digit HS codes

Finally, our fourth step assigns each of the k clusters to a 6-digit HS good classification in order to be able to match the sentences to external measures of product homogeneity (Rauch, 1999; Khandelwal, 2010) for use in robustness exercises. We begin by taking each usable Russian product code and using a concordance table (reference) to match it to its corresponding HS-6 classification. Next within each cluster, we identify the most frequent six-digit code and assign it to all products in that cluster.[54] Note that in this step we have used the Russian product codes contained in the procurement data that we do not believe are usable to adequately control for quality. Nevertheless, we believe that using aggregate information within each cluster like this is likely to yield a reliable HS-6 classification for each cluster.

[insert summ stats of how often it works correctly in the classified sample]

[54] In the case of ties, we choose a random six digit HS code from the most options deemed most frequent.

D Russian Procurement Totals Over Time

Table D.3: Total Procurement in Russia By Type of Mechanism Used

	2011	%	2012	%	2013	%	2014	%	2015	%	2011-2015	%
Electronic Auctions	2,242.43	46.4	3,333.64	54.6	3,418.72	58.0	2,849.01	51.7	2,809.82	51.5	14,653.62	52.7
Single Supplier	1,154.82	23.9	1,330.29	21.8	1,256.14	21.3	974.81	17.7	1,197.50	21.9	5,913.56	21.3
Request for Quotations	178.18	3.7	174.17	2.9	170.08	2.9	65.35	1.2	56.30	1.0	644.07	2.3
Open Tender	898.21	18.6	1,263.04	20.7	1,044.36	17.7	1,332.28	24.2	979.20	17.9	5,517.08	19.8
Other Methods	357.21	7.4	5.99	0.1	5.44	0.1	285.87	5.2	415.64	7.6	1,070.16	3.8
Total Procurement	4,832.86		6,109.13		5,896.75		5,509.34		5,460.47		27,808.56	
Russian Non-Resource GDP	42,950.27		51,150.31		54,453.02		60,198.96		66,250.61		275,003.16	
Procurement / Non-Resource GDP (%)	11.3		11.9		10.8		9.2		8.2		10.1	

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 rubles. 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).