

NBER WORKING PAPER SERIES

INDIVIDUALS AND ORGANIZATIONS AS SOURCES OF STATE EFFECTIVENESS, AND
CONSEQUENCES FOR POLICY

Michael Carlos Best
Jonas Hjort
David Szakonyi

Working Paper 23350
<http://www.nber.org/papers/w23350>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2017

We thank John M. Abowd, Daron Acemoglu, Nick Bloom, Richard Blundell, Raj Chetty, Allan Collard-Wexler, Francesco Decarolis, Pascaline Dupas, Ben Faber, Fred Finan, Matthew Gentzkow, Josh Gottlieb, Caroline Hoxby, Amit Khandelwal, Brad Larsen, Torsten Persson, Andrea Prat, Imran Rasul, Jimmy Roberts, Jonah Rockoff, Ori Shelef, Andrei Shleifer, Eric Verhoogen, 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, University of British Columbia, Central European University, Columbia, University of Copenhagen, Duke, the Econometric Society, Higher School of Economics, McGill, University of Michigan, NBER Development Economics Program meeting, NEUDC, Northwestern, University of Oslo, 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". Any remaining errors are ours alone. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Michael Carlos Best, Jonas Hjort, and David Szakonyi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Individuals and Organizations as Sources of State Effectiveness, and Consequences for Policy
Michael Carlos Best, Jonas Hjort, and David Szakonyi
NBER Working Paper No. 23350
April 2017
JEL No. H11,O2

ABSTRACT

How much of the variation in state effectiveness is due to the individuals and organizations responsible for implementing policy? We investigate this question and its implications for policy design in the context of public procurement, using a text-based product classification method to measure bureaucratic output. We show that effective procurers lower bid preparation/submission costs, and that 60% of within-product purchase-price variation across 16 million purchases in Russia in 2011-2015 is due to the bureaucrats and organizations administering procurement processes. This has dramatic policy consequences. To illustrate these, we study a ubiquitous procurement policy: bid preferences for favored firms (here domestic manufacturers). The policy decreases overall entry and increases prices when procurers are effective, but has the opposite impact with ineffective procurers, as predicted by a simple endogenous-entry model of procurement. Our results imply that the state's often overlooked bureaucratic tier is critical for effectiveness and the make-up of optimal policies.

Michael Carlos Best
Stanford Institute for Economic Policy Research
John A. & Cynthia Fry Gunn Building
366 Galvez Street
Stanford
CA 94305
mbest@stanford.edu

David Szakonyi
Department of Political Science
George Washington University
2115 G St. NW
Suite 440
Washington, DC 20052
dszakonyi@gwu.edu

Jonas Hjort
Graduate School of Business
Columbia University
3022 Broadway, Uris Hall 622
New York, NY 10027
and NBER
hjort@columbia.edu

A online appendix is available at <http://www.nber.org/data-appendix/w23350>

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). The NREGA employment guarantee scheme supports poor workers and helps complete important infrastructure projects in some Indian states, but is largely unused in others (Gulzar & Pasquale, 2017). 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 workers and organizations that implement state 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 state—prices paid for goods procured—and use administrative data covering the universe of public procurement in Russia from 2011 through 2015. We show that under a standard policy regime that treats all suppliers equally, 60% of the variation in prices paid is attributable to the bureaucratic apparatus. These differences in effectiveness across procurers can fundamentally alter the impact of policy changes. When a policy favoring domestic manufacturers is introduced, supplier entry and purchase prices improve in procurement processes run by ineffective procurers, but worsen in processes run by effective procurers. Our results demonstrate that increasing bureaucratic effectiveness could save the state billions of dollars every year, but also that policies tailored to the effectiveness of the individuals and organizations running the state enterprise can act as a partial substitute for improving effectiveness.

Public procurement in Russia is an ideal setting in which to study micro level state effectiveness for two reasons. First, procurement, which makes up roughly 8 percent of worldwide GDP (Schapper *et al.*, 2009), is one of the few state activities whose output is relatively well-defined, measurable, and comparable across the entire public sector.² Second, Russia is unusual among low- and middle-income countries in that detailed datasets spanning its massive and diverse bureaucracy are available.

There are three parts to our empirical analysis. We start by developing a text-based machine learning method that assigns procurement purchases of off-the-shelf goods 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 (e.g. ministries, schools or hospitals) and vice versa. This provides us with thousands

¹See for example Dal Bo *et al.* (2013); Duflo *et al.* (2013); Ashraf *et al.* (2014a,b); Hanna & Wang (2015); Callen *et al.* (2016a,b); Deserranno (2016); Bertrand *et al.* (2016). The literature has focused mostly on front-line public sector workers; for an excellent overview, see Finan *et al.* (forthcoming). On variation in state capacity, see, among many others, La Porta *et al.* (1999); Besley & Persson (2009, 2010); Acemoglu *et al.* (2015); Bai & Jia (2016); Dell *et al.* (2017).

²Newly available procurement data are therefore being used to investigate key open questions on state effectiveness. Like Bandiera *et al.* (2009) and Ferraz *et al.* (2015) we focus on purchases of items that are precisely defined (“off-the-shelf” goods), for which procurers’ mandate is simply to acquire the items at the lowest possible price while following the government’s policy rules. Such purchases make up over half of total public procurement in Russia (see Online Appendix Table OA.3). See also Lewis-Faupel *et al.* (2015); Coviello & Gagliarducci (forthcoming); Coviello *et al.* (2017).

of quasi-experiments that we can use to estimate the causal effect of specific individuals and organizations on prices paid. To do so, we combine the variance decomposition method of *Abowd et al.* (1999, 2002) with split-sample and shrinkage tools to correct for sampling error (*Finkelstein et al.*, 2016; *Kane & Staiger*, 2008). In the third part of the paper, instead of holding the policy environment constant and studying variation in output across procurers, we hold the bureaucrats and organizations constant and vary whether a particular procurement policy—bid preferences for domestic manufacturers—applies. Using the fact that the policy applies to an evolving set of goods and is in effect only parts of each year in a difference-in-differences approach, we estimate its average impact on prices paid, supplier entry and the share of contracts going to domestic firms. We then show how the impact of the policy depend on the effectiveness of the implementing bureaucrat and organization.

To guide our empirical analysis, we develop a stylized model of public procurement auctions. In the model, suppliers decide whether to enter an auction by trading off their expected profits against the costs imposed on auction participants by the bureaucracy. Ineffective bureaucracies impose high participation costs and as a result attract fewer bidders and pay higher prices. Introducing bid preferences for domestically manufactured goods, which we assume to be more expensive to produce on average, encourages entry by domestic manufacturers but discourages entry by foreign manufacturers. 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 net effect is the opposite.

We present four main empirical findings. First, we show that the individuals and organizations in the bureaucratic apparatus together account for 60 percent of the variation in the prices the government pays for its inputs in a standard policy regime that treats all suppliers equally. Bureaucrats and organizations each account for roughly equal shares of this overall effect. A battery of tests gives no indication that the log-linearity and “conditional random mobility” assumptions needed to interpret our estimated bureaucrat and organization effects causally are violated.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 observed performance range, improving the performance of existing bureaucrats, or lifting organization-wide characteristics such as management or “organizational culture” towards the high end of the observed range. Our estimates imply, for example, that moving the worst-performing 20 percent of bureaucrats and organizations to 50th percentile–effectiveness would save the Russian government 37.3 percent of the approximately USD 185 billion it spends on procurement annually.

Second, we show what effective procurers do differently. We correlate the effectiveness estimates with a rich set of indicators on how different procurers’ purchases are conducted and their intermediate outcomes. We find that effective bureaucrats and organizations impose lower entry costs on suppliers, as predicted by our conceptual framework, and that they attract different types of entrants.⁴

³The same tends to hold in the labor economics literature on workers and firms in the private sector (see e.g. *Mendes et al.*, 2010; *Card et al.*, 2013, 2015; *Goldschmidt & Schmieder*, 2015; *Shelef & Nugyen-Chyung*, 2015; *Alvarez et al.*, 2016; *Bloom et al.*, 2016, 2017). In the public sector there are additional institutional reasons to expect these assumptions to hold.

⁴By shedding light on the channels through which successful auctioneers in the public sector achieve lower prices, this evidence complements an innovative paper by *Lacetera et al.* (2016) studying auctioneers in U.S. used-car wholesale auctions.

Third, our difference-in-differences analysis of Russia’s “buy local” policy shows that, on average, the asymmetric procurement rules procurers are asked to implement achieve the government goal of channeling demand to domestic manufacturers. They do so at no cost to the government in that average prices paid are unaffected. In this sense, our results suggest that industrial policies of the form we study may be more successful in countries with low average bureaucratic effectiveness.⁵

Fourth, interacting the “buy local” policy with our estimates of bureaucrats’ and organizations’ effectiveness reveals that the average treatment effect masks considerable heterogeneity across the range of policy implementer effectiveness. The prices achieved by ineffective procurers *decrease* by about 15 percent when preferences apply, while the prices achieved by effective procurers increase by a similar magnitude. As our conceptual framework shows, heterogeneous policy effects across the range of procurer effectiveness can arise through the same underlying channel as the variation in general procurement effectiveness: individuals’ and organizations’ influence on entry costs. Intuitively, tilting the playing field in the auction in favor of local manufacturers is particularly effective at encouraging their entry when baseline participation is lowest since this is when a new entrant faces fewest incumbent competitors. Overall our results indicate that whether the benefit of distorting competition in public procurement—increased entry by favored firms—outweighs the costs fundamentally depends on how effective procurers are at attracting bidders. More generally, the results show that the design of optimal procurement policy depends on the effectiveness of the procurers at implementing policy.

This paper contributes to the literatures on the causes and consequences of state effectiveness; context-specific policy design; and methods for estimating productivity. We demonstrate that state effectiveness is to a considerable extent embodied in the bureaucratic apparatus. In this sense our paper is most closely related to an innovative study by [Bertrand *et al.* \(2016\)](#), which analyzes how the incentives of elite bureaucrats in India matter for perceived performance and aggregate outcomes, and the literature on the role of management practices in public organizations (see e.g. [Rasul & Rogger, forthcoming](#); [Bloom *et al.*, 2015a,b,c](#)).⁶ The approach we take—decomposing the *total variation* across the entire public sector in *a particular form* of public sector output—has three advantages. First, it allows us to compare workers and organizations pursuing a single objective, assuaging concerns that state actors that perform well in one dimension may perform worse in unobserved dimensions. Second, it allows us to quantify policy implementers’ influence on public sector output *relative to that of other determinants*. Finally, our approach yields measures of individuals’ and organizations’ task-specific effectiveness that can be used to study how optimal policy design depends on bureaucratic effectiveness.⁷

⁵The average treatment effect of Russia’s “buy local” program contrasts with the effect of similar policies found in higher state effectiveness contexts. For example, a five percent bid preference for small businesses in Californian road construction procurement is estimated to increase average costs by 1–4 percent ([Marion, 2007](#); [Krasnokutskaya & Seim, 2011](#)).

⁶[Burgess *et al.* \(2012\)](#); [Duflo *et al.* \(2013\)](#); [Khan *et al.* \(2016\)](#); [Callen *et al.* \(2016a,b\)](#); [Deserranno \(2016\)](#) also present important evidence on how performance incentives affect public sector workers’ performance. There is also an influential literature on the characteristics of front-line public sector workers and public goods provision (see footnote 1).

⁷In this sense our paper is related to [Yao & Zhang \(2015\)](#) who estimate the share of the variance in cities’ economic growth in China attributable to mayors. Their study belongs to an important body of work analyzing how public sector managers and politicians matter for aggregate economic outcomes ([Jones & Olken, 2005](#); [Bertrand *et al.*, 2016](#); [Xu, 2017](#)). While it is difficult to rule out that leaders who improve aggregate economic outcomes perform worse in other dimensions and to compare leaders’ influence on aggregate outcomes to that of other determinants, aggregate economic outcomes are uniquely important.

We show that the type of policy we focus on—“preferencing” a specific group of firms—is a much less costly way to steer government demand towards domestic manufacturers when the entities in charge of procurement are *less* effective. This part of our analysis builds on recent work starting to unpack how policy should be tailored to context (Laffont, 2005; Duflo *et al.*, 2016; Best *et al.*, 2015; Hansman *et al.*, 2016). We extend this literature by considering a new dimension of context—the individuals and organizations that implement government policy; by focusing on a state activity—procurement—that allows us to hold the exact task that the entities we compare engage in constant⁸; and by considering the range of state effectiveness observed across Russia’s entire public sector.

The type of procurement policy we focus on is common world-wide and extensively studied in U.S. settings.⁹ However, because the Russian version of the policy is turned off parts of each year and applies to a subset of goods, we can quasi-experimentally identify its average treatment effect (ATE), as well as conditional treatment effects (CTEs) by policy implementer effectiveness. In decomposing an estimated ATE into CTEs that are specific to individual entities of a given level of a particular trait, we follow the literature on heterogeneous treatment effects (see e.g. Heckman & Smith, 1997; Angrist, 2004; Deaton, 2010; Heckman, 2010). However, we are not aware of previous studies that consider treatment effects that condition on an unobserved (and therefore estimated) characteristic such as effectiveness. Our approach represents a new way to use an ATE of a government policy estimated in one setting to predict the effects of the policy in another (see also Vivalt, 2016; Dehejia *et al.*, 2016).¹⁰

Finally, we develop a new approach to measuring workers’ and organizations’ effectiveness, or productivity. Our starting point is the recognition that increasingly available types of data allow researchers to (a) construct direct measures of productivity, and (b) to partition the exact activities workers and organizations engage in into comparable categories. We build on 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.¹¹ We show that the AKM method can be used to estimate measures of productivity that are free of the confounds that arise from existing productivity-estimation methods’ (i) comparison of workers and/or organizations (e.g. firms) engaged in different tasks and (ii) use of wage and profits data as outcomes.¹²

⁸We follow novel papers by Mironov & Zhuravskaya (2016) and Andreyanov *et al.* (2016) in studying procurement in Russia. They focus on procurement contracts inferred from observed bank transfers and awards of large contracts through “subjective” channels respectively. We instead focus on purchases made via the blind, anonymous, descending auctions used for purchases of off-the-shelf goods in Russia.

⁹See McAfee & McMillan (1989); Marion (2007); Krasnokutskaya & Seim (2011); Athey *et al.* (2013); Bhattacharya *et al.* (2014) on asymmetric auction rules and theory’s ambiguous prediction for how prices respond, and, among others, Samuelson (1985); Levin & Smith (1994); Bulow & Klemperer (1996); Gentry & Li (2014); Branzoli & Decarolis (2015) on participation costs, entry, and prices achieved in auctions more generally.

¹⁰The heterogeneity in CTEs we find resonates with recent studies comparing program effects across branches of companies or private-versus-public status of the implementing agency (see Bold *et al.*, 2013; Allcott, 2015).

¹¹Abowd *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.

¹²Wages do not necessarily reflect productivity (Eckhout & Kircher, 2011; Card *et al.*, 2015), but are important objects in and of themselves. Existing applications of the AKM method have used samples that include workers performing many different tasks. Carneiro *et al.* (2012) and Cardoso *et al.* (2016) show the potential importance of accounting for differences in tasks. On the organization/firm side, conventional methods estimate productivity from revenue or profits data and thus risk conflating productivity with mark-ups and quality differentiation (see e.g. Goldberg & De Loecker, 2014).

To apply our method in the Russian procurement setting, we first use machine learning methods and the text of procurement contracts to assign purchases to narrow product categories while maintaining generality by not restricting the sample to very specific types of goods. This allows us to compare procurers purchasing the same good.¹³ We then show how fixed-effect estimates of individual and organization effects can be corrected for sampling error (Neyman & Scott, 1948; Lancaster, 2000). Specifically, we adapt split-sample (see e.g. Finkelstein *et al.*, 2016; Silver, 2016) and “shrinkage” (Kane & Staiger, 2008; Chetty *et al.*, 2014a) methods to a two-dimensional fixed effects context.¹⁴ Finally, we show that even in settings in which worker (and/or organization) mobility does not link all organizations—a challenge that will arise in many settings where AKM can be used to estimate productivity—(i) a suitable normalization of the fixed effects allows us to estimate a lower bound on the shares of the variance in productivity explained by workers and organizations and (ii) that the combined productivity effect of pairs of workers and organizations can nevertheless be identified.¹⁵

The rest of the paper is organized as follows. Section 2 presents an endogenous entry procurement model with variation in bureaucratic effectiveness that guides our analysis. Background on the Russian public procurement 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 sell an item to the government. The government uses an auction to award the contract and determine the price paid. 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. Then, in Sub-section 2.2, we show how introducing bid preferences for favored suppliers can have opposite consequences depending on whether state effectiveness is high or low.

¹³The method we develop ensures that within-category quality differences are minimal. The difficulty of categorizing goods accurately so as to ensure like for like comparisons has long dogged several literatures. In foregoing conventional methods and instead using text analysis to classify goods, our study relates to Hoberg & Phillips (2016). Their text classification method has similarities to ours, but differs in that they classify *firm* similarity based on text listing the various goods firms produce, whereas we classify *good* similarity based on text listing words describing items. We also carry out a battery of tests that relax the within-category homogeneity assumption.

¹⁴To our knowledge, two-dimensional shrinkage estimators like the ones we develop have not been used before.

¹⁵In the type of setting/data AKM has been applied to in the existing literature—private sector wages—worker mobility is often high enough that almost all workers and firms belong to the biggest connected set, particularly if workers and firms engaged in different tasks are compared. This paper is to our knowledge the first to estimate effects attributable to workers and organizations in the *public* sector, where mobility is lower. Limited connectivity issues likely arise in many settings where our productivity-estimation method can be fruitfully applied, e.g. where some, but not all, firms engage in the same activity.

2.1 A simple model of procurement 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 maximum 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.

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 $0 < \mu < 1$.¹⁶ 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 maximum allowable price of 1. If neither supplier chose to enter, the procurer randomly picks a supplier and awards her the contract at a price of 1.¹⁷ Finally, if both suppliers enter, the auction takes place.

The 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. Denoting the bidding strategy of supplier i with fulfillment cost x by $b_i(x)$, we have $b_F(x) = b_L(x) = x$ (see e.g. Milgrom, 2004; Krishna, 2010). 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 \bar{d}_F , and bidder L entering if her fulfillment cost is below a threshold \bar{d}_L .¹⁸

We outline the equilibrium here, relegating a detailed characterization and the proofs of propositions to Online Appendix OA. 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 \bar{d}_F, \bar{d}_L as

$$U_i(v; \bar{d}_F, \bar{d}_L) = m_i(v; \bar{d}_F, \bar{d}_L) - q_i(v; \bar{d}_F, \bar{d}_L) v \quad (1)$$

where $m_i(v; \bar{d}_F, \bar{d}_L)$ is the expected payment supplier i receives if she enters with fulfillment cost v , and $q_i(v; \bar{d}_F, \bar{d}_L)$ is the probability that supplier i receives the contract if she enters when her fulfillment cost

¹⁶The positions of the upper and lower bounds of the distribution are innocuous. Uniformity is a simplifying assumption that allows us to derive simple, closed-form expressions.

¹⁷A more realistic assumption would be that auctions in which no firms enter have to be re-run at some cost. Our assumption makes the model static, simplifying the exposition. The qualitative results are unlikely to depend on this choice.

¹⁸This 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.

is v . The probabilities of winning are given by

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

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 as¹⁹

$$U_i(v; \bar{d}_F, \bar{d}_L) = \int_v^1 q_i(x; \bar{d}_F, \bar{d}_L) dx$$

$$m_i(v; \bar{d}_F, \bar{d}_L) = \int_v^1 q_i(x; \bar{d}_F, \bar{d}_L) dx + q_i(v; \bar{d}_F, \bar{d}_L) v$$

The entry thresholds are pinned down by 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

$$U_F(\bar{d}_F; \bar{d}_F, \bar{d}_L) - c = \frac{1}{2}(1 - \bar{d}_F) \frac{1 - \bar{d}_L}{1 - \mu} \quad (3)$$

$$U_L(\bar{d}_L; \bar{d}_F, \bar{d}_L) - c = \frac{1}{2}(1 - \bar{d}_F)(1 - \bar{d}_L) \quad (4)$$

In this equilibrium, the expected number of entrants is

$$E[n] = G_F(\bar{d}_F) + G_L(\bar{d}_L) \quad (5)$$

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

$$E[p] = E_{v_F}[m_F(v_F; \bar{d}_F, \bar{d}_L)] + E_{v_L}[m_L(v_L; \bar{d}_F, \bar{d}_L)] + [1 - G_F(\bar{d}_F)][1 - G_L(\bar{d}_L)] \quad (6)$$

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

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

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

$$\frac{dE[n]}{dc} < 0 \quad \& \quad \frac{dE[p]}{dc} > 0 \quad (7)$$

Proof. See Online Appendix OA.2. □

2.2 Introducing bid preferences for favored firms

In the previous sub-section, while the suppliers were asymmetric, the government treated them symmetrically. In this sub-section, we introduce bid preferences favoring domestic manufacturers. Specifically,

¹⁹Strictly, $U_i(v; \bar{d}_F, \bar{d}_L) = U_i(1; \bar{d}_F, \bar{d}_L) + \int_v^1 q_i(x; \bar{d}_F, \bar{d}_L) dx$. However, since a supplier with fulfillment cost of 1 never makes a profit, $U_i(1; \bar{d}_F, \bar{d}_L) = 0$.

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 still to bid her true value $b_L(x) = x$.

In this case, the probability of winning is

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

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

but otherwise all the steps in characterizing the equilibrium are as before (see Online Appendix OA.3).

The following proposition summarizes the impact of introducing bid preferences favoring local products, emphasizing how the effects depend on the entry costs procurers impose on sellers.

Proposition 2. *Bid preferences favoring local manufacturers have opposite effects for procurers who impose high and low entry costs. Preferences lead procurers who impose high entry costs to attract more bidders and pay lower prices, but lead low entry cost procurers 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) Let $p(c, \gamma)$ be the expected price when preferences are given by $\gamma \in (0, 1]$ and participation costs are $c \in [0, \bar{c}]$, and let $\bar{\gamma}_p = \arg \min_{\gamma} p(\bar{c}, \gamma) < 1$ be the γ that minimizes prices for the buyer with the highest participation cost \bar{c} . Then for every $\gamma \in (\bar{\gamma}_p, 1)$, there exists a $\tilde{c}_p(\gamma) \in [0, \bar{c}]$ such that $p(c, \gamma) - p(c, 1) < 0$ for all $c > \tilde{c}_p(\gamma)$ and $p(c, \gamma) - p(c, 1) > 0$ for all $c < \tilde{c}_p(\gamma)$.

Similarly, let $n(c, \gamma)$ be the expected number of participants when preferences are given by $\gamma \in (0, 1]$ and participation costs are $c \in [0, \bar{c}]$, and let $\bar{\gamma}_n = \arg \max_{\gamma} n(\bar{c}, \gamma) < 1$ be the γ that maximizes participation for the buyer with the highest participation cost \bar{c} . Then for every $\gamma \in (\bar{\gamma}_n, 1)$, there exists a $\tilde{c}_n(\gamma) \in [0, \bar{c}]$ such that $n(c, \gamma) - n(c, 1) > 0$ for all $c > \tilde{c}_n(\gamma)$ and $n(c, \gamma) - n(c, 1) < 0$ for all $c < \tilde{c}_n(\gamma)$.

Moreover, (ii)

$$\frac{\partial p(c, \gamma) - p(c, 1)}{\partial c} < 0 \quad \& \quad \frac{\partial n(c, \gamma) - n(c, 1)}{\partial c} > 0$$

Proof. See Online Appendix OA.4. □

Intuitively, without preferences ($\gamma = 1$), high entry cost procurers depress entry and hence raise prices. They do so particularly for local manufacturers, since local manufacturers tend to have higher fulfillment costs and hence lower expected profits from participation.²⁰ Then, when preferences are introduced, this lowers expected profits for foreign manufacturers 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 entry cost procurers, who

²⁰Formally, we show in Online Appendix OA.3 that the entry thresholds of the foreign and local bidder satisfy $\bar{d}_F - \gamma \bar{d}_L = \sqrt{2\gamma c \mu}$. Hence, the gap between the foreign and local bidders' entry thresholds is increasing in c .

were suppressing entry by local bidders the most in the absence of preferences. For high cost procurers, the net effect is to increase participation and lower prices. Conversely, for low cost procurers, the net effect is to decrease participation and increase prices.

3 Public Procurement 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.²¹ Each government entity 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 in 1997, and on January 1, 2006, the procurement rules and regulations governing tender processes at all levels of government were harmonized by Federal Law No. 94-FZ (Yakovlev *et al.*, 2010; Krylova & Settles, 2011). New provisions assigned criminal and administrative liability for individuals and legal entities violating anti-monopoly legislation. 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. This website is our main data source.

3.2 Procurement of off-the-shelf goods through auctions

Public procurement makes up about 10 percent of Russia's non-resource GDP. We restrict attention to purchases of off-the-shelf goods through electronic auctions because it is then possible to compare procurers purchasing the exact same good (after applying our good classification method), and because bureaucrats and organizations may affect procurement outcomes along multiple dimensions in more subjective procurement mechanisms, making comparison difficult. Electronic (open) auctions are used for 53.5 percent of Russian procurement during our data period.²²

Since July 10, 2010, all auctions have been conducted through one of five designated websites. 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 out the steps involved in a procurement process; we now go through these.

Each purchase starts with an auction announcement. Our data, described in detail in Section 4, contains 5,054,498 announcements. The announcement contains technical details on the item(s) to be pur-

²¹The Soviet Union, like other socialist states (see e.g. Bai & Jia, 2016), operated a centralized bureaucracy (see e.g. Chere-mukhin *et al.*, 2016). Since 1991, the Russian bureaucracy has become very decentralized (Enikolopov & Zhuravskaya, 2007).

²²The other three main procurement mechanisms are open tenders (19.8 percent), open requests for quotations (2.3 percent), and single source procurement (21.3 percent). Online Appendix Table OA.3 shows usage of these methods over time.

chased, a maximum allowable price for the lot, the required security deposit (between 0.5 and 5 percent of the maximum price), other participation requirements, and the date of the electronic auction. In order to participate in an auction, suppliers must first obtain accreditation. This requires not being in a state of bankruptcy, not being sanctioned under administrative law, not having substantial unpaid taxes, and not being listed in the registry of suppliers who have committed violations of procurement rules during the last two years. To participate, suppliers must also submit the security deposit. Finally, suppliers must prepare a formal application, consisting of two parts. The first part describes the good that they are offering to fulfill the procurement order. The second part contains information on the supplier itself (name, address, etc.). Importantly, until the auction is concluded only the electronic trading platform has access to the second part of the application and the identities of the suppliers.

A five-member procuring commission designated to oversee the auction receives and evaluates the first part of the application before the auction is held. Applications to participate in auctions are denied if the supplier is not accredited, cannot pay the security deposit, or if its proposal does not comply with the requested item specifications. In the event that only one supplier is approved to participate, the auction is declared “not held”, the procuring commission receives the second part of the supplier’s application, and a contract is drawn up with that supplier at the maximum allowable price. This is a relatively common occurrence; in 1,344,825 cases, or 27 percent of the purchases we observe, 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. The eligible suppliers remain anonymous and are each assigned a participant number. All participants log in to the online platform and participate in a descending second-price auction. When a participant enters a bid lower than the current winning bid, information on the amount of the 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. The procuring commission checks the applications to make sure the suppliers’ accreditations, licenses, 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.

3.3 The role of bureaucrats and organizations in procurement

Public procurement purchases are made by and for a public sector entity that we will refer to as an *organization*. The 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 the municipal, regional or federal level. 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 be either in-house employees of the organization or employees of an external public agency whose bureaucrats conduct procurement auctions with and for multiple organizations. Each regional authority sets rules dictating that the organizations under its jurisdiction use either an in-house or external bureaucrat for a given type of purchase as defined by the maximum allowable price of the contract and the nature of the item. In some cases such rules change over time.²³

After Russia declared independence in 1991, the Soviet Union's network of civil service academies collapsed (Huskey, 2004), leaving academies to fend for themselves in a new market for higher education. As a result, the education and labor market for procurement bureaucrats is extremely decentralized.²⁴ Individuals interested in working in public procurement seek out educational and employment opportunities much as they would in the private sector. 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 acts) governing procurement. In all cases we are aware of, the procurement bureaucrats are paid a flat salary.

Since 2014, the division of labor between a procuring organization and a potential external bureaucrat has been specified by law. The organization must submit all technical documentation, and choose and justify a maximum allowable price. After this documentation is posted online, the organization and bureaucrat together designate the procuring commission to oversee the auction process. 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.²⁵ 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 purchases 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 final outcomes.

Bureaucrats and organizations engaged in procurement of off-the-shelf goods have one simple mandate: to acquire the items at the lowest possible price while following the government's policy rules. Other procurement policy goals the government may have—e.g. influencing which type of firms win contracts—manifest themselves in the policy rules the procurers are asked to follow.

Both public procurement and Russia are generally associated with corruption (OECD, 2016; Transparency International, 2016; Szakonyi, 2016). Corruption in procurement can take many forms, but al-

²³External procurement agencies can be organized by a given authority (for example an education or health ministry/department), at the federal, regional, or municipal level. Part of the motivation for allowing the creation of 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 decentralized management of procurement in Russia and coordination required to co-purchase means that such joint purchases are very rare. Note that we control for the factors that authorities with an external procurement agency use to determine whether an item can be purchased by an in-house bureaucrat or must be purchased by external bureaucrats—the type of good and/or maximum allowable price of the contract—in our empirical analysis below.

²⁴The 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). Examples of private academies offering trainings in the procurement sector include ArtAleks <http://artaleks.ru/> and the Granit Center <http://www.granit.ru/>.

²⁵The one exception to this are “Kazennyie organizations”, which can delegate the process to a centralized bureaucrat.

most all of them will result in a higher purchase price for the government, and as such will be captured in our measure of bureaucratic effectiveness. For example, collusion between potential suppliers, or between potential suppliers and procurers, is likely to manifest itself in the price ultimately paid for the items procured. If such collusion is associated with specific procurers, either because they are especially corrupt or because suppliers take advantage of certain procurers, our empirical procedure will assign a lower effectiveness score to those procurers.²⁶

Our data allow us to directly address the two forms of corruption or incompetence that could undermine our estimates of bureaucratic effectiveness. First, it is possible that procurers who achieve low prices systematically purchase poor quality goods. In sections 4 and 5 we therefore develop and apply a methodology designed to ensure that within-category quality differences are minimal. We also carry out a battery of tests that relax the within-category homogeneity assumption.

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

In summary, one form of corruption or incompetence that would invalidate our effectiveness estimates—quality differences—can be directly addressed using text fields available on each item in our data, and the other—non-delivery—is observable and very rare. We thus believe that our procurer effectiveness estimates capture what a government cares about in the first instance: the price paid for goods of specified quality that are satisfactorily delivered.²⁸ Throughout the paper we remain agnostic about the extent to which some procurers pay higher prices than others because they are prone to forms of corruption that manifest themselves in the prices paid and the extent to which they do so because they are of lower ability.²⁹ In the framework in Section 2, the two sources of supplier entry costs have the same

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

²⁷Less than one percent of the auctions in our sample suffered from “bad execution”. The data also include information on early and late delivery. As discussed in Section 5, these are correlated with estimated procurer effectiveness, but can only explain (or “offset”) a tiny fraction of the dispersion in price effectiveness we estimate.

²⁸From a longer term perspective, governments may also care whether the “right” firms win contracts, but this possibility is beyond the scope of our paper insofar as allocative objectives extend beyond a preference for domestic manufacturers.

²⁹*Bandiera et al.* (2009) find that 83 percent of waste in Italian public procurement purchases is due to low bureaucratic ability rather than corruption. They study purchases conducted via procurement mechanisms that allow a role for subjective judgment by the procurers. The reason why Russia’s modern-day procurement laws force procurers to use the blind, electronic auctions we study for purchases of off-the-shelf goods (see Online Appendix Table OA.3) is that the government believes that such auctions are harder to corrupt than procurement mechanisms that allow a role for subjective judgment, such as the “open request for quotations” studied by *Andreyanov et al.* (2016). In general there may of course be both advantages and disadvantages to allowing procurers flexibility in the allocation of contracts (*Hart et al.*, 1997; *Fisman & Gatti*, 2002; *Schargrodsky & Di Tella*, 2003; *Lewis-Faupel et al.*, 2015; *Meng et al.*, 2015; *Andreyanov et al.*, 2016; *Duflo et al.*, 2016; *Mironov*

impact on the main equilibrium outcomes we focus on, namely supplier participation and the price paid.

3.4 Preferences for domestically manufactured 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. Firms that produce their goods in Russia received a 15 percent bid discount for parts of 2011 through 2015.

The “buy local” regime worked as follows. Each year from 2011 to 2014 a list of goods for which preferences for domestic manufacturers was to apply was drawn up. The presidential order defining this list was passed in May or June and remained in effect for the remaining part of the calendar year, 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 only time during the period a list was in effect for more than a year. As such, preferences were never applied to purchases conducted the first period of each year from 2010-2014. Organizations filing procurement requests for any goods on this annual list were required to publicly inform potential suppliers that preferences would be applied. Preferred goods spanned numerous categories, including automobiles, clocks, various types of food products, medical equipment, pharmaceuticals, and textile and furs. The country of origin of a good was defined as the country where the good was “completely produced”, or where it underwent “significant reprocessing”.

For the preferences to apply, 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.

4 Data

4.1 Procurement auctions data

Our data on public procurement auctions and final contracts comes from the Unified Register of Federal and Municipal Contracts located at <http://zakupki.gov.ru/>. We use data on the universe of electronic auction requests, review protocols, auction protocols, and contracts from January 1, 2011 through December 31, 2015. In all, we have information on 5,054,498 auction requests. Figure 1 presents a mapping between our data and the sequence of procurement procedures described in Sub-section 3.2.

A great deal of previous research in economics has faced the challenge of assigning items to good categories so as to ensure that quality and other differences within categories are minimal. Broadly, three approaches have been taken: using hedonic regressions to estimate consumers’ demand for and/or suppliers’ costs of producing good attributes (Griliches, 1971; Rosen, 1974; Epple, 1987); using product codes provided by e.g. customs agencies to partition goods (Rauch, 1999; Schott, 2004); or restricting

& Zhuravskaya, 2016; Coviello *et al.*, forthcoming).

attention to products that are by nature especially homogeneous (see e.g. Syverson, 2004; Hortacsu & Syverson, 2007; Foster *et al.*, 2009). However, these approaches typically achieve good homogeneity at the cost of losing generality.³⁰ With our data, we are in the common situation that our most detailed information on the goods is in unstructured text fields in the contracts procurers sign with suppliers rather than encoded into product codes.³¹ We thus use text analysis methods from the machine learning literature to assign goods to homogeneous categories (see also Gentzkow & Shapiro, 2010; Hansen *et al.*, 2014; Hoberg & Phillips, 2016).

Our method proceeds in three steps. First, we transform the good descriptions in our contract data into vectors of word tokens to be used as input data in subsequent steps. Second, we develop a transfer learning procedure. The procedure uses good descriptions and their corresponding 10-digit Harmonized System product codes in data on the universe of Russian 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 “width” to the categories created in the second step. Details are in Appendix OA.6.

4.2 Procurement of pharmaceuticals data

We collect additional data on procurement requests for pharmaceuticals, a type of good that is by nature homogeneous and where items’ country of origin can be inferred using brand names (Bronnenberg *et al.*, 2015). Russia’s government regulates the pharmaceutical market to ensure that certain drugs are available and affordable, compelling manufacturers of these drugs to register them in a List of Vital and Essential Medicinal Drugs (LVEMD). This list includes information on each drug’s International Nonproprietary Name (INN)³²; 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 (trademark) name, active ingredient (INN), dosage (mg, g, mkg), active units (IU), concentration (mg/ml, mkg/ml), volume (ml), and units (tablets, packages). This matching allows us to construct barcode-level identifiers for drugs that we can use as alternatives to our text classification good categories, and to identify the manufacturer (and thus country of origin) for each pharmaceutical procured. We restrict the pharmaceuticals sample to purchases of drugs we can match to LVEMD.³³

³⁰Generality suffers both from restricting attention to very specific types of goods, and, in a methodological sense, from use of a method that is successful at creating homogeneous categories only being possible with unusual types of data.

³¹Note that we use the contract text since it describes the goods *purchased*—the procurer’s output—rather than text from requests, which are one of the procurer’s inputs into the purchase.

³²INN is a globally recognized term to denote the chemical substance of the medicine, see <http://www.who.int/medicines/services/inn/en/>, accessed October 8, 2015.

³³Cases in which we are unable to match a drug to the LVEMD can arise both because the medicine is not classified by the Russian government as “essential” (i.e., covered by the LVEMD) and because sufficient information on dosage and quantity is not available in the procurement contract.

4.3 Supplier firm data

We also collected data on all firms that participate in any stage of the procurement process. The primary dataset on firms is *Ruslana* from the Bureau Van Dijk (BVD) agency. *Ruslana* covers the vast majority of registered firms in Russia that file financial information. All firms are by law required to submit accounting data on an active basis. All statistics are standardized by the Russian Ministry of Finance and provided to agencies such as BVD for dissemination to end-clients.

5 Individuals and Organizations as Sources of State Effectiveness

In a given policy environment—here, standard procurement rules that treat all suppliers equally—prices vary dramatically across purchases within good categories. We begin our empirical analysis by investigating how much of this variation in state effectiveness is due to the bureaucratic apparatus.

5.1 Empirical model

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

$$p_i = \mathbf{X}_i\beta + \tilde{\alpha}_{b(i,j)} + \tilde{\psi}_j + \varepsilon_i \quad (10)$$

where \mathbf{X}_i is a vector of item-level controls, including log quantity, good fixed effects³⁴, month fixed effects, and interactions between 2-digit HS product categories, years, region, and lot size³⁵; $\tilde{\alpha}_{b(i,j)}$ is the bureaucrat effect, $\tilde{\psi}_j$ is the organization effect; and ε_i is a residual. If bureaucrats and organizations are important drivers of prices achieved, then we expect $\text{Var}(\tilde{\alpha}_{b(i,j)})$ and $\text{Var}(\tilde{\psi}_j)$ to be large relative to the overall variance in prices.

In estimating and interpreting this empirical model, we face four challenges. First, identification of procurer effects is possible only within sets of organizations connected by bureaucrats moving between them. Second, estimated procurer effects can only be interpreted causally if mobility is conditionally random. Third, we can compare procurers engaged in the *same* task only if our method for measuring and defining the specific good purchased in each auction is adequate. Fourth, the estimated fixed effects represent an accurate measure of a procurer’s influence on auction outcomes only if we can appropriately address finite sample biases. We discuss the first two issues and then present our results; after presenting the results we return to the third and fourth issues.

³⁴Hereafter we refer to the goods categories constructed using the method described in Sub-section 4.1 as “goods”.

³⁵By lot size we mean the maximum allowable price for all the items to be purchased in a given auction. We divide the maximum allowable price into bins so as to allow our estimates of procurer effectiveness to capture the impact on prices of the procurers’ choice of the exact maximum price posted. The interactions help address, for example, concerns that systematic variation in the average prices of different types of goods across space, in combination with differences across bureaucrats and organizations in the items purchased, confound our estimates of bureaucrat and organization effectiveness. Russian regions are highly heterogeneous (Enikolopov & Zhuravskaya, 2007; Acemoglu *et al.*, 2011; Yakovlev & Zhuravskaya, 2014).

5.2 Connectivity and estimation

Separate identification of the bureaucrat and organization effects is made possible by the fact that some bureaucrats make purchases for multiple organizations, and some organizations use multiple bureaucrats to make purchases. 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 $\tilde{\alpha}_{b(i,j)}$ ’s and $\tilde{\psi}_j$ ’s. Within each connected set, the bureaucrat and organization effects are identified relative to other procurers in the set 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 $E \left[\text{Var} \left(\alpha_{b(i,j)} \right) | s \right]$, but not the corresponding unconditional variance, and similarly for the organization effects.

Faced with this issue, previous work on private sector workers and firms has tended to restrict attention to the largest connected set, normalizing an arbitrary firm effect to 0, and estimating unconditional variances.³⁶ However, due to the decentralized nature of public procurement in Russia, lower worker mobility in the public sector, and our focus on bureaucrats performing the same task, our data contains 28,147 connected sets, and the largest connected set contains only 10,854 of the 95,420 organizations in the full sample. To maintain generality and representativeness, we conduct our analysis in two samples. Our *analysis sample* contains all connected sets containing at least three bureaucrats and organizations after making the following restrictions. 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 pairs. We also require that all bureaucrats and organizations purchase at least five items. In our second sample we restrict attention to the largest connected set in the analysis sample. [Table 1](#) 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 auctions, as well as item-level characteristics, such as quantity and price per unit.

In the analysis sample, it is natural to normalize the bureaucrat and organization effects 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. We rewrite the model in [equation \(10\)](#) as

$$p_i = \mathbf{X}_i \boldsymbol{\beta} + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i \quad (11)$$

We use [\(11\)](#) to decompose the variation in prices into its constituent parts as follows

$$\begin{aligned} \text{Var}(p_i) = & \text{Var} \left(\alpha_{b(i,j)} \right) + \text{Var} \left(\psi_j \right) + \text{Var} \left(\gamma_{s(b,j)} \right) + 2\text{Cov} \left(\alpha_{b(i,j)}, \psi_j \right) + \text{Var} \left(\mathbf{X}_i \boldsymbol{\beta} \right) \\ & + 2\text{Cov} \left(\alpha_{b(i,j)} + \psi_j, \gamma_{s(b,j)} \right) + 2\text{Cov} \left(\gamma_{s(b,j)}, \mathbf{X}_i \boldsymbol{\beta} \right) + \text{Var} \left(\varepsilon_i \right) \end{aligned} \quad (12)$$

³⁶An 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.

As shown in Online Appendix OA.5, the effects in this augmented model are related to the underlying bureaucrat and organization effects by

$$\alpha_b = \tilde{\alpha}_b - \bar{\alpha}_{s(b)} \quad \psi_j = \tilde{\psi}_j - \bar{\psi}_{s(j)} \quad \gamma_{s(b,j)} = \bar{\alpha}_{s(b,j)} + \bar{\psi}_{s(b,j)}$$

where $\bar{\alpha}_{s(b)}$ is the mean bureaucrat effect in the connected set containing bureaucrat b , and similarly $\bar{\psi}_{s(j)}$ is the mean organization effect in organization j 's connected set. 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:

$$\begin{aligned} \text{Var}(\tilde{\alpha}_b) &\equiv \text{E}[\text{Var}(\tilde{\alpha}_b|s(b))] + \text{Var}(\text{E}[\tilde{\alpha}_b|s(b)]) \\ &= \text{Var}(\alpha_b) + \text{Var}(\text{E}[\tilde{\alpha}_b|s(b)]) \geq \text{Var}(\alpha_b) \end{aligned} \quad (13)$$

$$\text{Var}(\tilde{\psi}_j) = \text{Var}(\psi_j) + \text{Var}(\text{E}[\tilde{\psi}_j|s(j)]) \geq \text{Var}(\psi_j) \quad (14)$$

$$\begin{aligned} \text{Var}(\tilde{\alpha}_b + \tilde{\psi}_j) &\equiv \text{E}[\text{Var}(\tilde{\alpha}_b + \tilde{\psi}_j|s(b,j))] + \text{Var}(\text{E}[\tilde{\alpha}_b + \tilde{\psi}_j|s(b,j)]) \\ &= \text{Var}(\alpha_b + \psi_j) + \text{Var}(\gamma_{s(b,j)}) \end{aligned} \quad (15)$$

Equations (13)–(15) show that consistent estimates of the variances of the bureaucrat and organization effects in (11) provide lower bounds on the variances of the true bureaucrat and organization effects in (10), respectively, and that we can construct the variance of the total effect of bureaucrats and organizations using our estimated bureaucrat and organization effects and connected set intercepts.

5.3 Interpretation: do the estimated procurer effects represent causal effects?

Our variance decomposition method uses movements of organizations between bureaucrats and between goods, and movements of bureaucrats between organizations and goods, to identify how specific bureaucrats and organizations affect prices. Identification therefore relies on these movements being orthogonal to the error term in equation (11). To illustrate the possible sources of endogenous mobility, we follow Card *et al.* (2013) and write the error term as consisting of five random effects:

$$\varepsilon_i = \eta_{b(i,j)j} + \theta_{b(i,j)g} + \kappa_{jg} + \xi_{b(i,j)} + \zeta_j + \nu_i \quad (16)$$

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 pairs. $\xi_{b(i,j)}$ and ζ_j 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 specific organizations work especially well (or poorly) with specific bureaucrats. Similarly, it is possible that some 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 effectiveness of an organization or bureaucrat over time. Such drift could for example reflect the procurer

learning how to achieve low prices, or potential bidders learning about the desirability of participating in auctions managed by a particular procurer.³⁷ The transitory term captures any remaining components of the error term.

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

$$\mathbf{p} = \mathbf{X}\boldsymbol{\beta} + \mathbf{B}\boldsymbol{\alpha} + \mathbf{J}\boldsymbol{\psi} + \mathbf{S}\boldsymbol{\gamma} + \boldsymbol{\varepsilon} \quad (17)$$

where B is the $N \times N_b$ design matrix indicating the bureaucrat conducting each purchase $[b_1, b_2, \dots, b_{N_b}]$; J is the $N \times N_j$ design matrix indicating the organization purchasing each item $[j_1, j_2, \dots, j_{N_j}]$; and S is the design matrix of connected set dummies $[s_1, s_2, \dots, s_{N_s}]$. \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 (17) by OLS will then identify the effects $\boldsymbol{\alpha}$, $\boldsymbol{\psi}$, and $\boldsymbol{\gamma}$ under the following assumptions:

$$\text{E}[b'_b \boldsymbol{\varepsilon}] = 0 \forall b; \quad \text{E}[j'_j \boldsymbol{\varepsilon}] = 0 \forall j; \quad \text{E}[g'_g \boldsymbol{\varepsilon}] = 0 \forall i; \quad \text{E}[\tilde{\mathbf{X}}' \boldsymbol{\varepsilon}] = 0 \quad (18)$$

which together imply that $\text{E}[s'_s \boldsymbol{\varepsilon}] = 0 \forall s$.³⁸ 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 (18). Similarly, especially effective bureaucrats and organizations can specialize in the purchase of certain types of goods. What (18) does rule out is systematic sorting based on unmodelled match effects. Such forms of endogenous mobility are a priori unlikely in the institutional context of Russian public procurement (see Section 3). Nevertheless, we now explore the possibility, following the existing literature, especially [Card *et al.* \(2013\)](#).

First, bias can arise if organizations choose bureaucrats to work with based on match effects (as in e.g. [Mortensen & Pissarides, 1994](#)). Using (16), an organization that switches from working with bureaucrat 1 to bureaucrat 2 can expect the prices it pays to change by

$$\text{E}[p|b = 1] - \text{E}[p|b = 2] = \alpha_1 - \alpha_2 + \text{E}[\varepsilon_i|b = 1] - \text{E}[\varepsilon_i|b = 2] \quad (19)$$

$$\begin{aligned} &= \alpha_1 - \alpha_2 + \text{E}[\eta_{1j}] - \text{E}[\eta_{2j}] + \text{E}[\theta_{1g}] - \text{E}[\theta_{2g}] \\ &\quad + \text{E}[\zeta_j|b = 1] - \text{E}[\zeta_j|b = 2] + \text{E}[\nu_i|b = 1] - \text{E}[\nu_i|b = 2] \end{aligned} \quad (20)$$

If organization-bureaucrat match effects influence organizations' choice of bureaucrats, $\text{E}[\eta_{1j}] \neq \text{E}[\eta_{2j}]$. To test for this possibility, we construct an event study tracking organizations that switch bureaucrats. We define an employment spell as a sequence of at least two purchases an organization-bureaucrat pair conduct together with less than 400 days between purchases.³⁹ Wherever possible, we then match

³⁷We 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} .

³⁸As the dimensions of the matrices of fixed effects involved are large, rather than inverting a high-dimensional matrix, we solve the OLS normal equations directly using the `lfe` package for R written by [Gaure \(2015\)](#).

³⁹Appendix figure [OA.2](#) requires spells to contain at least three days, and the results are similar.

an employment spell (event time ≤ 0) with the earliest future spell (event time > 0) involving the same organization but a different bureaucrat. This change of bureaucrats then constitutes an event (event time = 0). We classify the two bureaucrats involved in the event using the average price they achieve in purchases they make for *other* organizations during the quarter that the spell ends (for the earlier spell) or starts (for the later spell). We assign this bureaucrat-average price to the relevant quartile of the distribution of all bureaucrats' average prices in the same quarter.

Panel A of Figure 2 presents the results. The horizontal axis displays event time, i.e. purchase dates. On the vertical axis we display the average prices paid by the pair at a given point in event time, residualizing out month and good fixed effects. We see that prices paid change sharply when an organization switches to a less or more effective bureaucrat. In particular, the price changes associated with switching bureaucrat quartiles appear symmetric: organizations switching from a bureaucrat in the first quartile of prices to a bureaucrat in the fourth quartile experience a price change that is of opposite sign but very similar magnitude to organizations switching from the fourth to the first quartile. These last two observations together are compelling evidence against the possibility of strong or moderate sorting on match effects η_{bj} .⁴⁰

Another possibility is that organizations that become better (or worse) at procurement over time systematically switch to a different type of bureaucrat, or vice versa for bureaucrats. In that case $E[\zeta_j|b = 1] \neq E[\zeta_j|b = 2]$ in the example in (19). However, we do not see any systematic time trend in the trajectories of switchers in Panel A of Figure 2, suggesting that drift in effectiveness and organizations switching bureaucrats are uncorrelated.

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 organizations to replace their bureaucrats. This would lead us to overstate the difference in the bureaucrat effects since $E[\nu_i|b = 1] > E[\nu_i|b = 2]$. However, Panel A of Figure 2 shows no systematic "Ashenfelter dip" just before a bureaucrat switch, suggesting that the transitory error ν_i is not correlated with organizations switching bureaucrats.

It is also possible that procurers specialize in goods for which they are better at achieving low prices. In the example in (19), 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 $E[\theta_{1g}] < E[\theta_{2g}]$. 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 Online Appendix Figure OA.1. Each panel shows the same general patterns as in Figure 2. In addition to alleviating any concerns due to unmodeled match effects between organizations or bureaucrats and goods, this helps rule out the possibility of strong or moderate correlation between drift in the procurer effects or the transitory error and procurers switching goods.

Panel B of Figure 2 is identical to Panel A, except that we now depict the change in the number of suppliers that participate in an organization's auctions when the organization switches bureaucrats; the intermediate outcome that we hypothesize is the primary channel through which procurer effectiveness

⁴⁰If there was sorting on match effects, we would expect all switchers to display price drops and those moving from the first to the fourth quartile to display a smaller price increase than organizations moving in the opposite direction.

ultimately affects the price achieved. As in Panel A, the patterns in Panel B show no indication that the identifying assumptions in (18) are violated. Taken together, the evidence presented in this sub-section strongly suggests that the bureaucrat and organizations effects estimated using (12) can be interpreted as individual and organizational sources of the prices the government pays for its inputs.

5.4 The bureaucracy's influence on output

We now present this paper's first main result: our estimate of the extent to which individuals and organizations in the state's bureaucratic tier affect public sector output. Table 2 implements the variance decomposition in equation (12) in the analysis sample. The first column shows estimates of the variances from using the raw fixed effects from (11). Relative to the total variation in prices paid (s.d. = 1.646, controlling for good and month of the purchase), the standard deviations of the bureaucrat and organization effects in the first column are large, at 1.031 and 1.068 respectively. However, the two are negatively correlated so that the joint effect of bureaucrats and organizations has a standard deviation of 0.876. This estimate becomes 1.036 if we add in the connected set fixed effects to capture the total effect of procurers both within and across connected sets, as seen in rows 9 and 10.⁴¹ In Table 3 we repeat the variance decomposition using only the largest connected set. The results are very similar to those in Table 2. The standard deviation of log prices is 1.773, of which 1.063 can be attributed to bureaucrats and organizations. Overall, our estimates thus imply that bureaucrats and organizations jointly explain 60 percent of the standard deviation of log prices paid. Bureaucrats and organizations each explain about half of this total effect.

Column 3 of tables 2 and 3 presents results from an analogous variance decomposition, except that we now focus on variance in the number of suppliers that participate in procurement auctions. The results are similar to those for prices: the standard deviations of the procurer effects are large—somewhat larger for bureaucrats than organizations; and bureaucrat and organizations jointly explain about half of the total variation in supplier participation. This suggests that an important channel through which individual procurers matters for prices paid is that effective bureaucrats and procurers lower entry costs and attract more suppliers to their auctions, consistent with the conceptual framework in Section 2.

The model we have estimated assumes that the price achieved is approximately log-linear in the bureaucrat and organization effects. In the Appendix we probe this assumption in two ways: by examining patterns in the size of residuals across the bivariate distribution of the estimated bureaucrat and organization effects, and by reestimating equation (11) with fixed effects for each bureaucrat-organization *pair* added. We find no systematic patterns in the residuals and that the improvement in the model's fit from adding pair effects is very small, indicating that our log-linear model is a good approximation to the true, underlying production function.

⁴¹The variance of the bureaucrat and organization effects can be computed either across specific procurement purchases ("across items") or across pairs of bureaucrats and organizations ("across pairs"); we show both but focus mostly on the former since the across item-variance is arguably a more precise measure (because it weighs bureaucrat-organization pairs more the more times a pair appears together in the data). As discussed in Sub-section 5.2, it is well-known that the estimated covariance term in AKM models is downward biased (Andrews *et al.*, 2008). We therefore do not emphasize the estimated covariance between bureaucrats and organizations. Note, however, that the total variance explained by bureaucrats and organizations *combined* should not suffer from limited mobility bias.

As seen from the standard errors in columns 2 and 4 of tables 2 and 3, the estimated variance and covariance terms are highly statistically significant.⁴² One way to illustrate their magnitude is to consider what they imply would happen if bureaucrats and/or organizations were moved from one percentile of the effectiveness distribution to another, for example because of changes in recruitment practices, training of existing bureaucrats, or improved management systems. As seen in Figure 3, our estimates imply that moving all bureaucrats and organizations above the 80th percentile in (in)effectiveness to 50th percentile-effectiveness would save the Russian government 37.3 percent of its annual procurement expenses. If the proportionate savings we estimate for such a decrease in bureaucratic ineffectiveness in the analysis sample were achieved on *all* Russian public procurement, the government would save about USD 69 billion a year (see Online Appendix Table OA.3). Similarly, Figure 3 shows that moving only bureaucrats above 80th percentile-effectiveness to 50th percentile-effectiveness would save the government 9.7 percent of procurement expenditures. In the Appendix we compare our results to existing estimates of the extent to which individuals and organizations affect output in other settings.

5.5 Like-for-like comparison

A possible concern is that the differences in unit prices we attribute to procurer effectiveness may reflect not only differences in price per quality-adjusted unit, but also differences in the quality of the goods purchased. To investigate this possibility, we first show that our findings are remarkably similar in a sub-sample of goods that is by nature very homogeneous—medicines.⁴³

As described in Sub-section 4.2, we extract each drug’s active ingredient, dosage, and packaging from LVEMD. We use these characteristics to assign medicines to barcode-level bins. These bins, rather than the text analysis method, are used to define goods categories and thus to determine which procurers to compare, in the pharmaceuticals sample. With these goods categories in hand, we make the same connectivity restrictions as in the full sample (see Sub-section 5.2). As seen in columns 4 and 5 of Table 1, the pharmaceuticals analysis sample is similar to the full pharmaceuticals sample.

When we reestimate (11) on the pharmaceuticals sample, we find that, as in the overall sample, about half of the variance in prices that is not explained by which good is being bought or when is attributable to the bureaucrats and organizations making the purchases. These results are shown in Table 4.

Another way to investigate if procurers we label “effective” purchase lower quality goods than procurers we label “ineffective” is to restrict the sample to the items our text analysis classification method is able to assign a 10-digit product code to.⁴⁴ As seen in Column 6 of Table 5, the results from our variance decomposition exercise are essentially unaffected by this sample restriction.

Finally, we show that (i) our results are robust to restricting attention to the most homogeneous types of goods in the analysis sample, and (ii) that the results do not change as we allow the sample to include

⁴²We discuss how these standard errors are estimated in Sub-section 5.6.

⁴³This approach is inspired by influential earlier papers that focus on sectors producing especially homogeneous goods (Syverson, 2004; Hortacsu & Syverson, 2007; Foster *et al.*, 2009; Bronnenberg *et al.*, 2015).

⁴⁴The algorithm developed in Step 2 of the procedure outlined in Sub-section 4.1 assigns a 10-digit code to 63 percent of the items in our analysis sample with high confidence. The remaining items in the analysis sample are also clustered into homogeneous bins, but we cannot confidently assign a 10-digit code to these items.

more and more heterogeneous types of goods. We split the sample into quintiles of good homogeneity as defined by the commonly-used measure of scope for quality differentiation developed by [Sutton \(1998\)](#). We then reestimate (11) on successive subsamples. The first five columns of Table 5 shows the results. Column 5 includes all observations we are able to match to the [Sutton \(1998\)](#) ladder.⁴⁵ As we move from right to left, we restrict the sample to more and more homogeneous goods. As expected, the variance of average prices paid decreases with good type homogeneity. However, the estimated share of the variance explained by bureaucrats and organizations remains largely unchanged across the columns. In Online Appendix Table OA.4 we repeat this exercise using an alternative measure of scope for quality differentiation developed by [Khandelwal \(2010\)](#) and find the same result.

Overall, the tests discussed in this sub-section indicate that our measures of micro level state effectiveness are not confounded by differences in quality or other characteristics of the goods procurers that on average pay low versus high prices purchase.⁴⁶

5.6 Finite sample issues

A separate set of estimation issues arise from sampling error in finite samples. 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 ([Neyman & 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. 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](#)).⁴⁷

We address the finite sample issues in three ways. First, when calculating standard errors for our variance decomposition, we bootstrap so that we can take into account the patterns of correlation in the residuals. We construct partial residuals $\epsilon_i = p_i - \mathbf{X}_i\hat{\beta}$ 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 limitations⁴⁸, but makes bootstrapping feasible with our large dataset.

Our second method for dealing with sampling error is a non-parametric approach, similar to [Finkel-](#)

⁴⁵We are able to match 73 percent of the items assigned an 10-digit HS code in Step 2 of text analysis method with the [Sutton \(1998\)](#) ladder. We thank Eric Verhoogen for sharing the Sutton ladder data with us.

⁴⁶Another possibility that organizations endogenously respond to the effectiveness of the bureaucrats available to them by purchasing more/fewer, or different types of, goods. This would lead us to underestimate the true variance in procurer effectiveness and its consequences.

⁴⁷We are estimating models with three sets of high-dimensional fixed effects (for bureaucrats, organizations, and goods). (The models also contain month dummies to control for common time trends, but there are few enough of these month effects that “month-connectedness” is not an issue). To our knowledge, identification results for models with more than two sets of fixed effects are not yet available ([Gaure, 2013](#)), and providing such results is beyond the scope of this paper.

⁴⁸The procedure imposes clustering at the bureaucrat-organization level in the standard errors. Moreover, since we use the partial residuals ϵ_i rather than reestimating the full model on each iteration, we do not account for correlation between bureaucrat and organization assignment and \mathbf{X} . The procedure drastically speeds up computation.

stein *et al.* (2016) & Silver (2016). We randomly split our sample in half, stratifying by bureaucrat-organization pair. We then estimate equation (11) 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. The errors in the two estimates should be uncorrelated, so we can create split-sample estimates of the relevant variance terms as follows:

$$\begin{aligned}\widehat{\text{Var}}^{SS}(\alpha_b) &= \text{Cov}(\hat{\alpha}_b^1, \hat{\alpha}_b^2) & \widehat{\text{Var}}^{SS}(\psi_j) &= \text{Cov}(\hat{\psi}_j^1, \hat{\psi}_j^2) \\ \widehat{\text{Var}}^{SS}(\gamma_s) &= \text{Cov}(\hat{\gamma}_s^1, \hat{\gamma}_s^2) & \widehat{\text{Var}}^{SS}(\alpha_b + \psi_j) &= \text{Cov}(\hat{\alpha}_b^1 + \hat{\psi}_j^1, \hat{\alpha}_b^2 + \hat{\psi}_j^2)\end{aligned}$$

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.* (2014a). 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 $\text{Var}(\hat{\alpha}) = \sigma_\alpha^2 + \sigma_\mu^2$ and the variance of our estimated organization effects is $\text{Var}(\hat{\psi}) = \sigma_\psi^2 + \sigma_\omega^2$.

Our bootstrap method to calculate standard errors 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 $\hat{\sigma}_\alpha^2 = \text{Var}(\hat{\alpha}) - \text{E}_b[s_b^2]$, where expectations are taken across bureaucrats and with weights $1/s_b^2$. The procedure for constructing the variance of the organization effects is analogous. 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 $\lambda_b = \arg \min_{\lambda} \text{E}[\alpha_b - \lambda \hat{\alpha}_b] = \sigma_\alpha^2 / (\sigma_\alpha^2 + \sigma_{\mu_b}^2)$ and analogously for λ_j , and our shrinkage estimators replace these terms with their sample analogues $\hat{\alpha}_b^{Sh} = \lambda_b \hat{\alpha}_b$ and $\hat{\psi}_j^{Sh} = \lambda_j \hat{\psi}_j$.

The results from our variance decomposition for prices and participation using procurer effect estimates that are corrected for sampling error biases using the split-sample and shrinkage methods are shown in Table 6, alongside the estimates from Table 2. The split-sample estimates in columns 2 and 5 are very similar to the raw fixed effects estimates in columns 1 and 4. Using the shrunk procurer effects yields standard deviations of the bureaucrat and organization effects are about 30–45 percent smaller, relative to the total variation in prices and participation. However, the standard deviation of the joint effect of bureaucrats and organizations remains very similar to the raw fixed effect estimates, whether we focus on the split-sample or the shrunk estimates. We conclude that our first main finding—bureaucrats and organizations jointly explain half of the variation in procurement effectiveness in Russia, of which about half is due to bureaucrats and half to organizations—is unchanged when we correct our estimates for finite sample biases. Our estimates of the separate effect of bureaucrats and organizations is moderately smaller when we do so.

5.7 Correlates of bureaucratic effectiveness

What do effective policy implementers do differently? In this sub-section we relate variation in individual bureaucrats and organizations estimated fixed effects, α_b and ψ_j , to observed variation in behavior

and intermediate auction outcomes. As discussed in Section 4, our data contains detailed information on the evolution of each procurement process, from the initial 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 bivariate regressions of the bureaucrat and organization effects on all procurement process-related variables in our data, and then use a LASSO procedure to select which subset of the variables to include in a multivariate regression (Tibshirani, 1996).⁴⁹

Figures 4 and 5 show the results. The left panel of each figure shows regression coefficients from a series of bivariate regressions of the bureaucrat (in Figure 4) and organization (in Figure 5) effects on each of the selected observables alone. The right panel shows the coefficients from the multivariate regression of the procurer effects on all of the LASSO-selected variables. To facilitate comparison, all variables are standardized to have unit standard deviation. The coefficients can thus be interpreted as the effect in standard deviations of the bureaucrat/organization effects of a one standard deviation change in the measure of procurer behavior. Of course, the relationships displayed in Figures 4 and 5 need not be causal, in part because we do not observe everything different procurers do differently.

Four key findings emerge from Figure 4. First, effective buyers encourage many and diverse applicants. Organizations who attract more applicants pay lower prices, and bureaucrats who pay lower prices are more likely to award contracts to firms from other regions of Russia. Second, successful buyers make the auctions accessible and predictable by setting low required deposits, and writing contracts that don't require further modification after being signed. Third, more experienced procurers and in-house bureaucrats pay lower prices.⁵⁰ Fourth, geography and level of government impact prices strongly. Organizations that are further from their regional capital, and organizations at regional and municipal (as opposed to federal) level pay higher prices.

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.

6 Individual and Organizational Sources of Heterogeneous Policy Impacts: the Case of Bid Preferences for Favored Firms

In Section 5 we held the policy environment constant. We varied the bureaucrat and organization in charge of procurement, exploiting the thousands of quasi-experiments created by the movement of organizations across bureaucrats, and vice versa, to estimate how individual procurers affect public sector output. In this section we instead hold constant the procurers in charge of a purchase. We vary whether

⁴⁹More 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.

⁵⁰For bureaucrats, early and late delivery of goods is correlated with estimated price effectiveness in the direction one would expect, but these can only explain (or "offset") a tiny fraction of the dispersion in price effectiveness we estimate. Importantly, bureaucrats' success rate at making purchases (the fraction of auction announcements that actually result in a purchase) is uncorrelated with prices.

a particular policy that is commonly used worldwide applies and study whether the impact of the policy depends on the procurers in charge of the purchase. We do this for two reasons. First, if, as we argued above, participation costs are a key way in which bureaucrats and organizations affect procurement outcomes, then our conceptual framework in section 2 makes testable predictions for how the impacts of policy changes should depend on bureaucratic effectiveness. Furthermore, the second overall goal of this paper is to determine if there are policy design implications of micro level sources of state effectiveness.

6.1 Average impact of bid preferences

We first estimate the *average* treatment effect of Russia’s “buy local” policy described in Section 3. The policy comes into effect each year in the late spring, covering a subset of goods that varies year-on-year (albeit moderately so). These forms of variation allow us to estimate the policy’s ATE through a difference-in-differences strategy. Because there must be a minimum of one bidder in the auction offering a Russian-made good and a minimum of one bidder offering a foreign-made good for the policy to apply our estimates should be interpreted as Intent to Treat (ITT) effects.⁵¹

We estimate the Intent to Treat (ITT) effect of the preferences policy in regressions of this form:

$$y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Preferred}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt} \quad (21)$$

where y_{igt} is an outcome for the purchase of item i , which is a good of type g , in month t . Preferred_{gt} is a dummy indicating that g is on the preferences list in the year month t falls within, and PolicyActive_t is a dummy indicating that the year’s list of preferred goods has been published and the policy activated. \mathbf{X}_{igt} are the controls we use in our estimation in the previous section, but for clarity we separate the good and month fixed effects, μ_g and λ_t that capture time-invariant differences across goods and aggregate price trends, respectively. ε_{igt} is an error term we allow to be clustered by month and good.

Table 7 shows the results of estimating (21) in the analysis sample. Recall that procurers’ mandate is simply to acquire the items they purchase at the lowest possible price while following the government’s policy rules. In Column 1 we see that the average effect of the preferences policy on the log price achieved, controlling for month and good fixed effects and quantity, is a precisely estimated zero.⁵² This is despite the modest decrease in firms participating in procurement auctions when the policy applies seen in Column 3.⁵³ In columns 2 and 4 we see that the results are similar if equation (21) is estimated on the largest connected set sample.

Our estimates are valid estimates of the policy’s ITT under the parallel trends assumption that the time trend of prices paid for preferred goods would have mirrored that of unpreferred goods had

⁵¹In the analysis sample we do not observe goods’ country of origin.

⁵²Because y_{igt} is the price *paid*, not the winning bid, in this regression, the estimated ATE captures the “automatic” savings achieved by the government (in auctions won by a foreign manufacturer) from paying the winning supplier less than its bid. More generally, both the entry and bidding (conditional on entry) behavior of favored and non-favored firms are expected to respond endogenously to preference programs of this form. As is well-known, how prices are affected is therefore theoretically ambiguous in general (see e.g. McAfee & McMillan, 1989; Marion, 2007; Krasnokutskaya & Seim, 2011; Athey *et al.*, 2013; Bhattacharya *et al.*, 2014).

⁵³Note that the estimate in Column 3 is only significant at the 10 percent level.

the policy not been implemented. In our setting, the policy switches on and off multiple times, so a violation of the parallel trends assumption is a priori unlikely. There are two main reasons why the assumption might be violated. First, secular trends in prices may be different for the two groups of goods. Second, seasonality in prices might be different across the two groups. Figure 6 presents time trends of average prices in the two groups, allowing us to directly evaluate the assumption visually. We see no evidence of differences in either secular trends or seasonality.

In the analysis sample we do not observe goods' origin and so we cannot assess if the preferences policy achieves the government's goal of channeling demand to domestic manufacturers. In the pharmaceuticals sample, however, we do observe where pharmaceuticals are manufactured. Since all pharmaceuticals are preferenced, we cannot use variation across products in the application of the policy in our analysis. To repeat the difference-in-differences analysis in the pharmaceuticals subsample, we exploit the fact an auction must feature at least one domestic and at least one foreign manufacturer for the policy to apply and redefine Preferred_g as equal 1 if the drug purchased is made both in Russia and abroad.

Table 8 shows the results of estimating equation (21) in the pharmaceuticals subsample. The estimated effect of the policy on prices and participation are similar to those in the analysis sample, but as Column 3 shows, we find an increase in the likelihood that the winner is a domestic manufacturer. It thus appears that the preferences policy achieves the government's goal of purchasing more Russian-made goods. This finding is noteworthy since shifting demand towards domestic manufacturers comes at no direct cost to the government, as we saw in Table 7. This result contrasts with those from studies of similar preference policies in the U.S. (see e.g. Marion, 2007; Krasnokutskaya & Seim, 2011). In this sense, our estimates of the average impact of Russia's "buy local" policy point toward the possibility that industrial policies of this form are more successful in countries like Russia where bureaucrats and organizations are on average likely *less* effective than in advanced countries, foreshadowing our findings in the next sub-section.⁵⁴

6.2 Bureaucratic effectiveness and heterogeneity in the impact of bid preferences

The framework in Section 2 predicts that the introduction of bid preferences for favored firms should affect the prices achieved by low and high entry cost procurers differently. In particular, the framework predicts that the beneficial effect on prices of higher entry by favored firms should dominate the effect of lower entry by non-favored firms for high entry cost procurers—who have low baseline entry rates—but that the latter effect should dominate for low entry cost procurers. In Section 5 we saw that effective and ineffective procurers in Russia pay markedly different prices for the same goods in a standard policy regime that treats all firms the same. We also saw that bureaucrats and organizations that achieve lower prices are the ones who impose lower entry costs on firms. We thus hypothesize that bid preferences lead to a decrease in prices when administered by procurers of low estimated effectiveness, but an increase in prices when administered by procurers of high estimated effectiveness.

⁵⁴Pinning down the longer-term welfare consequences of channeling demand to potentially less productive firms is beyond the scope of this paper.

To test this hypothesis, we interact $\text{Preferred}_{gt} \times \text{PolicyActive}_t$ in equation (21) with the estimated procurer effects from Section 5 as follows:

$$\begin{aligned}
 y_{igt} = & \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta\text{Preferred}_{gt} \times \text{PolicyActive}_t + \gamma\text{Preferred}_{gt} \times \hat{\alpha}_b \\
 & + \zeta\text{Preferred}_{gt} \times \hat{\psi}_j + \eta\text{PolicyActive}_t \times \hat{\alpha}_b + \theta\text{PolicyActive}_t \times \hat{\psi}_j \\
 & + \pi\text{Preferred}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \zeta\text{Preferred}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt}
 \end{aligned} \tag{22}$$

Tables 9 (using the full analysis sample and the largest connected set) and 10 (using the pharmaceuticals subsample) show the results.⁵⁵ Several important findings emerge from these two tables. First, we see that the zero average price effect found in Table 7 combines a price *increase* when the preferences policy is administered by effective bureaucrats and organizations and a price *decrease* among ineffective procurers. The introduction of the “buy local” policy thus results in convergence of the price performance of effective and ineffective bureaucrats, and effective and ineffective organizations. The estimates in columns 3 and 4 of Table 9 and column 2 of Table 10 suggest that the explanation has to do with the entry costs associated with procurers at the high and low end of the performance range. Ineffective bureaucrats and organizations see an increase in firms participating in procurement auctions when the preferences policy applies, while participation decreases in auctions administered by effective procurers.

In Figure 7 we investigate how the impact of bid preferences for favored firms on prices and participation varies with bureaucratic effectiveness non-parametrically. The four panels depict the treatment effect of the preferences policy on prices and participation for bureaucrats and organizations of each decile of effectiveness, relative to the treatment effect for decile 1 (the most effective procurers). On the horizontal axis we plot the average effectiveness within the relevant decile, and on the vertical axis the corresponding (relative) treatment effect estimate. The estimated treatment effects decrease in magnitude throughout the observed range. The decline appears relatively linear, and somewhat flatter in the middle part of the distribution for bureaucrat effectiveness than organizational effectiveness. An important take-away from Figure 7 is that heterogeneity in how the preferences policy affects procurement outcomes is not concentrated among especially effective or ineffective procurers, but seen throughout the distribution of effectiveness.

A possible concern with our heterogeneous policy effects analysis is that our estimates may be picking up differences in seasonality or mean reversion across different types of bureaucrats and organizations. This would require time patterns across different bureaucrats and organizations to match the timing of the policy in very unusual ways, but is nevertheless a possibility. Appendix Table OA.6 performs two series of placebo tests to assuage these concerns. In Panel A we consider moving the timing of the policy in each year forward by increasing numbers of months. We see that as the timing shifts earlier, the magnitude of the triple-difference coefficients decreases, as we would expect, since we are moving untreated observations in the spring into the treatment group and treated observations in the winter into the control group. In Panel B we consider a placebo in which we imagine that the preference policy did not apply in increasing parts of early 2015 (even though it did). The estimated placebo effects

⁵⁵In tables 9 and 10 we use the shrunken estimates of the bureaucrat and organization effects. Online Appendix Table OA.5 uses the raw fixed effect estimates and shows very similar results.

are insignificant in all but one of the twelve estimates.

In the pharmaceuticals sample we can test if the differential change in the prices achieved by effective and ineffective procurers under the preferences policy is accompanied by a corresponding differential change in the probability that Russian manufacturers win procurement contracts. As seen in column 3 of Table 10, this is not the case. This suggests that, from the perspective of a government trying to minimize the prices it pays for its goods while simultaneously steering government demand towards domestic manufacturers, a “buy local” procurement policy of the form used in Russia is a more desirable policy tool when the procurers administering the policy are less effective.

One way to illustrate how much more desirable the preferences policy is when bureaucratic effectiveness is low is to continue our counterfactual example from Section 5. Bureaucrats and organizations above the 80th percentile in (in)effectiveness paid 17 percent less under the preferences policy for the same goods. Suppose procurers in this group were moved to 50th percentile-effectiveness. Our estimates imply that the preferences policy would then have led them to spend 3 percent *more* than they do in the absence of the policy.

Our findings in this sub-section are consistent with the argument that individuals’ and organizations’ influence on auction entry costs are a key driver of bureaucratic effectiveness. As our model in Section 2 shows, heterogeneity in auction entry costs across procurers predicts that heterogeneity in the effects of the preferences policy will follow precisely the patterns that we see: namely, that the policy will lead to improvements in prices and participation when implemented by procurers with low participation and high prices at baseline, while the opposite will be the case with ineffective procurers.

Overall our findings in this section provide the first direct evidence of the magnitude of the potential benefits of designing government policy with the effectiveness of those who implement policy—individuals and organizations in the bureaucratic tier of the state—in mind. Our estimates suggest that if policymakers want to steer demand towards domestic manufacturers using bid preferences, the size of the optimal bid preference to apply is higher for a less effective bureaucratic apparatus.

7 Conclusion

In this paper we have presented evidence that, contrary to the mechanistic view of the bureaucracy taken by the existing literature, the individuals and organizations tasked with implementing policy are important sources of state effectiveness. Bureaucrats and public sector organizations together account for 60 percent of the variation in quality-adjusted prices paid by the Russian government for its inputs. Consistent with our simple endogenous entry model of procurement, 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 favors bids from domestic manufacturers, we find that the induced increase in entry by domestic suppliers outweighs the tilting of the playing field against foreign manufacturers 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. First, they suggest that there are huge returns to the state from employing more bureaucrats at the high end of the observed performance range, training bureaucrats better, or improving organization-wide characteristics such as management quality. For example, our estimates imply that if the worst 20 percent of bureaucrats and organizations had 50th percentile effectiveness, government savings would be 37.3 percent. The large magnitude of the procurer effects we estimate suggests that the political leadership and front-line provider questions studied in the existing literature on the state enterprise may have been over-emphasized relative to the middle, bureaucratic tier of the state.

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 differences in policy implementer effectiveness across the two settings is essential. We show how bureaucratic effectiveness can be estimated in baseline data and then used in the estimation of heterogeneous treatment effects to guide such extrapolation.

Finally, our findings imply that policies that are suboptimal when state effectiveness is high may be second-best optimal when state effectiveness is low. For example, our heterogeneous treatment effect estimates imply that the Russian bid preference policy saved the least effective 20 percent of bureaucrats and organizations 17 percent of expenditures, but that if this group of ineffective procurers had been of median effectiveness then the policy would have led them to spend 3 percent more. Such dependence of policies' impact on state effectiveness may be part of the reason why many policies work well in some countries or regions and poorly in others. An important take-away is that policies should be designed with the effectiveness of the individuals and organizations that will implement the policies in mind. Achieving the *best* policy outcomes likely requires both maximizing the effectiveness of the bureaucratic apparatus and choosing policies that are tailored to the effectiveness of their implementers. However, our results also suggest that the returns to tailoring policy to implementers' effectiveness are likely to be especially large when bureaucratic effectiveness is low.

References

- ABOWD, JOHN M., KRAMARZ, FRANCIS, & MARGOLIS, DAVID. 1999. High Wage Workers and High Wage Firms. *Econometrica*, **67**, 251–333.
- ABOWD, JOHN M., CREECY, ROBERT H., & KRAMARZ, FRANCIS. 2002. *Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data*. Census Bureau Technical Paper TP-2002-06.
- ACEMOGLU, DARON, HASSAN, TAREK A., & ROBINSON, JAMES A. 2011. Social structure and development: a legacy of the holocaust in Russia. *The Quarterly Journal of Economics*.
- ACEMOGLU, DARON, GARCIA-JIMENO, CAMILO, & ROBINSON, JAMES A. 2015. The Role of State Capacity in Economic Development. *The Political Economist*.
- 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., GILL, LEONARD, SCHANK, THORSTEN, & UPWARD, RICHARD. 2008. High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias. *Journal of the Royal Statistical Society*, **171**(3), 673–697.
- 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. 2014a. *Do-gooders and Go-getters: Career Incentives, Selection, and Performance in Public Service Delivery*. mimeo, Harvard.
- ASHRAF, NAVA, BANDIERA, ORIANA, & JACK, KELSEY. 2014b. No Margin, No Mission? A Field Experiment on Incentives for Public Services Delivery. *Journal of Public Economics*.
- ATHEY, SUSAN, COEY, DOMINIC, & LEVIN, JONATHAN. 2013. Set-Asides and Subsidies in Auctions. *American Economic Journal: Microeconomics*.
- BAI, YING, & JIA, RUIXUE. 2016. Elite Recruitment and Political Stability: The Impact of the Abolition of China’s Civil Service Exam. *Econometrica*.
- BAJARI, PATRICK, HOUGHTON, STEPHANIE, & TADELIS, STEVE. 2014. Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs. *American Economic Review*.
- BANDIERA, ORIANA, PRAT, ANDREA, & VALLETTI, TOMMASO. 2009. Active and Passive Waste in Government Spending: Evidence from a Policy Experiment. *American Economic Review*, **99**(4), 1278–1308.
- BARABASHEV, ALEXEI, & STRAUSSMAN, JEFFREY D. 2007. Public service reform in Russia, 1991–2006. *Public Administration Review*, **67**(3), 373–382.
- BERTRAND, MARIANNE, & SCHOAR, ANTOINETTE. 2003. Managing with style: The effect of managers on firm policies. *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. 2010. State capacity, conflict and development. *Econometrica*, **78**, 1–34.
- BESLEY, TIMOTHY, & PERSSON, TORSTEN. 2009. The origins of state capacity: Property rights, taxation, and politics. *American Economic Review*, **99**, 1218–44.
- BEST, MICHAEL CARLOS, BROCKMEYER, ANNE, KLEVEN, HENRIK, SPINNEWIJN, JOHANNES, & WASEEM, MAZHAR. 2015. Production vs Revenue Efficiency With Limited Tax Capacity: Theory and

- Evidence From Pakistan. *Journal of Political Economy* 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.
- 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*.
- BLOOM, NICHOLAS, PROPPER, CAROL, SEILER, STEPHAN, & REENEN, JOHN VAN. 2015c. The Impact of Competition on Management Practices in Public Hospitals. *The Review of Economic Studies*.
- BLOOM, NICHOLAS, BENDER, STEFAN, CARD, DAVID, REENEN, JOHN VAN, & WOLTER, STEPHANIE. 2016. *Management practices, workforce selection and productivity*. Unpublished paper.
- BLOOM, NICHOLAS, SONG, JAE, PRICE, DAVID, GUVENEN, FATIH, & VON WACHTER, TILL. 2017. *Firming up inequality*. Unpublished paper.
- 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.
- BRANZOLI, NICOLA, & DECAROLIS, FRANCESCO. 2015. Entry and Subcontracting in Public Procurement Auctions. *Management Science*.
- 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, ROBIN, HANSEN, MATTHEW, OLKEN, BENJAMIN, POTAPOV, PETER, & SIEBER, STEFANIE. 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, & REZAEI, ARMAN. 2016a. *Personalities and Public Employee Performance: Experimental Evidence from Pakistan*. mimeo UCSD.
- CALLEN, MICHAEL, GULZAR, SAAD, HASANAIN, ALI, & KHAN, YASIR. 2016b. *The Political Economy of Public Employee Absence: Experimental Evidence from Pakistan*. mimeo UCSD.
- CARD, DAVID, HEINING, JÖRG, & KLINE, PATRICK. 2013. Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics*, **128**, 967–1015.
- CARD, DAVID, CARDOSO, ANA RUTE, & KLINE, PATRICK. 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.
- CHEREMUKHIN, ANTON, GOLOSOV, MIKHAIL, GURIEV, SERGEI, & TSYVINSKI, ALEH. 2016. The Industrialization and Economic Development of Russia through the Lens of a Neoclassical Growth Model.

The Review of Economic Studies.


- CHETTY, RAJ, FRIEDMAN, JOHN N., & ROCKOFF, JONAH E. 2014a. Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates. *American Economic Review*, **104**, 2593–2632.
- CHETTY, RAJ, FRIEDMAN, JOHN N., & ROCKOFF, JONAH E. 2014b. Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood. *American Economic Review*, **104**, 2533–2679.
- CHONG, ALBERTO, LA PORTA, RAFAEL, DE SILANES, FLORENCIO LOPEZ, & SHLEIFER, ANDREI. 2014. Letter Grading Government Efficiency. *Journal of the European Economic Association*.
- COVIELLO, DECIO, & GAGLIARDUCCI, STEFANO. forthcoming. Tenure in Office and Public Procurement. *American Economic Journal: Economic Policy*.
- COVIELLO, DECIO, MORETTI, LUIGI, SPAGNOLO, GIANCARLO, & VALBONESI, PAOLA. 2017. Court Efficiency and Procurement Performance. *Scandinavian Journal of Economics*.
- COVIELLO, DECIO, GUGLIELMO, ANDREA, & SPAGNOLO, GIANCARLO. forthcoming. The Effect of Discretion on Procurement Performance. *Management Science*.
- DAL BO, ERNESTO, FINAN, FRED, & ROSSI, MARTIN. 2013. Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *Quarterly Journal of Economics*.
- DEATON, ANGUS. 2010. Instruments, Randomization and Learning about Development. *Journal of Economic Literature*, 424–455.
- DECAROLIS, FRANCESCO. 2014. Awarding Price, Contract Performance and Bids Screening: Evidence from Procurement Auctions. *American Economic Journal: Economic Policy*.
- DECAROLIS, FRANCESCO, & PALUMBO, GIULIANA. 2015. Renegotiation of Public Contracts: An Empirical Analysis. *Economic Letters*.
- DEHEJIA, RAJEEV, POP-ELECHES, CRISTIAN, & SAMII, CYRUS. 2016. *From Local to Global: External Validity in a Fertility Natural Experiment*. NBER Working Paper No. 21459.
- DELL, MELISSA, LANE, NATHAN, & QUERUBIN, PABLO. 2017. *The Historical State, Local Collective Action, and Economic Development in Vietnam*. Unpublished paper.
- DESERRANNO, ERIKA. 2016. *Financial Incentives as Signals: Experimental Evidence from the Recruitment of Village Promoters in Uganda*. mimeo Northwestern University.
- DUFLO, ESTHER, GREENSTONE, MICHAEL, PANDE, ROHINI, & RYAN, NICHOLAS. 2013. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics*, **128**, 1499–1545.
- DUFLO, ESTHER, GREENSTONE, MICHAEL, PANDE, ROHINI, & RYAN, NICHOLAS. 2016. *The Value of Regulatory Discretion: Estimates from Environmental Inspections in India*. NBER Working Paper.
- ECKHOUT, JAN, & KIRCHER, PHILIPP. 2011. Identifying Sorting-In Theory. *The Review of Economic Studies*.
- ENIKOLOPOV, RUBEN, & ZHURAVSKAYA, EKATERINA. 2007. Decentralization and political institutions. *Journal of Public Economics*.
- EPPLE, DENNIS. 1987. Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products. *Journal of Political Economy*, **95**, 59–80.
- FERRAZ, CLAUDIO, FINAN, FEDERICO, & SZERMAN, DIMITRI. 2015. *Procuring Firm Growth: The Effects of Government Purchases on Firm Dynamics*. mimeo, UC Berkeley.
- FINAN, FEDERICO, OLKEN, BENJAMIN A., & PANDE, ROHINI. 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*.

- FISMAN, RAYMOND, & GATTI, ROBERTA. 2002. Decentralization and Corruption: Evidence from U.S. Federal Transfer Programs. *Public Choice*.
- 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.
- GENTRY, MATTHEW, & LI, TONG. 2014. Identification in auctions with selective entry. *Econometrica*, **82**(1), 315–344.
- GENTZKOW, MATTHEW, & SHAPIRO, JESSE. 2010. What Drives Media Slant? Evidence From US Daily Newspapers. *Econometrica*, **78**, 35–71.
- GOLDBERG, PINELOPI KOUJIANOU, & DE LOECKER, JAN. 2014. Firm Performance in a Global Market. *Annual Review of Economics*.
- GOLDSCHMIDT, DEBORAH, & SCHMIEDER, JOHANNES. 2015. *The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure*. NBER Working Paper No. 21366.
- GRILICHES, ZVI (ed). 1971. *Price Indexes and Quality Change*. Cambridge: Harvard University Press.
- GULZAR, SAAD, & PASQUALE, BENJAMIN. 2017. Politicians, Bureaucrats, and Development: Evidence from India. *American Political Science Review*, **111**, 162–183.
- HANNA, REMA, & WANG, SHING-YI. 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.
- HANSMAN, CHRISTOPHER, HJORT, JONAS, & LEÓN, GIANMARCO. 2016. *Interlinked Firms and the Consequences of Piecemeal Regulation*. Unpub. paper.
- HART, OLIVER, SHLEIFER, ANDREI, & VISHNY, ROBERT W. 1997. The Proper Scope of Government: Theory and an Application to Prisons. *The Quarterly Journal of Economics*.
- 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.
- JONES, BENJAMIN F., & OLKEN, BENJAMIN A. 2005. Do leaders matter? National leadership and growth since World War II. *The Quarterly Journal of Economics*.
- KANE, THOMAS J., & STAIGER, DOUGLAS O. 2008. *Estimating teacher impacts on student achievement: An experimental evaluation*. NBER working Paper.
- KHAN, ADNAN, KHWAJA, ASIM, & OLKEN, BEN. 2016. Tax Farming Redux: Experimental Evidence on

- Performance Pay for Tax Collectors. *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, ELENA, & SEIM, KATJA. 2011. Bid Preference Programs and Participation in Highway Procurement Auctions. *American Economic Review*, 2653–2686.
- KRISHNA, VIJAY. 2010. *Auction Theory (second edition)*. Elsevier.
- KRYLOVA, DINA V., & SETTLES, ALEXANDER. 2011. *Procurement Reform in Russia: Implications for the Fight Against Corruption*. Economic Reform Feature Service.
- LA PORTA, RAFAEL, DE SILANES, FLORENCIO LOPEZ, SHLEIFER, ANDREI, & VISHNY, ROBERT W. 1999. The quality of government. *The Journal of Law, Economics, and Organization*.
- 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.
- MARION, JUSTIN. 2007. Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions. *Journal of Public Economics*, 1591–1624.
- MAS, ALEXANDRE, & MORETTI, ENRICO. 2009. Peers at Work. *American Economic Review*.
- MCAFEE, PRESTON R., & MCMILLAN, JOHN. 1989. Government procurement and international trade. *Journal of International Economics*, 26(3), 291–308.
- MENDES, RUTE, VAN DEN BERG, GERARD J., & LINDEBOOM, MAARTEN. 2010. An Empirical Assessment of Assortative Matching in the Labor Market. *Labour Economics*, 17, 919–29.
- MENG, XIN, QIAN, NANCY, & YARED, PIERRE. 2015. The Institutional Causes of Famine in China, 1959–61. *The Review of Economic Studies*.
- 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.
- MIRONOV, MAXIM, & ZHURAVSKAYA, EKATERINA. 2016. Corruption in Procurement: Evidence from Financial Transactions Data. *American Economic Journal: Economic Policy*.
- 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.
- NEYMAN, JERZY, & SCOTT, ELIZABETH L. 1948. Consistent Estimates Based on Partially Consistent Observations. *Econometrica*, 16, 1–32.
- OECD. 2016. *Preventing Corruption in Public Procurement*. 2016 version.
- PANG, SINNO JIALIN, & YANG, QIANG. 2010. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22, 1345–1359.
- RASUL, IMRAN, & ROGGER, DANIEL. forthcoming. Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service. *Economic Journal*.

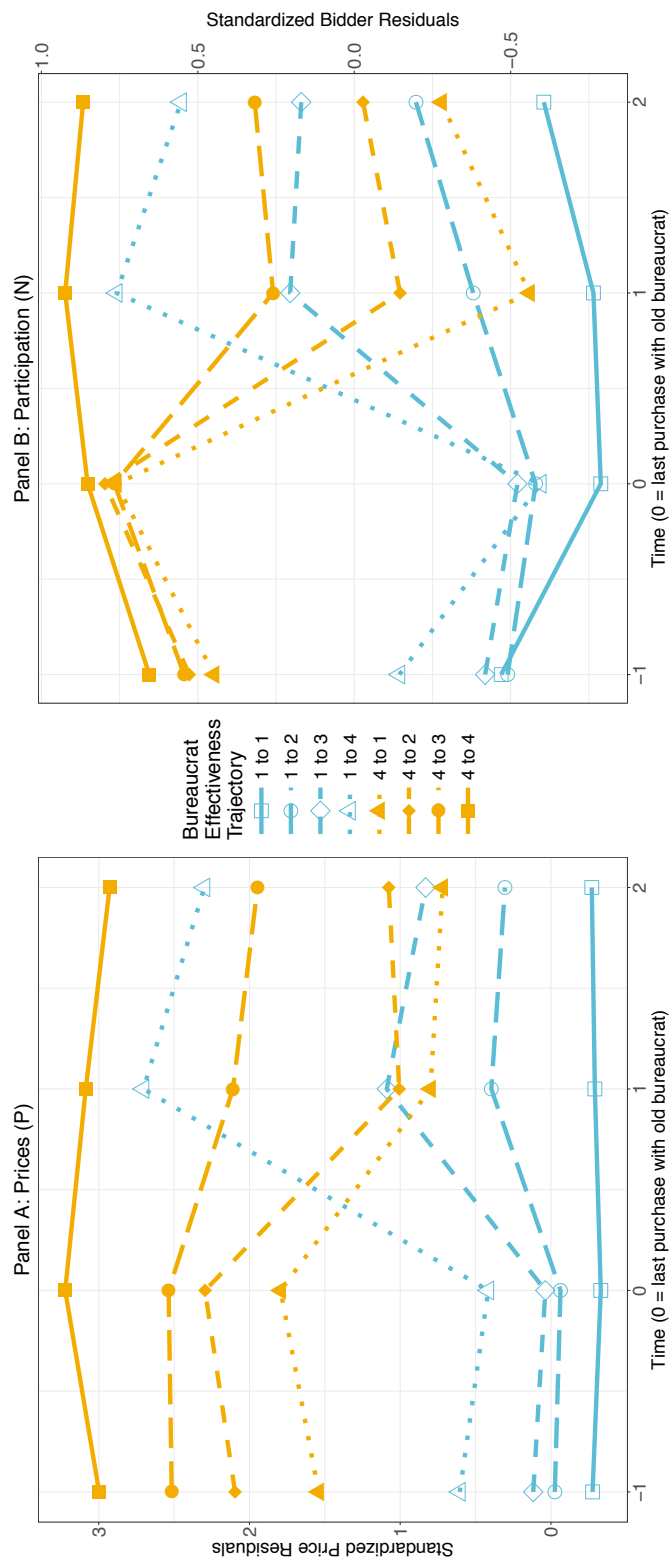
- RAUCH, JAMES E. 1999. Networks Versus Markets in International Trade. *Journal of International Economics*, 7–35.
- RIFKIN, RYAN, & KLAUTAU, ALDEBARO. 2004. In Defense of One-Vs-All Classification. *Journal of Machine Learning Research*, 5, 101–141.
- ROBERTS, JAMES W., & SWEETING, ANDREW. 2015. *Bailouts and the Preservation of Competition*. National Bureau of Economic Research.
- RODRIG, DANI. 2009. Diagnostics before Prescription. *Journal of Economic Perspectives*.
- ROSEN, SHERWIN. 1974. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82, 34–55.
- SAMUELSON, WILLIAM F. 1985. Competitive bidding with entry costs. *Economics Letters*, 17(1), 53–57.
- SCHAPPER, PAUL R., MALTA, JOAO NUNO VEIGA, & GILBERT, DIANE L. 2009. Analytical Framework for the Management and Reform of Public Procurement. In: THAI, KHI V. (ed), *International Handbook of Public Procurement*. CRC Press.
- SCHARGRODSKY, ERNESTO, & DI TELLA, RAFAEL. 2003. The Role of Wages and Auditing during a Crackdown on Corruption in the City of Buenos Aires. *Journal of Law and Economics*.
- SCHOTT, PETER K. 2004. Across-Product Versus Within-Product Specialization in International Trade. *Quarterly Journal of Economics*, 119, 647–678.
- SCULLEY, D. 2010. Web-Scale K-Means Clustering. *Proceedings of the 19th International Conference on World Wide Web*.
- SHELEF, ORIE, & NUGYEN-CHYUNG, AMY. 2015. *Competing for Labor through Contracts: Selection, Matching, Firm Organization and Investments*. mimeo, Stanford.
- SILVER, DAVID. 2016. *Haste or Waste? Peer Pressure and the Distribution of Marginal Returns to Health Care*. Mimeo: UC Berkeley.
- SUTTON, JOHN. 1998. *Technology and Market Structure: Theory and History*. Cambridge: MIT Press.
- SYVERSON, CHAD. 2004. Market Structure and Productivity: A Concrete Example. *Journal of Political Economy*.
- SZAKONYI, DAVID. 2016. *Businesspeople in Elected Office: Identifying Private Benefits from Firm-Level Returns*. SSRN Working Paper.
- TIBSHIRANI, ROBERT. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*.
- TORREY, LISA, & SHAVLIK, JUDE. 2009. Transfer Learning. In: *Handbook of Research on Machine Learning Applications*. IGI Global.
- TRANSPARENCY INTERNATIONAL. 2016. *Corruption Perceptions Index*. 2016 version.
- VIVALT, EVA. 2016. *How Much Can We Impact Evaluations Inform Policy Decisions?* Mimeo: Stanford.
- XU, GUO. 2017. *The Costs of Patronage: Evidence from the British Empire*. mimeo LSE.
- 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.
- YAKOVLEV, EVGENY, & ZHURAVSKAYA, EKATERINA. 2014. The Unequal Enforcement of Liberalization: Evidence from Russia's Reform of Business Regulation. *Journal of the European Economic Association*.
- YAO, YANG, & ZHANG, MUYANG J. 2015. Subnational leaders and economic growth: evidence from Chinese cities. *J Econ Growth*.

FIGURE 1: PROCUREMENT PROCESS FLOW-CHART



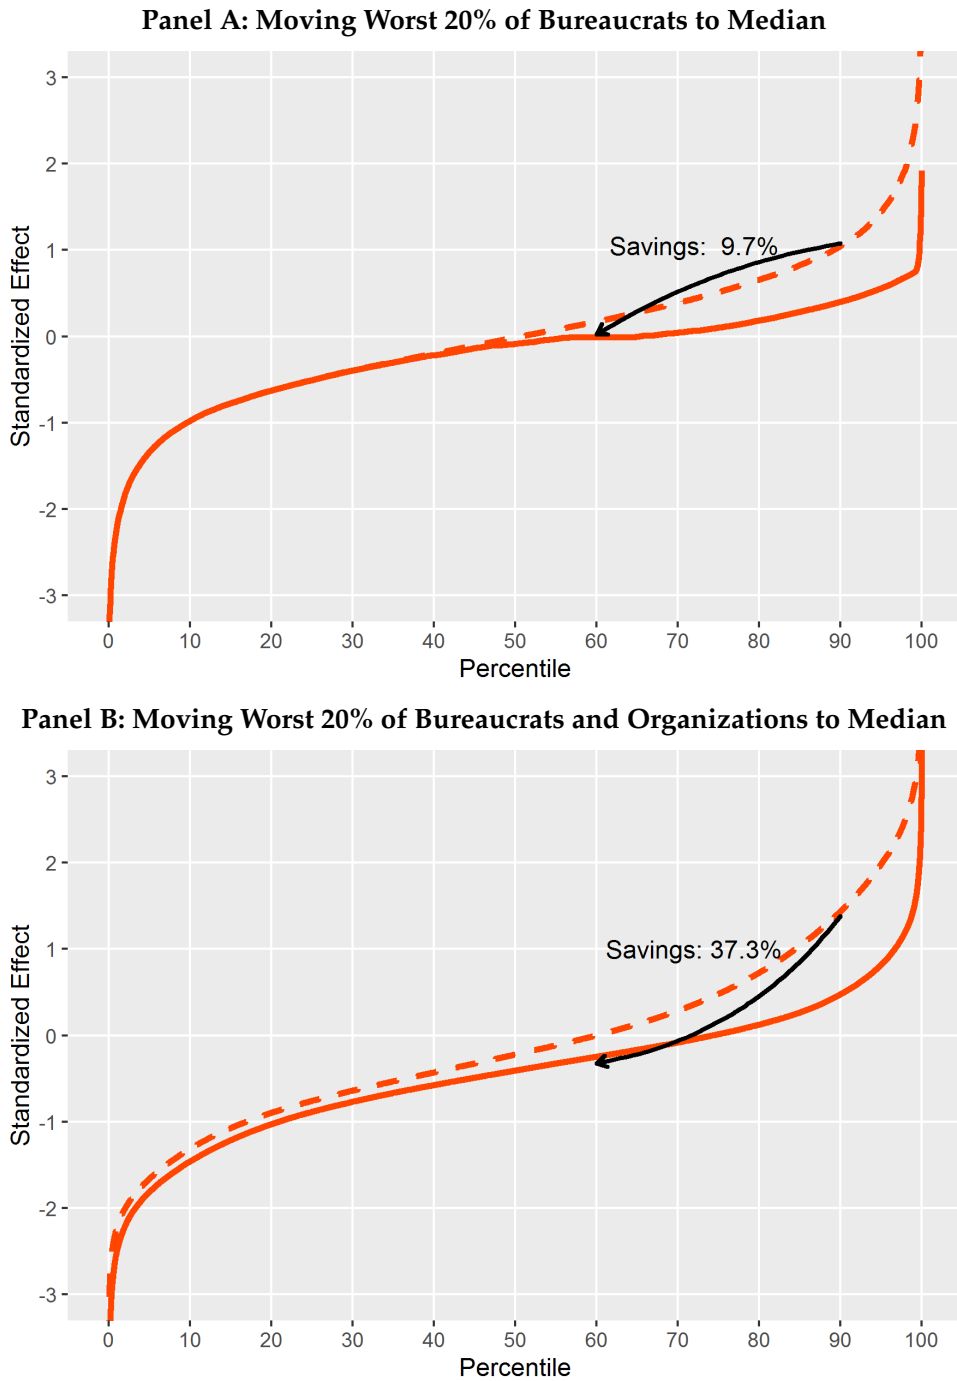
This figure lays out the stages of the process public procurement purchases of off-the-shelf goods through electronic auctions follow in Russia. Numbers are based on all purchases made under laws 94 and 44 in 2011-2015. The stages are described in detail in Sub-section 3.2.

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




The figure shows the evolution of prices and participation around the times when organizations switch which bureaucrat makes purchases with and for them. Each point on the horizontal axis represents a date when a given bureaucrat-organization pair makes a purchase together, with event time = 0 being the last occasion on which the organization works with the “old” bureaucrat, and event time = 1 being the first occasion on which the organization works with the “new” bureaucrat. The vertical axis on Panel A measures average residualized prices (log) paid by the bureaucrat-organization pair, while Panel B measures the average number of residualized bidders participating in the auction by the bureaucrat-organization pair. Prices and bidders separately are residualized by regressing each outcome on good and month fixed effects. We create a balanced panel in which we require each bureaucrat-organization pair to work together on two separate dates and each bureaucrat to work with at least one other organization in the quarter containing event time = 0 (for the “old” bureaucrat the organization works with before the switch) or event time = 1 (for the “new” bureaucrat the organization works with after the switch). Bureaucrats are classified into quartiles according to the average (residualized) prices (Panel A) or average (residualized) bidders (Panel B) they achieve with the other organizations they work with in the quarter containing event time = 0 (for the “old” bureaucrat) or the quarter containing event time = 1 (for the “new” bureaucrat).

FIGURE 3: CRUDE COUNTERFACTUALS




The figure shows the impact of two counterfactual scenarios on the distribution of our estimated price effects. Panel A considers moving all bureaucrats above the 80th percentile of their connected set's distribution of shrunken price effects down to their connected set's median. The dashed line shows the distribution of our shrunken estimates of the bureaucrat effects, while the solid line shows the distribution that would result from implementing the counterfactual. Panel B considers moving both all bureaucrats and all organizations above the 80th percentile of their connected set's distribution down to the median. The dashed line shows the distribution of bureaucrat-organization pair effects we estimate, while the solid line shows the distribution that would occur in the counterfactual scenario. Overlaid on both panels are the implied aggregate savings.

FIGURE 4: CORRELATES OF BUREAUCRAT EFFECTIVENESS




The figure shows the results of regressions of estimated bureaucrat effects $\hat{\alpha}_b$ from estimation of equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(i,b,j)} + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The right column shows the coefficients from a multivariate regression of the estimated bureaucrat effects on all the correlates that are selected by a LASSO regularization procedure with the regularization parameter that gives the minimum cross-validated error.

FIGURE 5: CORRELATES OF ORGANIZATION EFFECTIVENESS



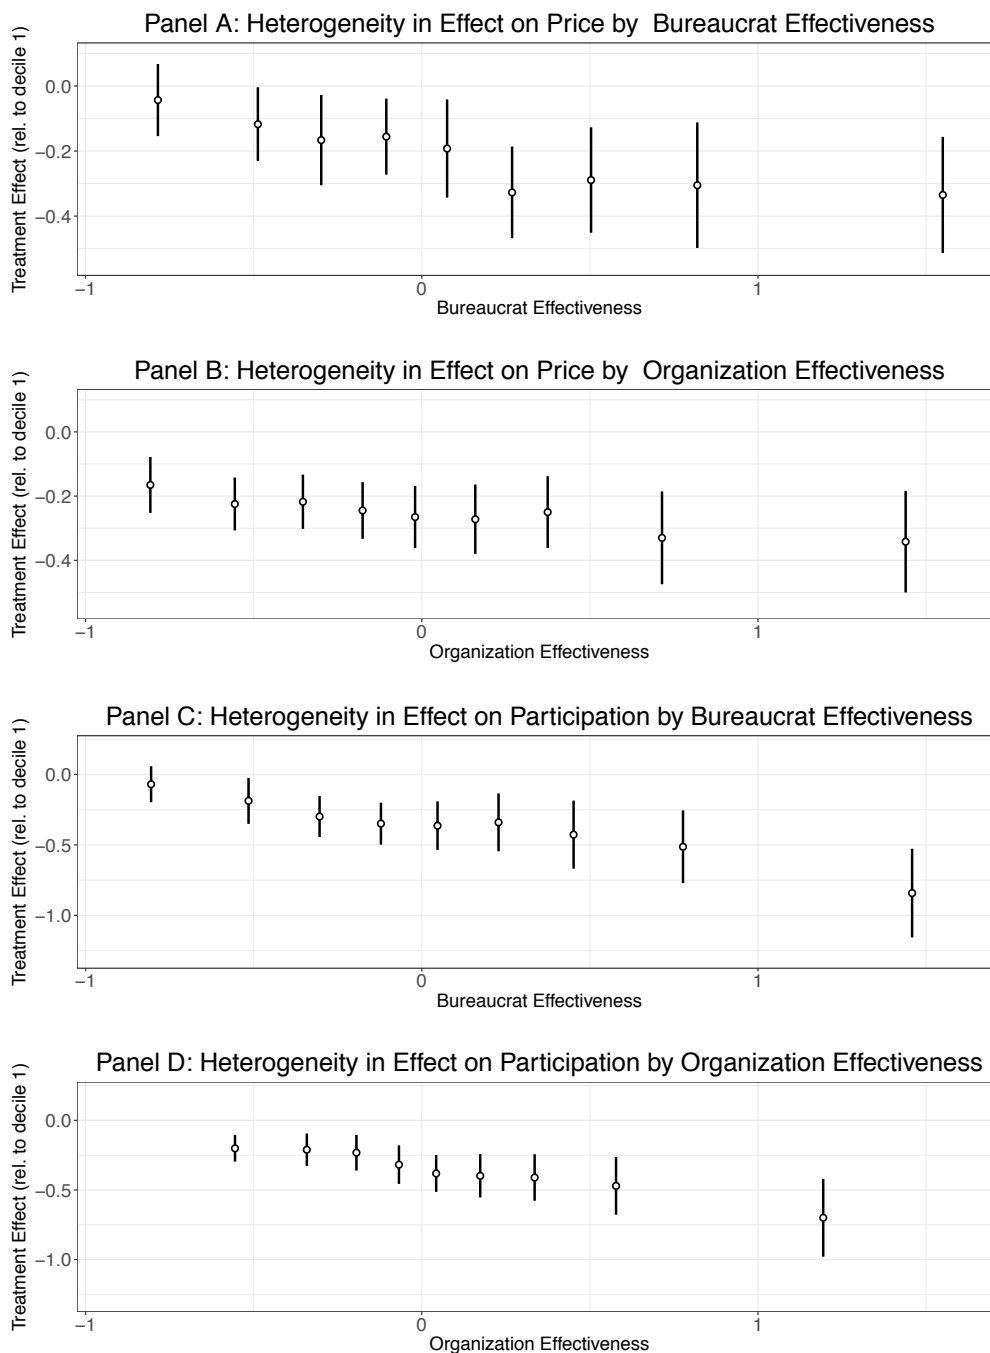
The figure shows the results of regressions of estimated organization effects $\hat{\psi}_j$ from estimation of equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_s(b,j) + \varepsilon_i$ on observable characteristics of the purchase procedure followed. The left column shows standardized bivariate regressions for each correlate individually. The right column shows the coefficients from a multivariate regression of the estimated organization effects on all the correlates that are selected by a LASSO regularization procedure with the regularization parameter that gives the minimum cross-validated error.

FIGURE 6: EFFECT OF BID PREFERENCES FOR DOMESTIC MANUFACTURERS: ASSESSMENT OF PARALLEL TRENDS



The figure shows a graphical analysis of the preferences policy over the period of study. The x-axis is measured in months, with dotted vertical lines indicating when the preference policy became active in 2011, 2012, 2013, and 2014 (the policy remained active for all of 2015). The solid vertical lines indicate when the policy was no longer active in each year. The top panel shows residualized prices averaged over month for each treatment group. Prices are residualized by regressing the log price on good fixed effects and the interaction between 2-digit HS Product categories, years, region, and lot size. We trim the top and bottom 1% from the residuals within each treatment group and month interaction. The bottom panel shows the difference between the two treatment groups on a scale for the y-axis equal to one-half of the standard deviation of the trimmed residualized prices.

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



The figure shows results from a non-parametric estimation of the triple-differences equation (22): $y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta \text{Preferred}_{gt} \times \text{PolicyActive}_t + \gamma \text{Preferred}_{gt} \times \hat{\alpha}_b + \zeta \text{Preferred}_{gt} \times \hat{\psi}_j + \eta \text{PolicyActive}_t \times \hat{\alpha}_b + \theta \text{PolicyActive}_t \times \hat{\psi}_j + \pi \text{Preferred}_{gt} \times \text{PolicyActive}_t \times \hat{\alpha}_b + \zeta \text{Preferred}_{gt} \times \text{PolicyActive}_t \times \hat{\psi}_j + \varepsilon_{igt}$. Bureaucrat and Organization effects are instead binned into deciles and the decile dummies interacted with the treatment indicators PolicyActive_t and Preferred_{gt} and their interaction. On the horizontal axis we plot the average effectiveness within the relevant decile. On the vertical axis we show the corresponding treatment effect estimate (relative to decile 1, which itself drops out). The top two panels present estimates from using log price as the outcome, while the bottom two panels present estimates from using the number of bidders as the outcome.

TABLE 1: SUMMARY STATISTICS

	All Products			Pharmaceuticals Subsample	
	Full Sample	Analysis Sample	Largest Connected Set	Full Sample	Analysis Sample
(1) # of Bureaucrats	123,099	35,774	12,980	6,015	3,070
(2) # of Organizations	95,420	43,385	10,854	3,653	2,131
(3) # of Connected Sets	28,147	605	1	1,476	122
(4) # of Bureaucrats with >1 Org.	14,742	11,008	3,490	1,168	1,146
(5) # of Organizations with >1 Bur.	57,382	36,859	9,758	2,334	1,713
(6) # of Federal Organizations	13,461	1,547	442	494	38
(7) # of Regional Organizations	25,980	15,126	4,749	2,661	1,742
(8) # of Municipal Organizations	55,979	26,712	5,663	498	351
(9) # of Health Organizations	14,378	9,355	2,105	3,266	2,050
(10) # of Education Organizations	50,616	26,008	6,778	126	69
(11) # of Internal Affairs Organizations	15,659	3,196	741	224	0
(12) # of Agr/Environ Organizations	1,682	447	51	2	0
(13) # of Other Organizations	13,085	4,379	1,179	35	12
(14) # of Goods	18,650	17,946	15,501	4,476	4,153
(15) # of Regions	90	90	62	75	69
(16) # of Auction Requests	1,733,422	1,149,496	368,072	85,240	63,540
(17) Mean # of Applicants	3.42	3.43	3.5	2.78	2.81
(18) Mean # of Bidders	2	2.01	2.07	1.87	1.89
(19) Mean Reservation Price (bil. USD)	34,512	38,059	59,488	56,398	56,760
(20) Quantity Mean	968	981	1,909	412	426
Median	20	20	20	40	40
SD	113,415	131,901	260,248	6,442	6,977
(21) Total Price Mean (bil. USD)	368	310	670	306	256
Median	8	7.1	11	13	13.9
SD	1,940	1,748	2,949	8,790	1,698
(22) Unit Price Mean (bil. USD)	324	275	597	726	86.7
Median	0.421	0.325	0.62	0.361	0.381
SD	52,394	61,274	27,275	114,142	1,413
(23) # of Observations	15,366,194	11,228,122	2,858,982	274,991	200,816
(24) Total Procurement Volume (bil. USD)	1,074	797	358	21.6	14.8

The table reports summary statistics for five samples. The All Products columns show statistics for purchases of all off-the-shelf goods, while the Pharmaceuticals Subsample columns restrict attention to purchases of medicines. Full Sample denotes all unpreferred auctions. Analysis Sample denotes all unpreferred auctions in connected sets that fulfill three restrictions: singleton bureaucrat-organization, bureaucrat-good, and organization-good pairs are removed; each procurer (bureaucrats and organizations) implements a minimum of five purchases; and connected sets have at least three bureaucrats and organizations. Largest Connected Set is the largest connected set from the Analysis Sample (as measured by the number of organizations). Organizations working in Education include schools, universities, pre-schools, and youth organizations. Organizations working in Internal Affairs include police, emergency services, local administration, taxes, and transportation. The Other category includes funds, monitoring agencies, and land cadasters, among many others. All sums are measured in billions of US dollars at an exchange rate of 30 rubles to 1 US dollar.

TABLE 2: SHARE OF VARIANCE OF PROCUREMENT PRICES AND PARTICIPATION EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS: FULL ANALYSIS SAMPLE

	Prices (P) (1)	(s.e.) (2)	Participation (N) (3)	(s.e.) (4)
(1) s.d. of Bureaucrat Effects	1.570	(0.0381)	1.257	(0.0244)
(2) s.d. of Organization Effects	1.372	(0.039)	0.979	(0.0257)
(3) s.d. of Connected Set Effects	1.000	(0.0115)	0.523	(0.0108)
(4) s.d. of Bur + Org Effects Within CS (across pairs)	1.258	(0.00519)	0.895	(0.00315)
(5) s.d. of Total Bur + Org Effects (across pairs)	1.364	(0.00247)	0.913	(0.00295)
(6) s.d. of Bureaucrat Effects (across items)	1.031	(0.0462)	0.919	(0.0418)
(7) s.d. of Organization Effects (across items)	1.068	(0.0496)	0.888	(0.0468)
(8) s.d. of Connected Set Effects (across items)	0.555	(0.035)	0.302	(0.0147)
(9) s.d. of Bur + Org Effects Within CS (across items)	0.876	(0.0154)	0.642	(0.00654)
(10) s.d. of Total Bur + Org Effects (across items)	1.036	(0.00126)	0.710	(0.00358)
(11) s.d. of Y	2.417		1.355	
(12) s.d. of Y good, month	1.646		1.241	
(13) Adjusted R-squared	0.955		0.837	
(14) Sample Size	11,228,122		11,228,122	

The table implements the variance decomposition in equation (12) using the estimates from equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_s(b,j) + \varepsilon_i$. The sample used is the Analysis Sample (All Products) summarized in Table 1. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Columns (2) and (4) shows standard errors for the estimates in columns (1) and (3) respectively, estimated by bootstrapping 100 times. In Column (3), the outcome variable, Participation (N), is the number of bidders. The s.d. of the bureaucrat and organization effects can be computed either across specific procurement purchases (“across items”) or across pairs of bureaucrats and organizations (“across pairs”); the across item s.d. weights bureaucrats-organizations pairs more the more times the pair appears together in the data.

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

	Prices (P) (1)	(s.e.) (2)	Participation (N) (3)	(s.e.) (4)
(1) s.d. of Bureaucrat Effects	1.803	(0.187)	1.282	(0.0564)
(2) s.d. of Organization Effects	1.499	(0.291)	0.970	(0.0689)
(3) s.d. of Bur + Org Effects (across pairs)	1.577	(0.00414)	1.155	(0.00281)
(4) s.d. of Bureaucrat Effects (across items)	1.287	(0.244)	0.883	(0.0813)
(5) s.d. of Organization Effects (across items)	1.241	(0.271)	0.792	(0.0823)
(6) s.d. of Bur + Org Effects (across items)	1.063	(0.00287)	0.656	(0.00216)
(7) s.d. of Y	2.683		1.364	
(8) s.d. of Y good, month	1.773		1.231	
(9) Adjusted R-squared	0.959		0.828	
(10) Sample Size	2,858,982		2,858,982	

The table implements the variance decomposition in equation (12) using the estimates from equation (11): $p_i = X_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. The sample used is the Largest Connected Set (All Products) summarized in Table 1. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i, j)$. Columns (2) and (4) shows standard errors for the estimates in columns (1) and (3) respectively, estimated by bootstrapping 100 times. In Column (3), the outcome variable, Participation (N), is the number of bidders. The s.d. of the bureaucrat and organization effects can be computed either across specific procurement purchases ("across items") or across pairs of bureaucrats and organizations ("across pairs"); the across item s.d. weights bureaucrats-organizations pairs more the more times the pair appears together in the data.

TABLE 4: SHARE OF VARIANCE OF PROCUREMENT PRICES AND PARTICIPATION EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS: PHARMACEUTICALS SUBSAMPLE WITH BARCODE INFORMATION

	Prices (P)	(s.e.)	Participation (N)	(s.e.)
	1	2	3	4
(1) s.d. of Bureaucrat Effects	0.330	(0.0168)	0.786	(0.00907)
(2) s.d. of Organization Effects	0.282	(0.0134)	0.572	(0.00457)
(3) s.d. of Connected Set Effects	0.387	(0.0245)	0.217	(0.0028)
(4) s.d. of Bur + Org Effects Within CS (across pairs)	0.225	(0.00723)	0.626	(0.0022)
(5) s.d. of Total Bur + Org Effects (across pairs)	0.279	(0.0181)	0.628	(0.00427)
(6) s.d. of Bureaucrat Effects (across items)	0.191	(0.0149)	0.520	(0.00915)
(7) s.d. of Organization Effects (across items)	0.194	(0.0149)	0.449	(0.00911)
(8) s.d. of Connected Set Effects (across items)	0.201	(0.0133)	0.0823	(0.00224)
(9) s.d. of Bur + Org Effects Within CS (across items)	0.162	(0.00672)	0.518	(0.00193)
(10) s.d. of Total Bur + Org Effects (across items)	0.251	(0.0152)	0.519	(0.00467)
(11) s.d. of Y	2.048		1.204	
(12) s.d. of Y good, month	0.404		1.066	
(13) Adjusted R-squared	0.997		0.849	
(14) Sample Size	200,816		200,816	

The table implements the variance decomposition in equation (12) using the estimates from equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. The sample used is the Analysis Sample (Pharmaceuticals) summarized in Table 1. The good fixed effects used when running equation (11) are here barcode-level and constructed using the active ingredient, dosage, and packaging as described in section 4.2 (instead of our text analysis method). Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i,j)$. Columns (2) and (4) shows standard errors for the estimates in columns (1) and (3) respectively, estimated by bootstrapping 100 times. In Column (3), the outcome variable, Participation (N), is the number of bidders. The s.d. of the bureaucrat and organization effects can be computed either across specific procurement purchases (“across items”) or across pairs of bureaucrats and organizations (“across pairs”); the across item s.d. weights bureaucrats-organizations pairs more the more times the pair appears together in the data.

TABLE 5: SHARE OF VARIANCE OF PROCUREMENT PRICES AND PARTICIPATION EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS: RELAXING HOMOGENEOUS GOODS ASSUMPTION

	(1)	(2)	(3)	(4)	(5)	(6)
	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5	10-Digit Codes
(1) s.d. of Bur + Org Effects Within CS (across items)	0.789	0.801	0.863	0.865	0.847	0.817
(2) s.d. of Total Bur + Org Effects (across items)	0.927	0.966	1.053	1.007	1.027	1.008
(3) s.d. of log P	1.752	2.175	2.291	2.390	2.433	2.388
(4) s.d. of log P good, month	1.271	1.444	1.532	1.574	1.599	1.526
(5) s.d. of Bur+Org Within Efs / s.d. of log P good, month	0.621	0.555	0.564	0.549	0.530	0.535
(6) s.d. of Bur+Org Total Efs / s.d. of log P good, month	0.729	0.669	0.687	0.640	0.643	0.660
(7) Sample Size	1,097,233	2,275,959	3,231,115	4,300,461	5,222,931	7,055,150

The table implements the variance decomposition in equation (12) using the estimates from equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i, j)$. Column (6) uses the sub-sample consisting of all auctions for goods that our text analysis classification method is able to assign a 10-digit product code to. Column (5) uses the sub-set of the sample in Column (6) that we can match to the scope-for-quality-differentiation ladder developed by Sutton (1998). Column (4) removes the quintile with the highest scope-for-quality-differentiation according to the Sutton (1998) ladder, Column (3) the highest two quintiles, and so on.

TABLE 6: SHARE OF VARIANCE OF PROCUREMENT PRICES AND PARTICIPATION EXPLAINED BY BUREAUCRATS AND ORGANIZATIONS: ADDRESSING FINITE SAMPLE ISSUES

	Prices (P)			Participation (N)		
	Raw Fixed Effects (1)	Split-Sample (2)	Min MSE (3)	Raw Fixed Effects (4)	Split-Sample (5)	Min MSE (6)
(1) s.d. of Bureaucrat Effects	1.570	1.466	0.864	1.257	1.128	0.910
(2) s.d. of Organization Effects	1.372	1.356	0.816	0.979	0.860	0.661
(3) s.d. of Connected Set Effects	1.000	0.695	0.905	0.523	0.253	0.445
(4) s.d. of Bur + Org Effects Within CS (across pairs)	1.258	1.238	1.069	0.895	0.887	0.793
(5) s.d. of Total Bur + Org Effects (across pairs)	1.364	1.333	1.153	0.913	0.902	0.798
(6) s.d. of Bureaucrat Effects (across items)	1.031	0.988	0.664	0.919	0.794	0.744
(7) s.d. of Organization Effects (across items)	1.068	1.068	0.699	0.888	0.759	0.648
(8) s.d. of Connected Set Effects (across items)	0.555	0.506	0.527	0.302	0.250	0.274
(9) s.d. of Bur + Org Effects Within CS (across items)	0.876	0.859	0.819	0.642	0.629	0.612
(10) s.d. of Total Bur + Org Effects (across items)	1.036	0.997	0.974	0.710	0.677	0.670
(11) s.d. of Y	2.417	2.417	2.417	1.355	1.355	1.355
(12) s.d. of Y good, month	1.646	1.646	1.646	1.241	1.241	1.241
(13) Sample Size	11,228,122	11,228,122	11,228,122	11,228,122	11,228,122	11,228,122

The table implements the variance decomposition in equation (12) using the estimates from equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. The sample used is the Analysis Sample (All Products) summarized in Table 1. Each observation is an item procured by an organization j and a bureaucrat indexed by $b(i, j)$. In the left three columns the outcome is prices achieved, while in the right three columns the outcome is the number of bidders. Columns (1) and (4) shows estimates from using the raw fixed effects. Columns (2) and (5) shows estimates from using the split-sample method. This method randomly splits the sample in half, stratifying by bureaucrat-organization pair, and estimates bureaucrat and organization effects in each subsample. The variance components are then estimated as the covariances between the estimates from each subsample. Columns (3) and (6) shows estimates from using fixed effects estimated using the shrinkage method to minimize the mean-squared-error of predictions. This method uses the bootstraps to estimate the sampling error in each bureaucrat effect s_b^2 and each organization effect s_j^2 , and the signal variances of the bureaucrat and organization effects (σ_α^2 and σ_ψ^2 respectively). The minimum-mean-squared error predictor for each bureaucrat effect is then $[\hat{\sigma}_\alpha^2 / (\hat{\sigma}_\alpha^2 + s_b^2)] \cdot \hat{\alpha}_b$, where $\hat{\alpha}_b$ is the bureaucrat's fixed effect from the decomposition in Column (1), and analogously for the organization effects.

TABLE 7: AVERAGE EFFECT OF BID PREFERENCES POLICY FOR DOMESTIC MANUFACTURERS ON PROCUREMENT PRICES AND PARTICIPATION: FULL ANALYSIS SAMPLE

	Prices (P)		Participation (N)	
	Analysis Sample (1)	Largest Connected Set (2)	Analysis Sample (3)	Largest Connected Set (4)
log Standardized Quantity	-0.510*** (0.027)	-0.552*** (0.019)	0.032*** (0.003)	0.036*** (0.004)
Preferred (Good on list)	-0.050* (0.028)	-0.043 (0.047)	-0.145*** (0.040)	-0.195*** (0.053)
Preferred (Good on list) * Policy Active	-0.012 (0.025)	-0.007 (0.045)	-0.050 (0.047)	-0.029 (0.062)
Outcome Mean	5.69	6.26	1.64	1.68
Month, Good FEs	Yes	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes	Yes
Observations	21,017,045	3,973,832	21,017,045	3,973,832
R ²	0.592	0.620	0.276	0.272

*** p<0.01, ** p<0.05, * p<0.1 This table estimates the Intent to Treat (ITT) from equation (21): $y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta \text{Preferred}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt}$. In columns (1) and (3) the sample used is the combination of the Analysis Sample summarized in Column (2) of Table 1 and treated auctions that those procurers carried out. In columns (2) and (4) the sample used is the combination of the Largest Connected Set summarized in Column (3) of Table 1 and “treated” auctions that the procurers therein carried out. The first two columns estimate the ITT on the log price paid (P); the second two columns estimate the ITT on the number of bidders participating in the auction (N). An item has Preferred (Good on list) = 1 if the type of good appears on the list of goods covered by the preferences policy for that year. Policy Active = 1 during the part of the relevant year that the preferences policy was in effect. The Outcome Mean is the mean of the dependent variable in the control group, i.e. for goods that were not covered by preferences purchased during the period when the preferences policy was not active. Month and good fixed effects are included in all columns, as are interactions between 2-digit HS Product categories, years, region, and lot size. (We use “product” to distinguish the categories used in these interactions from the much more disaggregate goods categories used for the good fixed effects). Standard errors are clustered on month and good.

TABLE 8: AVERAGE EFFECT OF BID PREFERENCES FOR DOMESTIC PRODUCERS ON PROCUREMENT PRICES, PARTICIPATION, AND DOMESTIC PRODUCERS WINNING: PHARMACEUTICALS SAMPLE, ITT ANALYSIS

	Prices (P) (1)	Participation (N) (2)	Domestic Winner (3)
log Standardized Quantity	-0.039*** (0.003)	0.008*** (0.003)	0.009*** (0.002)
Preferred (Produced Russia+Abroad) * Policy Active	-0.007 (0.011)	-0.028 (0.028)	0.036*** (0.010)
Outcome Mean	6.27	1.89	0.32
Month, Active Ingredient FEs	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes
Observations	557,455	557,455	557,455
R ²	0.943	0.326	0.581

*** p<0.01, ** p<0.05, * p<0.1 This table estimates the Intent to Treat (ITT) from equation (21): $y_{igt} = \mathbf{X}_{igt}\beta + \mu_g + \lambda_t + \delta \text{Preferred}_{gt} \times \text{PolicyActive}_t + \varepsilon_{igt}$. The sample used is the combination of the Analysis Sample summarized in Column (4) of Table 1 and treated auctions that those procurers carried out. The first column estimates the ITT on the log price paid (P); the second column estimates the ITT on the number of bidders participating in the auction (N); and the third column estimates the ITT on a binary indicator for whether the medicine was produced by a domestic (Russian) manufacturer or not. Unlike in Table 7, we here define an auction as (potentially) preferred (Preferred (Produced Russia+Abroad)=1) if the relevant drug is produced both in Russia and abroad (all medicines are on the preferences list). Policy Active = 1 during the part of the relevant year that the preferences policy was in effect. The Outcome Mean is the mean of the dependent variable in the control group, i.e. for medicines that were produced either completely in Russia or abroad purchased during the period when the preferences policy was not active. Month and active ingredient fixed effects are included in all columns. Active ingredient denotes a higher category above the barcode-level that does not use information on dosage, packaging and manufacturer. Interactions between 2-digit HS Product categories, years, region, and lot size also included in all models. (We use "product" to distinguish the categories used in these interactions from the much more disaggregate goods categories used for the active ingredient fixed effects). Standard errors are clustered on month and active ingredient.

TABLE 9: HETEROGENEITY OF EFFECT OF BID PREFERENCES FOR DOMESTIC PRODUCERS ON PROCUREMENT PRICES AND PARTICIPATION BY BUREAUCRAT AND ORGANIZATION EFFECTIVENESS: ANALYSIS SAMPLE

	Prices (P)		Participation (N)	
	Analysis Sample (1)	Largest Connected Set (2)	Analysis Sample (3)	Largest Connected Set (4)
log Standardized Quantity	-0.481*** (0.023)	-0.544*** (0.017)	0.026*** (0.003)	0.031*** (0.003)
Preferred (Good on list)	0.044** (0.022)	0.005 (0.030)	-0.040 (0.025)	-0.077** (0.035)
Policy Active	0.023 (0.026)	-0.014 (0.043)	0.230** (0.108)	0.261** (0.124)
Bureaucrat FE	1.153*** (0.037)	1.247*** (0.025)	0.769*** (0.094)	0.911*** (0.089)
Organization FE	1.185*** (0.039)	1.425*** (0.030)	0.815*** (0.096)	0.941*** (0.090)
Preferred (Good on list) * Policy Active	-0.114*** (0.023)	-0.123*** (0.034)	-0.114*** (0.030)	-0.164*** (0.043)
Bureaucrat FE * Preferred (Good on list)	-0.026 (0.018)	-0.061** (0.026)	0.043 (0.042)	-0.028 (0.045)
Bureaucrat FE * Policy Active	-0.019 (0.017)	-0.051*** (0.019)	0.011 (0.096)	-0.087 (0.103)
Organization FE * Preferred (Good on list)	-0.008 (0.023)	-0.064** (0.031)	0.070 (0.043)	0.018 (0.050)
Organization FE * Policy Active	-0.023 (0.019)	-0.080*** (0.025)	0.045 (0.099)	-0.091 (0.106)
Bureaucrat FE * Preferred (Good on list) * Policy Active	-0.183*** (0.028)	-0.119*** (0.033)	-0.279*** (0.055)	-0.286*** (0.064)
Organization FE * Preferred (Good on list) * Policy Active	-0.164*** (0.029)	-0.111*** (0.040)	-0.307*** (0.060)	-0.295*** (0.075)
Outcome Mean	5.69	6.26	1.64	1.68
Month, Good FEs	Yes	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes	Yes
Connected Set FEs	Yes	Yes	Yes	Yes
Observations	15,957,594	3,973,832	15,957,594	3,973,832
R ²	0.645	0.692	0.372	0.364

*** p<0.01, ** p<0.05, * p<0.1 This table implements a triple-difference approach, interacting the Intent to Treat (ITT) from equation (21) with the estimated bureaucrat and organization effects from Section 5. In columns (1) and (3) the sample used is the combination of the Analysis Sample summarized in Column (2) of Table 1 and treated auctions that those procurers carried out. In columns (2) and (4) the sample used is the combination of the Largest Connected Set summarized in Column (3) of Table 1 and “treated” auctions that the procurers therein carried out. The first two columns estimate the triple-difference on the log price paid for each item (P); the second two columns estimate the triple-difference on the number of bidders participating in the auction (N). An item has Preferred (Good on list) = 1 if the type of good appears on the list of goods covered by the preferences policy for that year. Policy Active = 1 during the part of the relevant year that the preferences policy was in effect. The Outcome Mean is the mean of the dependent variable in the control group, i.e. for goods that were not covered by preferences purchased during the period when the preferences policy was not active. Month and good fixed effects are included in all columns, as are interactions between 2-digit HS Product categories, years, region, and lot size. (We use “product” to distinguish the categories used in these interactions from the much more disaggregate goods categories used for the good fixed effects). Standard errors are clustered on month and good.

TABLE 10: HETEROGENEITY OF EFFECT OF BID PREFERENCES FOR DOMESTIC PRODUCERS ON PROCUREMENT PRICES AND PARTICIPATION BY BUREAUCRAT AND ORGANIZATION EFFECTIVENESS: PHARMACEUTICALS SAMPLE

	Prices (P) (1)	Participation (N) (2)	Domestic Winner (3)
log Standardized Quantity	-0.031*** (0.001)	-0.005** (0.002)	0.008*** (0.002)
Policy Active	0.018 (0.012)	0.098 (0.085)	-0.014 (0.013)
Bureaucrat FE	0.809*** (0.050)	0.874*** (0.042)	0.007 (0.010)
Organization FE	0.803*** (0.048)	0.855*** (0.042)	-0.006 (0.012)
Preferred (Produced Russia+Abroad) * Policy Active	-0.010 (0.010)	-0.051** (0.025)	0.037*** (0.010)
Bureaucrat FE * Preferred (Produced Russia+Abroad)	0.056 (0.047)	0.148*** (0.030)	-0.017 (0.020)
Bureaucrat FE * Policy Active	-0.017 (0.048)	-0.024 (0.046)	0.001 (0.009)
Organization FE * Preferred (Produced Russia+Abroad)	0.018 (0.042)	0.162*** (0.031)	-0.009 (0.018)
Organization FE * Policy Active	-0.003 (0.049)	-0.037 (0.041)	0.008 (0.010)
Bureaucrat FE * Preferred (Produced Russia+Abroad) * Policy Active	-0.430*** (0.068)	-0.352*** (0.035)	-0.014 (0.022)
Organization FE * Preferred (Produced Russia+Abroad) * Policy Active	-0.402*** (0.072)	-0.277*** (0.033)	-0.014 (0.020)
Outcome Mean	6.27	1.89	0.32
Month, Active Ingredient FEs	Yes	Yes	Yes
Year×Product×Size×Region FEs	Yes	Yes	Yes
Connected Set FEs	Yes	Yes	Yes
Observations	406,424	406,424	406,424
R ²	0.954	0.430	0.587

*** p<0.01, ** p<0.05, * p<0.1 This table implements a triple-difference approach, interacting the Intent to Treat (ITT) from the equation (21) with the estimated bureaucrat and organization effects from Section 5. The sample used is the combination of the Analysis Sample summarized in Column (4) of Table 1 and treated auctions that those procurers carried out. The first columns estimate the triple-difference on the log price paid for each item (P); the second column estimates the triple-difference on the number of bidders participating in the auction (N); and the third column estimates the triple-difference on a binary indicator for whether the medicine was produced by a domestic (Russian) manufacturer or not. Unlike in Table 9, we here define an auction as (potentially) preferred (Preferred (Produced Russia+Abroad)=1) if the relevant drug is produced both in Russia and abroad (all medicines are on the preferences list). Policy Active = 1 during the part of the relevant year that the preferences policy was in effect. The Outcome Mean is the mean of the dependent variable in the control group, i.e. for medicines that were produced either completely in Russia or abroad purchased during the period when the preferences policy was not active. Month and active ingredient fixed effects are included in all columns. Active ingredient denotes a higher category above the barcode-level that does not use information on dosage, packaging and manufacturer. Interactions between 2-digit HS Product categories, years, region, and lot size also included in all models. (We use “product” to distinguish the categories used in these interactions from the much more disaggregate goods categories used for the active ingredient fixed effects). Standard errors are clustered on month and active ingredient.

A Appendix

A.1 Probing the log-linearity assumption

The model we have estimated assumes that the price achieved is approximately log-linear in the bureaucrat and organization effects. A direct piece of evidence in support of the log-linearity assumption comes from studying the distribution of the residuals across bureaucrat and organization effect deciles. If the log-linear specification was substantially incorrect, we would expect to see systematic patterns in the residuals. For example, positive match effects would lead the residuals to be large when the bureaucrat and organization are either both in the top or both in the bottom deciles of effectiveness. Figure A.1 shows a heat map of residuals for the analysis sample. The map reveals no clear patterns in the residuals.⁵⁶

As a further test of our log-linear model of prices, we reestimate equation (11) but include fixed effects for each bureaucrat-organization pair, allowing for arbitrary patterns of complementarity between bureaucrats and organizations (see also [Card *et al.*, 2013](#)). If there are indeed strong or moderate match effects that our model omits, then we expect this pair effect model to fit significantly better. The pair effect model does not fit the data much better than our baseline model: adding pair effects decreases the RMSE of the residuals from 1.322 to 1.285 and increases the R^2 from 0.955 to 0.957, and the pair effects have a much smaller variance than the procurer effects from the log-linear model (results available from the authors upon request).

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

A.2 Comparison to existing estimates of individuals' and organizations' effects on output

How do our results compare to existing estimates of the extent to which individuals and organizations affect output in other settings? While we are not aware of comparable estimates of the causal effects of workers and organizations on output in a developing-country government context, several studies are indirectly comparable. First, studying front-line service providers in rich countries, [Chetty *et al.* \(2014b\)](#) find that increasing the performance of 5th percentile American grade 3-8 teachers to 50th percentile would increase the present value of their students' lifetime incomes by 2.76 percent, and [Silver \(2016\)](#) finds that improving the performance of American emergency room doctors by one standard deviation would decrease time-of-care by 11 percent. We find that the same (relative) improvement in performance among Russian procurement officers would lower prices paid by 42.2 and 55.1 percent respectively.⁵⁷ However, teachers and doctors may differ from procurement officers in the complexity of the job performed, motivations, and many other dimensions.

Second, in studies of workers in the private sector performing a simpler task, [Mas & Moretti \(2009\)](#) and [Lacetera *et al.* \(2016\)](#) find, respectively, that increasing performance by one standard deviation would decrease cashier processing times in a U.S. supermarket chain and increase the probability of


⁵⁶Online Appendix Figure OA.3 shows the analogous figure for the largest connected set, again showing no signs of systematic patterns in the residuals.

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

cars being sold in U.S. used-car auctions by 11 and 4.3 percent, while in our case the improvement is 55.1 percent.⁵⁸ Of course, in the public sector, output is less easily measured and monitored, and so we expect greater scope for differences between bureaucrats.

Third, while their estimates do not have a causal interpretation, and they do not separate individual and organizational effects, *Bandiera et al. (2009)* find that Italian public bodies at the 90th percentile of performance pay 55 percent more than those at the 10th percentile for 21 generic goods. In our context, the bureaucrat-organization pair at the 10th percentile pays 75.3 percent less than the pair at the 90th percentile. Our effects are larger, possibly reflecting the fact that institutional constraints on bureaucrats and organizations are weaker in Russia than in Italy, or that we include a wider range of goods for which bureaucrats and organizations have a greater impact on prices. We are not aware of any existing papers estimating causal effects of individual organizations on output in either the private or public sector.

FIGURE A.1: MAGNITUDE OF PROCUREMENT PRICE RESIDUALS FOR PURCHASES BY BUREAUCRATS AND ORGANIZATIONS OF VARYING EFFECTIVENESS



The figure presents a heatmap of averages of the residuals from the estimation of equation (11): $p_i = \mathbf{X}_i\beta + \alpha_{b(i,j)} + \psi_j + \gamma_{s(b,j)} + \varepsilon_i$. The residuals are binned by vintiles of the estimated bureaucrat effect $\hat{\alpha}_b$ and organization effect $\hat{\psi}_j$ within each connected set of organizations. The sample used is the Analysis Sample (All Products) summarized in Table 1.

⁵⁸*Bertrand & Schoar (2003)* find that CEOs in the top quartile of performance achieve a return-on-assets that is about 200 percent higher than CEOs in the bottom quartile. In our context, bureaucrats in the bottom quartile save 84.6 percent relative to the top quartile due solely to the bureaucrat effects.