Andrei drives Bentley, Boris drives Lada: why corrupt states have clean agencies

by

Shaun C McGirr

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Political Science)
in The University of Michigan
2016

Doctoral Committee:

      Associate Professor Brian Min, Co-Chair
      Professor William Roberts Clark, Co-Chair, Texas A&M University
      Associate Professor Allen Hicken
      Professor Anne Pitcher

DEDICATION

Моей любимой, ненаглядной, и долгожданной Поле

ACKNOWLEDGEMENTS

Though this study was an individual endeavor and any errors are mine alone, it represents the incalculable contributions of others to my development as a scholar.

First and foremost I thank my wife, Polina Klyueva, for her understanding of the long struggle and patient resolve in helping me reach the end. We met as Fulbright scholars in Portland, Oregon, right before starting our respective programs in 2008. In that sense she has been with me on the entire journey as confidant, sounding board, and number one cheerleader. This dissertation is dedicated to her, with all my love.

For their infinite ability to withstand my constant questioning of everything since the early 1980s, I thank my parents Judy and Paul. While many of my stranger questions no doubt went unanswered, their encouragement cemented the scholarly drive within me. I am enormously privileged to have had their unconditional moral and financial support for every decision I made on the path to this achievement. Of course, their commitment to my development stems in part from their own parents, who remain a constant source of inspiration to me. Both my grandmothers (Joan and Joy) trained as nurses against the prevailing attitudes of society and family, and both my grandfathers (John and Hugh) were model Kiwi entrepreneurs who turned uncanny knowledge and scarce resources in to progress for their families. Only one, Joy McGirr, has lived to see the culmination of my many years in education, but all were unfailingly interested in what I studied and supportive of my path.

My teachers, coaches and mentors from many stages of my life have drawn out the gifts given to me by the people above, and made them real in the world. Particular

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF MAPS

Map

ABSTRACT

Andrei drives Bentley, Boris drives Lada:
why corrupt states have clean agencies

by

Shaun C McGirr

Co-Chairs: Brian Min (Michigan), William Roberts Clark (Texas A&M)

In societies where corruption is rife, why do bureaucrats in some agencies extract more than enough to get rich, while others do not? This dissertation explains why corruption varies across agencies within the same state, when the factors that explain variation across states (e.g. political competition) are weak or absent. In other words, *why in corrupt states aren't all agencies equally corrupt?*

My theory explains how these puzzling cases arise from the interaction between bureaucrats in a hierarchy. They must balance risky collusion necessary to grow the available rents, against competition to take a larger share of these spoils. Characteristics of agencies well-known in the bureaucratic politics literature change the terms of this trade-off, most importantly the degree of information asymmetry between bureaucrats authorizing corruption, and those carrying it out. The greater this asymmetry, the more difficult it is to maintain the collusion necessary to succeed at corruption. I test this and other implications using a novel dataset on Russian public procurement covering five million purchases by 110,000 agencies from 2011-2015.

<center>**CHAPTER I**</center>

Introduction: a puzzle in the study of corruption

Good governance is quite bland. In countries where politicians are incentivized to make policies acceptable to broad swathes of society, and bureaucrats implement these fairly, citizens make two generic complaints: 'it takes too long to get what I want from government' and 'I didn't get enough'. Bad governance, on the other hand, is fascinating in its intricacy and variety, and corruption is one such example. This dissertation explains a puzzle that results from taking variation in corruption within states as seriously as scholars have previously taken variation across states.

1.1   What is the puzzle?

Why in high-corruption countries are there agencies as 'clean' as the 'dirtiest' agencies in low-corruption countries? To sharpen the language further, *why in corrupt states aren't all agencies equally corrupt?*

This is puzzling because dominant explanations of variation in corruption focus on national-level attributes, such as competition between political parties,[1] and informed voters.[2] These certainly explain why Denmark experiences less corruption than Russia. Yet as I show below, some agencies in Russia and other high-corruption states

---

[1]Grzymala-Busse (2010)
[2]Persson, Tabellini and Brocas (2000)

<center>1</center>

have more of a corruption problem than others. Why? Existing theory misleads us to predict that extraction of rents by bureaucrats should not vary by agency, because a lack of effective control by robust oppositions and activist voters has left open the 'door to the safe'. Bureaucrats from *every* agency should be lining up to steal.

The problem with existing explanations is their underlying mechanisms do not allow variation by agency. To take one example, voters cannot 'kick out the rascals' from only one agency, they must eject a whole government. And as sham elections and bureaucracies fused with political parties are much more common than activist voters and robust competition, these control mechanisms are usually absent anyway.[3] This implies that much existing literature states only *necessary* but not *sufficient* conditions for agency-level corruption. We know why every agency in Russia can engage in corruption, but not why some hold back. If there is indeed significant variation in corruption across agencies, and this is systematic and not random, what we already know cannot be an explanation.

1.2 Existing evidence of this puzzle

There is empirical support for my argument that national-level explanations leave unanswered the puzzle of 'clean' agencies in corrupt states. Figure 1.1 is based on data from Transparency International's *Global Corruption Barometer 2013*, a survey of citizens in 100 countries across the globe about their experiences with corruption.[4] The question generating these data is: "In your contact or contacts with these institutions have you or anyone living in your household paid a bribe in any form in the past 12 months?" Institutions are grouped in sectors: Education, Judiciary, Medical, Po-

---

[3]Transparency International's headline result for the 2015 Corruption Perceptions Index (CPI) is "More than 6 billion people live in countries with a serious corruption problem". Source: `http://www.transparency.org/cpi2015`

[4]I thank Transparency International (2013) in Berlin for sharing these data, and their local affiliate in Russia for first showing me the detailed procurement data I use in subsequent chapters.

Sector–specific bribery rates vs overall public sector corruption,
Global Corruption Barometer and V–Dem 2013
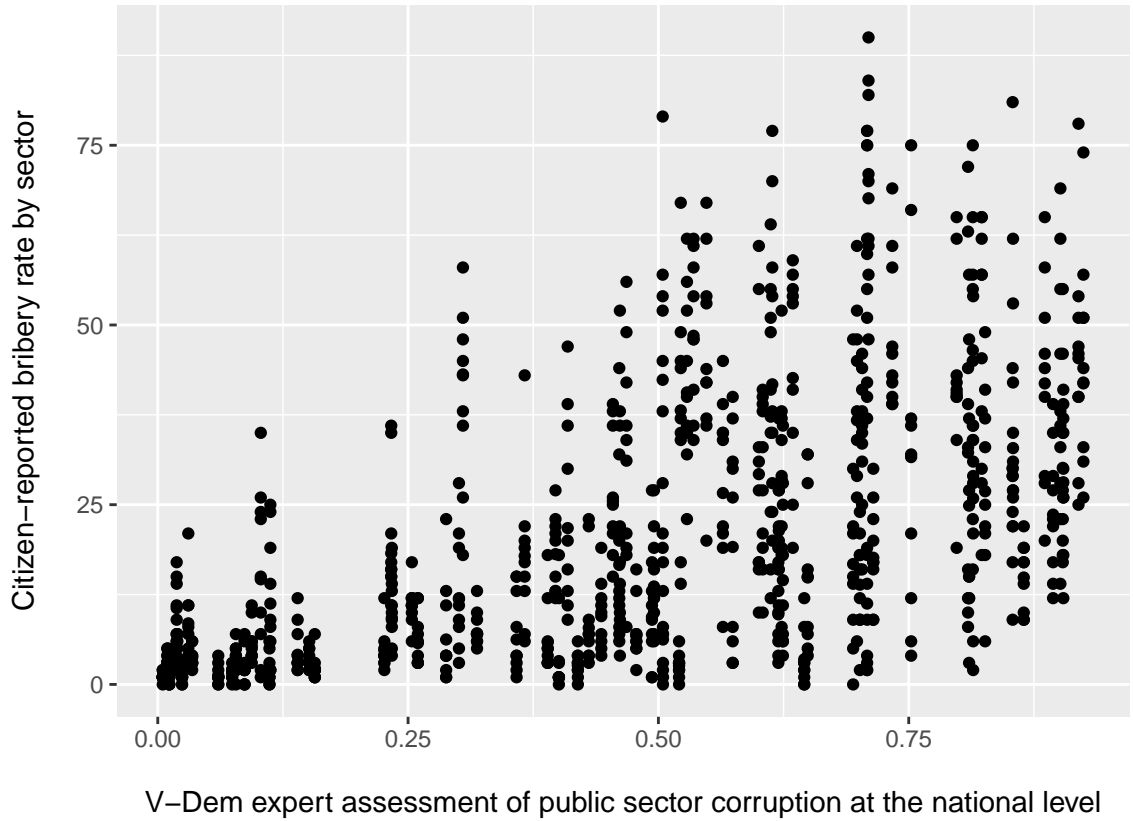
Figure 1.1: Sector-specific bribery rates in 100 states from the 2013 Global Corruption
Barometer, versus Varieties of Democracy overall public sector corruption
score for 2013. One dot represents one sector in a state (Education,
Judiciary, Medical, Police, Registry/Permit, Utilities, Tax, Land).

lice, Registry/Permit, Utilities, Tax, Land.[5] Pooling all sectors for a given state and

plotting these sector scores against that state's 2013 public sector corruption score

from the Varieties of Democracy (V-Dem) project (Coppedge et al., 2016) produces

a triangular pattern.

These data show that even the tautological argument 'national-level corruption

causes sub-national corruption' is misleading. 'Clean' states have no or few corrupt

---

[5]The use of 'sectors' in the survey is a concession to the fact states have different constellations
of agencies performing functions of the state. For now I set aside the difference between sectors and
agencies, and clarify my definition of the latter in the next chapter.

Sector–specific bribery rates vs index of electoral democracy, Global Corruption Barometer and V–Dem 2013
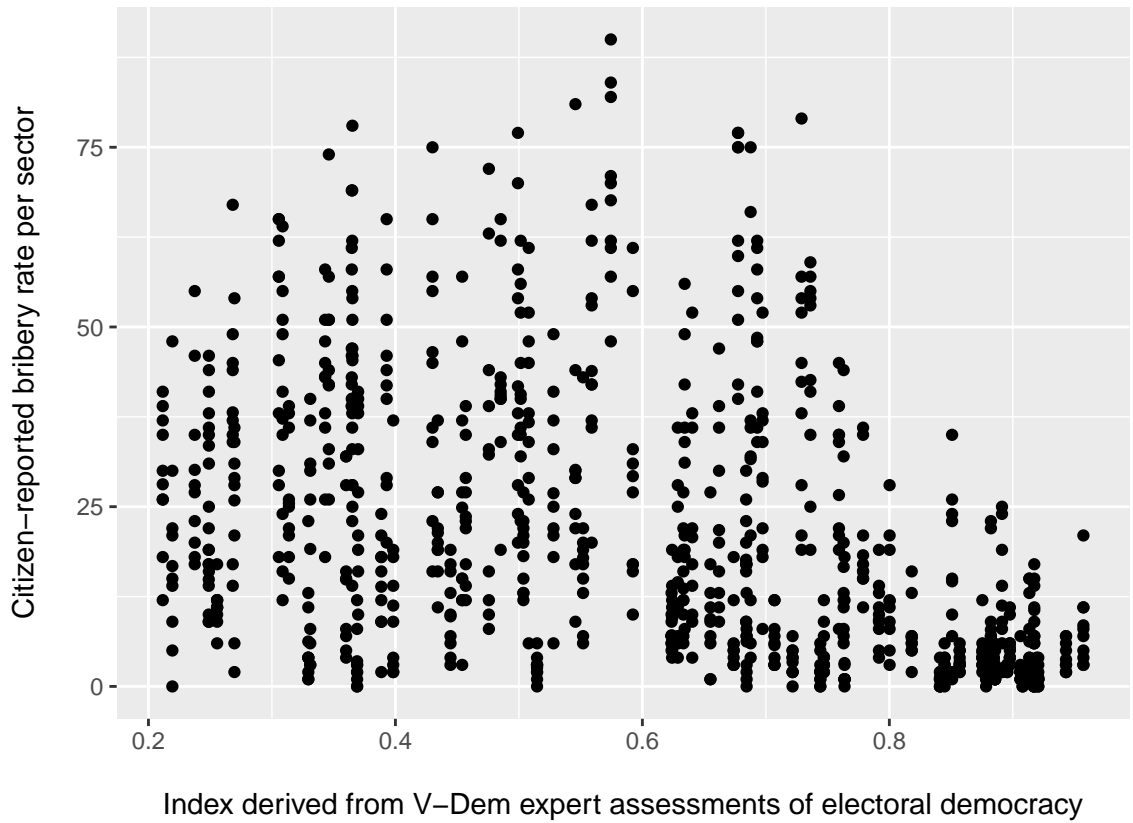
Figure 1.2: Sector-specific bribery rates in 100 states from the 2013 Global Corruption Barometer, versus Varieties of Democracy electoral democracy score for 2013. One dot represents one sector in a state (Education, Judiciary, Medical, Police, Registry/Permit, Utilities, Tax, Land).

sectors, as we would expect. Yet states with high scores have *both* highly corrupt sectors, and sectors no more corrupt than the average sector in a 'clean' state.

Other prominent explanations for cross-national variation in corruption face the same problem, sometimes worse, when used to explain why some corrupt states have clean sectors. Figure 1.2 shows the same sector rates plotted against their state's score on a summary index of electoral democracy from the V-Dem project. Only the highest-scoring democracies show low variation across sectors; states situated across 80% of the range of that measure have both clean and corrupt sectors.

Sector–specific bribery rates vs opposition autonomy,
Global Corruption Barometer and V–Dem 2013

Citizen–reported bribery rate by sector

V–Dem expert assessment of autonomy of opposition parties from ruling regime
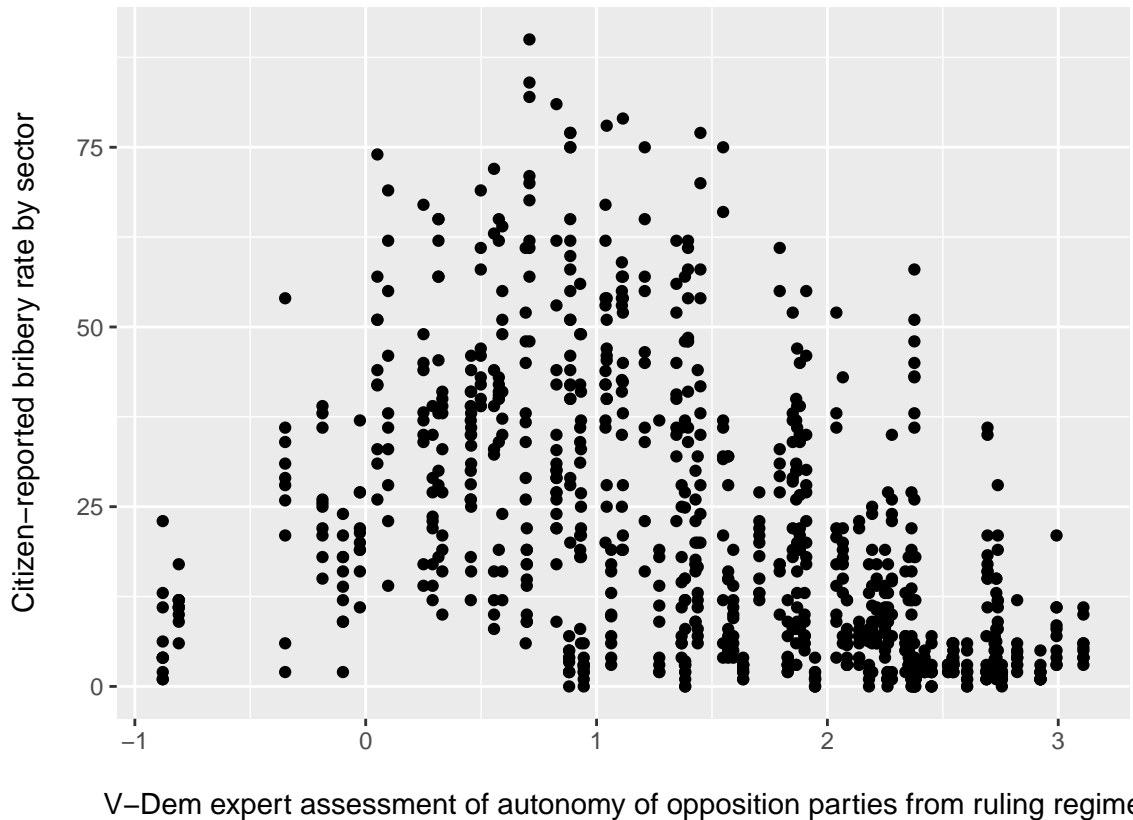
Figure 1.3: Sector-specific bribery rates in 100 states from the 2013 Global Corruption Barometer, versus Varieties of Democracy index of opposition party autonomy for 2013. One dot represents one sector in a state (Education, Judiciary, Medical, Police, Registry/Permit, Utilities, Tax, Land).

The same is true for another explanation in comparative politics, that 'robust opposition' (Grzymala-Busse, 2007) disciplines ruling parties, encouraging them to extract less from the state and in turn, restrict extraction by their bureaucrats. In figure 1.3 sectors are plotted against their state's score on the V-Dem measure of opposition party autonomy from the ruling regime, a proxy for robust competition. If anything the pattern is curvilinear: states with both very high and very low opposition autonomy tend to have only low-corruption sectors, whereas states with co-opted oppositions have clean and corrupt sectors.

This evidence suggests that an answer to my puzzle – *why in corrupt states aren't all agencies equally corrupt?* – cannot be derived from national-level explanations that focus on culture or history in a narrow sense. Nor can national-level correlates of corruption explain this within-state variation: if, for example, low economic development causes corruption then all agencies should be equally affected. In the next section I outline my agency-centric theory, and then explain why my empirical strategy focuses on all the agencies in one country, Russia.

1.3 My explanation of the puzzle

I depart from much existing literature, which tends to view corruption as a breakdown of cooperation between bureaucrats who otherwise seek to govern well. Instead, I propose that in corrupt states the *lack* of corruption in some agencies, observed above, is explained by failure to cooperate over governing poorly. My theory is conditional: agency corruption depends on opportunities for extraction *and* the availability of information about these opportunities. This improves on existing explanations, discussed in the next chapter, that directly equate opportunities with outcomes.

For an agency to be highly corrupt, there must be *robust collusion* between its bureaucrats. *Collusion* implies cooperation to extract state resources above the level preferred by some other actor; *robust* emphasizes the fragility of that cooperation in the face of a temptation to defect and extract from each other. The preconditions for this collusion vary across agencies, so sometimes it fails, producing low-corruption agencies even when 'average' corruption is high in a given state. The key feature of agencies that varies in my theory, explaining variation in corruption, is the degree of information asymmetry between senior bureaucrats whose permission is usually necessary to gain from corruption, and junior bureaucrats who carry it out. Counterintuitively based on most existing literature, corruption is more prevalent in agencies where this asymmetry is lower.

I develop this theory using a simple game-theoretic model between a senior and junior bureaucrat, which helps simplify existing explanations and demonstrate the utility of my approach. In the model the two bureaucrats decided whether and how to collude to engage in corruption under the nose of a manager. The moving part that explains clean agencies in corrupt states is the information asymmetry between the bureaucrats by virtue of the simple structure of their agency: one sets the rules while the other does the work. Unlike a standard principal-agent framework, in which the senior bureaucrat would delegate more the greater this asymmetry, delegation in my model carries a risk that makes her think twice. Watching over the agency is a manager who is not much concerned with the specific activities of the agency, as long as it does some of its job. This player is non-strategic, and simply represents an exogenous parameter which is the tolerance of the political masters of an agency for corruption. If the bureaucrats breach this limit they both face punishment. The wrinkle I introduce is to combine the delegation decision with a bargaining situation. After doing the work of corruption, the better-informed junior bureaucrat must report the 'spoils' to the senior (who may have authorized corruption), and can choose to conceal the extent of his success. As the risk of provoking the manager is increasing in the rents extracted through corruption, and the junior may effectively hide part of this risk, corruption is not always sustainable in equilibrium despite ample opportunities. The model predicts that in corrupt states, agencies where information asymmetries between bureaucrats are higher will be cleaner, despite the ample opportunities for corruption they present.

1.4   Testing my explanation against alternatives

My empirical strategy compares agencies on an activity that all engage in: procurement of goods and services. This makes agencies more comparable, despite their widely varying missions and purposes. Procurement also provides a setting for indi-

vidual bureaucratic decisions that maps neatly to my theory. I use comprehensive data from Russia's centralized public procurement system covering 6 million purchases by 140,000 individual agencies across 2011-2015. These permit construction of detailed measures of the key concepts in my theory, which I use to test its implications at the level of individual purchases and agencies, within Moscow, within all regions, and within and across regions simultaneously. I find broad support for my theory at both the purchase and agency levels.

Despite my focus on Russia for motivation and data, I build a general theory of cross-agency variation in corruption that should apply to any state. In that sense, Russia is but one of the great majority of states in the world where corruption is "not a bug, but an essential feature of the operating system"[6].

1.5   Structure of the dissertation

In Chapter 2, I carefully define the key terms of this study (corruption, bureaucrat, agency) and use these to develop the logic underlying a well-functioning procurement market. This becomes my foil for the rest of the dissertation. I then review the literature to understand what we already know about threats to such a market, crystallizing diverse studies in to common hypotheses that inform my own theory and the empirical approach. Finally, I tie together these hypotheses and my own insights in to the model described above.

In Chapter 3 I propose a research design for evaluating these implications. I settle on public procurement as a source of data for two key reasons: every agency engages in it, which reduces heterogeneity between agencies in their activities, and corruption in this process has fairly clear implications for what should be observed empirically. I generate a large dataset myself from the universal public procurement system in

---

[6]Brian Whitmore, *License To Steal - A Bug Or A Feature?* ,http://www.rferl.org/content/license-to-steal-a-bug-or-a-feature/24795746.html

Russia that operated from 2011-15, which contains about 6.5 million purchase records from almost 140,000 different agencies across 84 regions. For the remainder of Chapter 3 I develop several core measurement concepts that help to operationalize my theory for this context, and provide preliminary tests at the agency level in Moscow.

Chapter 4 continues the empirical work by zooming in to the purchase level to learn about procurement behavior at the micro-level. I estimate models of a well-functioning procurement market and show how the presence of certain 'red flags' at the purchase level change the relationship between variables in the marketplace. These measures are grounded in the behavior of Russian bureaucrats running procurement, and informed by detailed implications developed from my theory. This allows me to very closely map measures on to the discretionary decisions of bureaucrats, and demonstrate systematic variation in corruption at the purchase level within Moscow and across Russia. Then I aggregate these purchase level measures to the agency level, and show that the predictions of my theory are borne out in general across Russia's regions. The final set of tests in this chapter grow from the insight that while sharing the same broad dynamics in terms of their cross-agency variation, Russia's very diverse regions have different underlying baseline levels of corruption. Instead of relegating this variation to fixed effects I join the implications of my theory to those of the cross-national literature to test within-region and across-region variation in corruption for the first time.

The concluding chapter discusses my findings in the context of the literature and outlines the contributions the study makes on several fronts. I discuss future work to understand how agency structure and bureaucrats' incentives interact to produce the wide variation in corruption uncovered by seeing inside the state.

CHAPTER II

Theory: why corruption fails against the odds

"If you see the oppression of the poor, and the violent perversion of justice and righteousness in a province, do not marvel at the matter; for high official watches over high official, and higher officials are over them."
Ecclesiastes 5:8, New King James Version

In the previous chapter I motivated and illustrated the puzzle of this dissertation: why are there low-corruption agencies in high-corruption states? Here I review the relevant literature to understand what we know about the causes of corruption, and why these vary across agencies. This requires defining terms, most importantly *agency* and *corruption*. In the process of reviewing literature I extract implications from existing work, and build on these to create my own theory, which I formalize in a game-theoretic model at the end of the chapter.

2.1   Precis of theory

The intuition behind my theory is simple. As with any complex agency activity, bureaucrats can improve their returns from corruption by cooperating. This pools their knowledge about the opportunities for corruption, helping those more senior and further from the 'coal face' understand what is possible, and helping those more

junior understand what is permitted. What separates this cooperation from its analog in standard delegation models is its collusive nature. As corruption is illegal, discoverable, and punishable by the agency's political master, it must be hidden to some extent.[1] This leaves bureaucrats with a dilemma: cooperate to grow a lucrative stream of illicit benefits (carrying some risk of punishment), or forgo corruption altogether, avoiding risk. If corruption is understood as a simple transaction, e.g. a bribe for a scarce permit, cooperation between bureaucrats need only be minimal and their decisions can be easily understood by extending implications of Becker's (1968) crime model.[2] More lucrative corruption, especially when the agency is spending funds, requires more elaborate schemes than a bribe taken by a junior bureaucrat and shared with a supervisor. Willing suppliers must be found, and procedures manipulated in order to succeed. These more complex arrangements for extracting agency funds on a grand scale depend on the expertise of bureaucrats implementing them.

The more expertise required in an agency,[3] the larger the information asymmetry between those organizing corruption and those permitting it. The logic is the same as in any delegation model: proximity to a decision grants better information. As this asymmetry grows, so does the difficulty of 'dividing spoils': the same factor that improves total gains from corruption (i.e. expertise) also gives junior bureaucrats organizing corruption a means to hide the extent of their success. They can engage in corruption, extract X, but claim to have extracted only $\frac{X}{2}$ when reporting back and sharing spoils. It is difficult for a senior bureaucrat to verify this report, due to information asymmetry, and the inherent risk of auditing an illegal activity for any reason other than deterrence.[4] In the limit, those atop an agency have no idea what their subordinates do and cannot believe any report about the spoils.

---

[1] As explained below, this does not require an assumption that politicians want zero corruption. Furthermore, if anti-corruption is costly, every leader must permit some level of corruption.

[2] The bureaucrat receiving the bribe pays a 'cut' to their supervisor to stave off punishment.

[3] I explain further below what I mean by 'required'; in the case of procurement the mission/purpose of an agency makes certain purchases more or less necessary.

[4] An investigation could reveal that the senior bureaucrat permitted corruption.

Unlike a standard bargaining model, in which the party with less information must accept whatever is offered by the other party, another factor influences their decision here: the risk to all involved should corruption be discovered and punished. At some point this risk of punishment outweighs the expected gains from delegating corruption, and the fragile collusion between bureaucrats breaks down. Clean agencies are the result, by this logic, of a failure to successfully collude; corruption flourishes where the 'expertise problem' is less severe. As I show below, this is counter-intuitive because the agencies where expertise is most required are those where *opportunities* for corruption are indeed greatest, yet the expertise problem means these opportunities go wasting. The ideal agency for the corruption-seeking bureaucrat is not one buying military hardware, of which there are few sellers (and therefore low price competition) but one buying generic goods over which bureaucrats can sustain robust collusion.

This intuition leads to a model fusing two well-known classes of model: delegation and bargaining. First, a senior bureaucrat decides to what extent she *delegates* to a junior bureaucrat permission to engage in corruption. Second, after corruption plays out (or not) the bureaucrats must *bargain* over any spoils. Delegation comes with two uncertain costs for the senior, realized in the bargaining stage: punishment by her own manager for the agency exceeding some threshold of tolerable corruption, and concealment of the true spoils by the junior bureaucrat. These force a departure from the 'delegate more, the less you know' logic in standard delegation models.

My theory contributes to the corruption literature as a 'broker' between disparate and poorly-connected fields. It builds on comparative politics, institutional economics and the political economy of bureaucracy. As I demonstrate below, each partially explains my puzzle in isolation, so I offer a simple framework for tying them together. My goal is not to explain all variation in corruption across all cases, but to resolve my puzzle about clean agencies in corruption-prone contexts.

2.2 Definitions

Before reviewing the literature I define several key terms: corruption, rents and rent-seeking, bureaucrat and agency. My goal is to outline where and why I deviate from standard definitions, not an exhaustive treatment of the respective debates. A clear definition is important for any measure, and particularly for measures of difficult-to-observe phenomena such as corruption.[5] I define cooperation and collusion in the specific context of my theory later, during exposition of the model.

2.2.1 Corruption

Corruption shares in common with many social science concepts significant conceptual flexibility, and debate over what constitutes corruption is long-lived. Since it reemerged as a topic of interest to social scientists in the 1960s (Nye, 1967; Scott, 1969), scholarly consensus has oscillated over whether a universal definition is more appropriate, or whether corruption should be defined in terms of locally-acceptable behavior. Early scholars recognized that a legalistic definition would depend too much on the specific legal regime in a given state, but that a definition couched strictly in terms of behavior would require extensive auxiliary definitions and examples to make it concrete. Nye swung toward the latter with his definition:

> "Corruption is behavior which deviates from the formal duties of a public
> role because of private-regarding (personal, close family, private clique)
> pecuniary or status gains; or violates rules against the exercise of cer-
> tain types of private regarding influence. This includes such behavior as
> bribery (use of a reward to pervert the judgment of a person in a position

---

[5]My general approach follows the recommendations of (McMann et al., 2016) for constructing valid measures. First, assess *content validity* in the process of defining the measure. Second, judge *data generation validity* by the clarity and reproducibility of the process. Third, evaluate *convergent validity* with evidence that the measure distinguishes between important cases and correlates as expected with other measures. The second and third steps are left to the following chapters.

of trust); nepotism (bestowal of patronage by reason of ascriptive relationship rather than merit); and misappropriation (illegal appropriation of public resources for private-regarding uses)." (Nye, 1967, p. 419)

This definition covers many activities of elected and appointed officials (politicians and bureaucrats, respectively) that scholars have labeled corruption. What these activities all have in common is that some private consideration interferes with the official's public job. Subsequent literature, culminating with Rose-Ackerman (1999) has condensed Nye's formulation to "the misuse of public office for private gain".[6] This now-standard definition brings to mind examples that most could agree are corruption, for example the issuing of a license to an obviously unqualified candidate in exchange for a bribe. This is useful for giving the concept relevance, as any observer can think of an example of such misuse. What this definition fails to do is specify where 'misuse' begins, or how much 'private gain' is material, or which 'public offices' are covered. I do not aim to resolve these difficulties, but to point them out and build a case for an alternative definition.

My case hinges on two claims supported by my analysis of the literature below. The first is based on theoretical parsimony: corruption defined as broadly as above encompasses too many different processes. Vote-buying, kickbacks, nepotism, or simple direct theft of public funds all qualify as 'misuse' but involve different sets of actors in varying settings. Relying on 'misuse' to tie these together is appropriate for very broad research questions, for example how public perception of corruption affects confidence in political systems; citizens are likely to group these activities together when answering those questions. For narrower research questions such as mine, a better definition of corruption is tied to the decision-making process under study, not subjective interpretation of what constitutes 'misuse'.

---

[6]Gerring (2001) calls this a *minimal definition*: it is the common core of almost every definition of corruption, despite their differences.

The second advantage of a narrower definition is the reduced distance between my theory and the available data. When corruption was primarily studied using national-level perception measures, it made more sense to use an all-encompassing definition because the available data were likewise blunt. While this allowed some progress in identifying correlates of corruption (see Treisman, 2007), several scholars have identified biases arising from this approach.[7] With more fine-grained data, we are freed to go beyond the standard definition and build narrower theories that explain the behavior of fewer actors (not 'all officials'), but with greater precision.

Following Rothstein and Teorell, I define corruption as *any violation of the norm of universality*, which states that when making decisions, "government officials shall not take into consideration anything about the citizen/case that is not beforehand stipulated in the policy or the law."[8] This strict definition offers no ethical gray line, which is appropriate for this context.[9] This dissertation studies corruption in one setting, public procurement, in which bureaucrats extract agency funds precisely by taking in to consideration factors outside the relevant procurement laws in their jurisdiction.[10] Instead, they might award the contract to a friend, or the firm paying the highest bribe, and public procurement laws never allow such factors to be taken into consideration.[11] As elaborated in the following empirical chapters, my measures of procurement corruption are consistent with this definition: they try to infer how

---

[7]The problems with perception-based measures are well known in the literature: see Olken (2009), Andersson and Heywood (2009) and Rose-Ackerman (2006) for an overview. Many result from lack of specificity over what constitutes corruption in survey questions, so respondents interpret questions differently based on their exposure to corruption. This is acceptable for studying perceptions of corruption, but not for comparing actual underlying corruption across states or agencies.

[8]Rothstein and Teorell 2008, pg. 170

[9]Note that this definition does not rule out discretion by bureaucrats when they decide a matter; it simply requires the bounds of this discretion be well-specified beforehand.

[10]This definition is equally useful for other types of corruption, for example 'petty' or 'retail' corruption (typically bribe solicitation) carried out by street-level bureaucrats without access to public funds. Despite their different contexts and motivations, these are both examples of the *deviation from specified procedure* at the heart of Rothstein and Teorell's definition.

[11]This definition also helpfully sidesteps the debate over whether corruption is efficient overall (Méon and Weill, 2010). Even those who view corruption most favorably, arguing that bribes serve as 'speed money' when firms face onerous regulation, would have difficulty arguing that bureaucrats controlling public procurement should choose suppliers on factors other than price and quality.

much bureaucrats depart from the norm of universality by capturing deviations in behavior from what is expected under 'normal conditions'. The theory I develop below explains variation in these measures, building on what existing literature implies should encourage vs discourage corruption according to this definition.

2.2.2 Rents and rent-seeking

A rent is 'a return in excess of a resource owner's opportunity cost', and in well-functioning markets, rents are 'short-lived...because competition will drive them to normal levels' (Tollison, 1982). Rents arise in two ways: changes to demand or supply create a temporary opportunity for additional profit;[12] or because an actor intervenes to create rents artificially through manipulating demand or supply.[13] As recognized by many besides Tollison, including Krueger (1974) and Tullock (1967), there is competition both for short-term natural rents signaled through the price system, and for artificial rents created by intervention. Competition over natural rents is motivated by standard profit-seeking, whereas competition to create, capture and defend artificial rents is labeled rent-seeking. Firms pursue both.

Rent-seeking and corruption are often treated as interchangeable terms, but it is useful to distinguish them, following Lambsdorff (2002). Rent-seeking is a *motivation:* all economic actors would prefer a long-term advantage engineered in their favor and protected from dissipation by competitive pressures. Any time a government subsidizes production of a particular good, or rewards procurement contracts to a more expensive local supplier over a cheaper international competitor, it creates such a rent. Rents do not necessarily lead to corruption under the definition I adopted above, because the official creating the rent may be empowered by policy to do so; this is the second case in Table 2.1.

---

[12]Tollison calls these natural rents as they arise and subside only through the price system.

[13]Supply can be manipulated, for example, by government restricting the set of actors allowed to produce a good. Demand for a service could be altered by subsidizing it to lower the cost to consumers, or by taxing alternatives.

Case	Example
No rent created, no corruption	Agency awards procurement contract to the bidder best meeting the specified requirements; over time any 'natural' rent dissipates due to competition
Rent created, but no corruption	Agency favors local bidder for procurement contract because of a specified policy tilting the playing field; direct beneficiary is a rent-seeking firm
Rent created with corruption, transfer not realized	Agency favors local bidder for procurement contract due to any factor outside the policy (e.g. personal connection); direct beneficiary captures entire rent
Rent created with corruption, transfer realized	Agency favors local bidder for procurement contract due to any factor outside the policy; direct beneficiary and decision-maker share the rent

Table 2.1: Distinguishing rents from corruption

Corruption, on the other hand, is a *technique* bureaucrats can use to capture rents for their own benefit. Under the Rothstein and Teorell definition, corruption occurs during the creation, capture or defense of rents when officials making decisions take in to consideration factors not specified beforehand. When officials award procurement contracts to a 'preferred supplier' when policy does not allow for this, or hire a relative over a more qualified candidate, they engage in corruption that creates a rent for the direct beneficiary of that decision. Corruption only benefits decision-makers, however, *if they realize a transfer of resources from the beneficiary.*[14] Whether the decision-maker benefits distinguishes between the third and fourth cases in Table 2.1.

I will clarify this distinction between the third and fourth cases when I operationalize these definitions in the next chapter. I will argue that certain behavior observed in public procurement makes the creation of rents through corruption more likely, but does not necessarily indicate that bureaucrats extracted any benefit. In the language of my empirical setting, the very existence of procurement creates 'corruption opportunities' from the perspective of bureaucrats. Whether these opportunities are

---

[14]The mode of transfer is most commonly modeled as a bribe to induce the decision favoring the beneficiary; all that is necessary for my theory is that a transfer is made possible by the corrupt creation, capture or defense of the rent. I rely on intuitions from prior literature that explain the interaction of bidders and bureaucrats, and do not model this part of the overall process.

'converted' is a function of *other* observed behavior that is consistent with bureaucrats trying to ensure a transfer of corruption-driven rents in to their own pockets.

The logic outlined above makes the creation or existence of rents a necessary, but not sufficient condition for corruption. There must be the opportunity to intervene and create an excess return, before officials can extract some of this for themselves, but the latter is by no means guaranteed.[15] This explains why states with more rent-prone economies tend to experience more corruption on average, and below I extend this logic to the agency-level variation observed in Chapter 1. In the case of procurement (which systematizes rent-creation opportunities) there is a further twist: the portion of an agency's budget *not* spent yields no potential corrupt benefit to the bureaucrats making spending decisions. This only heightens the temptation to spend it on rent-creation, which creates potential for corruption.

To summarize the distinctions drawn above, rents do not always lead to corruption. When they do, it is because officials seek to benefit, but this benefit is not always realized. From this perspective, rent-seeking motivates intervention in the economy, and corruption is a mechanism used by officials to generate returns for overstepping their bounds in creating those rents. Applied to procurement, spending the agency budget on certain goods and services creates the *opportunity* for bureaucrats to gain, while corruption provides the *means*. The theory I develop below explains motives.

2.2.3 Bureaucrat and agency

In my theory a *bureaucrat* is anyone empowered by government to make decisions about implementation of policy, and who is not elected.[16] This includes everyone from local clerks deciding who gets a permit for some activity (a 'street-level' bureaucrat), to the top unelected official running a ministry or department. What they share

---

[15]This excludes from my theory outright theft, in which an official directs budget funds directly to their pocket. Corruption only occurs when a decision is made about a citizen/case.

[16]Elected officials also make such decisions but my puzzle is about the variation in corruption among actors not directly subject to electoral pressures.

is power over the application of policy to citizens and firms, which puts them in a position to create rents and engage in corruption according to my definitions above.

An *agency* is any collection of bureaucrats who report to the same bureaucrat above them in a hierarchy. This flexibility allows agencies to be nested within each other, an approach first taken by Tirole (1986). For example the Federal Bureau of Investigation and the Antitrust Division are both agencies within the US Department of Justice, which is also an agency by my definition. The question of where to set the line on a government's organization chart, for the purposes of empirical analysis, is left to my research design in the next chapter. For now, I build a theory that can apply to the decision over rents and corruption at any level of the hierarchy within a government. The simplest agency is the one I model at the end of this chapter, and contains one junior bureaucrat accountable to one senior bureaucrat, who is the highest unelected official. This agency is empowered by an elected official, the manager, who plays a passive role in my model for the sake of simplicity.

2.2.4   Limitations and scope conditions

While the definitions I adopt depart subtly from some standards in the literature, I consider the gains from sharpening the set of actors, goals and behavior under study to be worthwhile. In one sense these definitions result in a narrower set of candidate phenomena to test my theory empirically: some clear types of 'misuse' fall outside corruption, and I intentionally exclude politicians. On the other hand, because my favored definition of corruption is *procedural* (about the way decisions are made over who gets which rents) I can avoid splitting hairs over, for example, whether a bureaucrat *asks* for a bribe or the citizen/firm *offers* it. For the purposes of explaining my puzzle, I only need there to be opportunities for bureaucrats to decide to create rents and a mechanism (corruption) for extracting some of this value.

Though the explanation I offer is general, and applies to any decision by which

a bureaucrat can create or distribute a rent, it is particularly suited to the setting for my empirical tests: public procurement. I defend the usefulness of this context in the next chapter, but at this stage it is worth noting that governments introduce procurement procedures precisely to *reduce* rents for firms, who want to preserve them wherever possible (Lambert-Mogiliansky, 2011). Firms can defend these rents by influencing decision-making bureaucrats, who in turn may try to extract a portion of the rents in exchange for a corrupt decision. As public procurement introduces tighter rules for agency spending, it threatens a lucrative source of rents, suggesting it is fertile ground for observing corruption, i.e. deviation from such rules.

2.3 Literature review

My puzzle arises from the limitations of prior work explaining variation in the prevalence of corruption. That literature is large and ever-expanding, and so my review is not exhaustive.[17] Instead, I favor deeper discussion of the literature more relevant to my puzzle about within-state variation, which results in less time spent on explanations of cross-national variation. I spend the most time discussing literature that helps explain the puzzling cases: clean agencies in corruption-prone states.

I begin my review by identifying what we already know about the demand for corruption by citizens and firms interacting with agencies. Then I ask the literature what encourages vs constrains corruption across several nested levels of analysis: the individual bureaucrat facing these external demands; the bureaucrats working together within an agency; the agency with respect to its political principal; and the political 'owners' of agencies within the wider political system. As I move from the individual bureaucrat deciding whether to engage in corruption in a specific case, to the individual voter deciding how to weigh corruption in her vote choice, I exclude

---

[17]For summaries of the key questions, answers and problems as they have evolved through time, see: Olken and Pande (2012), Svensson (2005), Aidt (2003) and Rose-Ackerman (1999).

several existing explanations as unhelpful for my puzzle, identify others as compelling alternatives to my own, and note the gaps that my theory should fill.

2.3.1 Why do citizens and firms demand corruption?

The definition of corruption I adopted above is procedural: corruption occurs when officials take in to account factors beyond their discretion when dealing with a specific case. This implies the allowable set of factors sometimes produces decisions that are inconvenient for others, creating demand for corruption in the short-term, or lobbying to change policy in the long-term.[18]

That agency decisions create demand to bend or change rules is unsurprising, given these decisions create and distribute rents. Until a firm receives a necessary permit it cannot benefit from the excess returns afforded only to permit-holders; once it receives the permit it would prefer the rules be altered to prevent the granting of further permits and the natural dissipation of rents.[19] In public procurement, larger bidders benefit if stringent pre-qualification requirements are attached to a tender announcement, because the added burden reduces competitive pressure, creating a rent. In response, a smaller bidder could attempt to compensate the bureaucrat for selectively relaxing some of the costly requirements. In both the permit and procurement examples, the rent exists for firms to seek *because* the bureaucrat implements some rule. The larger this potential rent, the more fierce will be the competition to encourage its creation, direct its allocation, and defend its existence.[20] If bureaucrats can link outside-the-rules access to rents (i.e. corruption) with obtaining a cut for themselves, they can profit. This logic generates two hypotheses that are common in the corruption literature:

---

[18]There is some evidence corruption and lobbying are substitutes (Campos and Giovannoni, 2006). My theory focuses on the bureaucrats in an agency, so I exclude the possibility of lobbying.

[19]Aidt and Dutta (2008) present a model in which increasing entry regulation reduces total output, and generates corruption as entrepreneurs bid in competition over licenses.

[20]As Tullock (1990) points out, rents are captured cheaply compared to their initial value, which can quickly decline. Agency procurement, however, creates new rents over time.

L1a: Rents are larger the more government intervenes in the economy[21]

L1b: Demand for corruption increases in the size of rents

These hypotheses are uncontroversial and well-supported by existing research, so I do not test them.[22] I introduce them as a baseline for later discussion, as they help explain why some states are more corruption-prone than others: for various reasons, in some states bureaucrats intervene in the economy more, which creates larger rents. There are two hypotheses instead of one, because unlike other literature, I do not draw an unconditional link between the scale of intervention and corruption. Large rents are only a *necessary* condition for extensive corruption; other factors discussed below can prevent bureaucrats from turning rents in to corruption.[23]

The relevance of this logic to my puzzle is that agencies specialize in particular kinds of interventions, based on their 'mission' or purpose (Wilson, 1989), so it is likely that the extent and nature of intervention in the economy varies by agency. This generates agency-level versions of the hypotheses above:

L1c: Agencies intervening more in the economy generate larger rents

L1d: Demand for agency corruption increases in the rents it generates

These two mechanisms linking availability of rents to *demand* for corruption lead to another hypothesis, that completes the link between rents and corruption-derived

---

[21]I number hypotheses generated directly from the literature review using *Ln*, and summarize these in a table at the end of this section, reserving *Hn* for those derived from my own theory.

[22]For evidence of the link between government intervention in the economy and corruption, see Mauro (1995), Keefer (2007) and Mohtadi (2003).

[23]Singapore provides an illustrative case. The state invests heavily in the economy and in certain sectors government intervention is much higher than the average for similarly wealthy states. Despite the significant rents created, however, Singapore is widely acknowledged to have 'solved' its early corruption problems and is a poster child for bureaucratic reform. This does not mean nobody benefits from the rents, simply that they are not created or distributed by bureaucrats making exceptions to the rules. This lack of exceptions reduces opportunities for corruption.

benefits to bureaucrats, and which I will test:

L1e: Bureaucrats' private returns on corruption increase monotonically with the opportunities for corruption generated by their agency's administration of rents

This hypothesis emphasizes that increasing opportunities for corruption are, at the very least, never *bad* for those looking to profit from corruption. If this is true, then the answer to my puzzle is simple: clean agencies in corrupt settings are those creating the least opportunities for corruption, because they make fewer rent-creating interventions. This contention is a foil for my theory, which partly contradicts it. First of all, even if *L1e* is true, we need to explain why agencies create more or less rents in the first place, otherwise the problem of inferences about cross-agency variation simply shifts. Second, *L1e* ignores the problem of what to do with any private returns generated by corruption: because these returns are the product of collusive cooperation (not just a tax) this division is not straightforward, and in my theory this strategic interaction modifies the relationship specified in *L1e*.[24]

A study by Hunt (2005) specifically addresses variation in agency-level demand for corruption, using household survey data from Peru to calculate corruption rankings for 21 public institutions. Those with the highest adjusted bribery rate were the judiciary and police, with 42% and 27% of reported bribe revenues respectively, despite only accounting for 2% of interactions each.[25] The explanation for this is consistent with *L1e* above: these institutions have both monopolistic and very interventionist positions in the economy; few others can seize assets or detain individuals. Bribery rates were near zero for some agencies, even in Peru, a mid-ranked state in most national-level corruption measures; this further motivates my puzzle. In contrast

---

[24]My theory predicts that conversion of rents (corruption 'opportunities') into private returns from corruption in an agency is weaker the more expertise is required to discharge its mission/purpose.

[25]Reports of bribes are weighted by the extent of respondent interaction, as "an institution might appear to be relatively honest simply because few households even interact with [its] officials".

to my framework, however, the study focuses on characteristics of agency 'clients' for explanatory power, leading to the conclusion that "bribery rates are higher at institutions whose clients have a bribe-prone profile" (p.5).

Hunt's study answers my puzzle with "clean agencies in corrupt states are those with less bribe-prone clients", but does not explain why citizens select in to the client pools of various agencies. Without holding the pool of those offering bribes fixed across agencies, the theory and empirical approach explain bureaucratic structure more than they do the choices of bureaucrats within that structure. In part this is natural consequence of data collected from the perspective of potential bribe-payers: we observe little else about the interaction besides the bribe and reported amount. When developing my empirical approach in the next chapter, I discuss the ideal characteristics of a dataset to tackle my puzzle, and advance public procurement as a candidate. Its clearest advantage over bribery data is that procurement is undertaken by every agency, and variation in 'intensity' of this corruption-prone activity is easily measured by the amount of spending, rather than rates of contact with officials, where selection is endogenous. This means variation in observed behavior is driven much less by the idiosyncratic nature of each agency's intervention in the economy, which is largely predetermined prior to bureaucrats' tenure.[26] In turn this forces me to develop a comprehensive theory of agency-level corruption from first principles, with a cause driven by they way bureaucrats interact with each other.

To summarize, agencies vary in corruption because they intervene differently in the economy, which creates different levels of demand for corruption by those affected. This relationship is not as simple at the agency level as it has typically been construed at the national level, because of the way bureaucrats within agencies *react* to demands for corruption. I turn now to explanations of these reactions.

---

[26]A Ministry of Justice would have difficulty justifying the purchase of nuclear equipment, but like any agency it needs office supplies.

2.3.2 How do individual bureaucrats respond to demand for corruption?

In the section above I outlined what existing research shows about the background level of demand for corruption in a society. Here I examine explanations of variation in the response of individual bureaucrats to these demands, and how these might vary systematically across agencies.[27]

The canonical treatment of an individual's decision to engage in criminal activity is Becker's decision-theoretic crime model (1968). Its key implication is that all else equal, increasing the costs of a crime decreases the probability individuals commit one. Potential criminals weigh these costs against the benefits they anticipate from crime and the opportunity cost of time spent on legal income-generating activities. Applied to bureaucrats, the balance can be tilted towards 'not corruption' by increasing the value of the inside option (their wage), decreasing the value of the outside option (through punishment) or increasing the probability of detection.

The implication of this basic model that has manifested most prominently in the corruption literature is the *efficiency wage hypothesis*, expressed by Becker and Stigler (1974) as "discourag[ing] malfeasance by raising the salaries of public enforcers".[28] Paying bureaucrats who implement policies a higher salary increases their future stream of income, lowering the temptation to engage in corruption that may (subject to detection and punishment) lead to loss of that future benefit. This logic can be summarized in the following hypothesis, and its extension to the agency level:

L2a: Paying a bureaucrat a higher income decreases the probability she engages in corruption for her personal benefit

L2b: Paying an agency's bureaucrats a higher income decreases the probability of

---

[27]As I alluded to above, the bureaucrat may have directly created demand for corruption by erecting arbitrary barriers; in that case the bureaucrat should be very responsive to demand.

[28]I use salary and wages interchangeably, as does the literature. This is unproblematic in the case of bureaucrats, who generally do not work overtime.

corruption for personal benefit in that agency

L2a suggests an important determinant of corruption is the remuneration of those who might engage in corruption. This association has been studied extensively in theoretical literature, and is a feature of most models where individual bureaucrats appear as actors, but empirical tests have been rare. Di Tella and Schargrodsky (2003) tested whether the effect of a corruption crackdown at public hospitals in Buenos Aires varied by wage rates, finding no effect during the initial crackdown, but a strong association between higher wages and lower corruption later. Niehaus and Sukhtankar (2013) extended this logic to posit that *any* future income stream should have the same effect on current corruption, *including future corruption*; they found strong evidence for such a 'golden goose effect' during an Indian wage reform.[29]

There are two obvious concerns in extending the efficiency wage hypothesis from the individual to agency level (i.e. *L2b*): reverse causation, and selection effects. The underlying causal argument is that higher wages reduce corruption by changing bureaucrats' calculus, but even if this were true on average for individuals in an agency, existing corruption likely causes wages at the agency level. When politicians set wages they could be wary of stretching budgets to further compensate those already earning 'on the side' from corruption, or they may actively punish agencies for revealed corruption that reflects poorly on the government. Both mechanisms bias estimates of *L2a* towards danger. The selection problem is the same as for any aggregation of individuals who might vary in their willingness to engage in some behavior: over time those agencies where bureaucrats engage in more corruption attract like-minded candidates, reinforcing the problem of reverse causation.[30]

---

[29]The same logic underpins the 'stationary bandit' theory in Olson (1993): political leaders restrain their own extraction from the economy if they value its future production highly enough.

[30]Several formal papers outline such a 'contagion' effect, in which only a small amount of initial corruption 'seeds' an agency, encouraging ever more (Cadot, 1987; Lui, 1986). Much of this literature sets up the principal's problem in terms of recruitment; she faces a continuum of potential agents who vary in their susceptibility to corruption. I avoid this formulation as the answer it provides to

These weaknesses translating the efficiency wage hypothesis to the agency level are not only theoretical. They also obstruct an empirical approach to my puzzle based on wage rates, whether observational or experimental.[31] Nevertheless, due to the prevalence of this hypothesis in the literature, I present a limited test of *L2b* in Chapter 4, showing that average wages in the Russian federal ministries are not correlated with my measures of agency-level corruption.[32] More constructively, as the Becker model is the progenitor of the efficiency wage hypothesis, I use its core tension in my own theory: *balancing between inside and outside options under the shadow of detection.* I extend its decision-theoretic implications to my game-theoretic context by making the success of crime (i.e. corruption) depend on collusion with another party who doubts the reported gains from rule-breaking. This illuminates the dual role of any senior bureaucrat toying with corruption: they are subject to detection and punishment by someone above, while playing the same role with respect to the junior bureaucrat below, who has the same incentive to conceal information from superiors.

In the context of procurement (the setting for my empirical tests), one model offers further detail on bureaucrats' responses to demands for corruption. Lambert-Mogiliansky and Sonin (2006) build a model in which returns to the bureaucrat for allowing 'adjustments' to already-submitted bids are linked to the extent of collusion among firms. The more collusion, the more government will need to pay for a given good; this in turn places more value on the bureaucrat's decision, and she can ask for higher bribes. In equilibrium, corruption by the bureaucrat even *enforces* collusion that results in higher prices: firms prefer transferring a substantial portion of the

---

my puzzle is trivially true: bureaucrats with lower 'susceptibility to corruption' (a trait *not* well-established empirically) choose cleaner agencies, which remain clean. Changes in corruption at the agency level are, by extension, due to changes in the composition of types within them.

[31]The theoretical holes due to reverse causation and selection effects become insurmountable problems for observational studies, while experiments at the agency level would be resource-intensive.

[32]I also include it as a regional-level control in the multi-level model the concludes Chapter 4.

rents they earn to the bureaucrat, over a price war. This model and its companions[33] generate useful hypotheses at the individual level, which can be extended to agencies:

L2c: Procurement market features that aid collusion between bidders increase the returns to corruption for bureaucrats by decreasing competition between bidders

L2d: Agencies facing procurement markets with lower levels of competition are at greater risk of corruption

My theory developed below does not include firms as actors, for two reasons. First, existing models already provide clear implications for how market structure (i.e. degree of collusion) affects bureaucrats' decisions about corruption. Instead of reinventing these, I take them as a starting point and focus instead on improving our understanding of how bureaucrats respond within the context of an agency. The key contribution of this work to mine is the importance of maintaining above-market purchase prices to rent-seekers on both side of the procurement transaction. This informs the measurement discussion in the following chapter. The second reason is simplicity. Thanks to this prior work on how firm behavior affects bureaucrats' corruption decision, I am able to develop a compelling explanation of my puzzle at the level of the agency, and characterize the procurement market it faces, without needing to model that market in detail.[34]

In sum, bureaucrats facing demands for corruption must weigh the costs and benefits of deviating from the prescribed procedure. The more she can coordinate with actors demanding corruption, the better off is the bureaucrat when the resulting rent is divided. Finally, unlike common crime, where the probability of detection is

---

[33]For example, see Compte, Lambert-Mogiliansky and Verdier (2005).

[34]In contrast to the difficulties transferring the efficiency wage hypothesis to the agency level, procurement market structure is more defensibly exogenous. Each agency has a mission/purpose that is quite fixed and implies it purchase a particular 'basket' of goods. The nature of this market varies by agency along these relatively clear mission-driven lines, so firm-level variation in behavior within this 'basket' is less important. I measure the competitiveness of these markets in Chapter 4.

plausibly exogenous in a single case, corruption is both overseen and enabled by the bureaucrat up the chain from where the transaction takes place. For this reason I discuss next how the incentives inside agencies explain variation in corruption.

2.3.3   How do agencies manage corruption internally?

Beyond the individual-centric model of the corruption decision discussed above, bureaucrats operate within a particular kind of organization, the agency.[35] This makes corruption different to crimes committed by private citizens. Were corruption no different, my agency-level puzzle could be answered simply by making one bureaucrat a 'representative agent'; agency-level corruption would be determined by the same things affecting this individual's decision. The literature I discuss below makes clear that agency features, in particular relationships between bureaucrats in a hierarchy, modify the implications of individual-level explanations.

To my knowledge only one study has asked a research question close to the puzzle motivating this dissertation. Recanatini, Prati and Tabellini (2005), in *Why Are Some Public Agencies less Corrupt than Others? Lessons for Institutional Reform from Survey Data*, studied agency-level corruption in eight (unnamed) developing states in Latin America and Africa. The data are from the *World Bank Institute Governance and Anti-corruption Diagnostic Surveys*, which ask citizens, firms, and public officials detailed questions about governance vulnerabilities. They posit that three agency-level features reduce corruption: the extent of internal and external auditing of decisions; adherence to meritocratic recruitment and promotion; openness about agency decision-making and performance. They measure each feature with several questions asked of bureaucrats and use the respective first principal components as independent variables. In the case of public procurement they find auditing and openness are associated with reduced corruption.

---

[35]I defined this above as "any collection of bureaucrats who report to the same bureaucrat above them in a hierarchy."

As with the Hunt (2005) study discussed above, this is an important contribution to the research on my puzzle, and much of the literature on 'what works' in designing institutions to prevent corruption is consistent with its findings. Ferraz and Finan (2008) and Olken (2007) show that audits are effective, as long as their schedule is not announced in advance and they provide high-quality, accurate information about specific individuals' responsibility; only this leads politicians and voters to 'kick the bums out'. Likewise, meritocratic recruitment is widely acknowledged as an important part of any low-corruption equilibrium at the national level (Charron, Dahlström and Lapuente, 2016; Sundell, 2014). And none of these mechanisms can function well without a sufficient level of information about agency workings being available to the government and external parties. These findings can be summarized in a hypothesis:

L3a: Agency corruption is lower the more it adopts random auditing, meritocratic staffing policy, and openness about its actions

Given these features represent how an agency chooses, *as much as it can*, to manage corruption, this work provides answers to my puzzle while raising further questions. Just as Hunt's study did not explain how differences in agency client pools arise, Recanatini, Prati and Tabellini acknowledge that neither their theory or empirics consider why agency features differ to begin with.[36] This shortcoming has plagued corruption studies for a long time: evidence of how policies can change bureaucrats' behavior is easier to find than evidence of why policies are adopted and enforced to varying degrees. Resolving such an issue is beyond the scope of this work, but I take the lessons of this previous work into consideration when building my theory and research design. I am able to generate plausible empirical implications explaining clean agencies in corrupt states with a model that includes, at the very

---

[36]While the problem of unobserved heterogeneity across agencies can be managed empirically, resolving my puzzle requires an explicit theory to explain as much of this heterogeneity as possible.

least, a simple treatment of auditing and recruitment. The bureaucrats in my model operate under the shadow of a manager who tolerates only so much corruption, and may launch an audit if he suspects this limit has been surpassed. Expertise, which is the main benefit of meritocratic recruitment, also plays a role in my theory: a more expert junior bureaucrat is more able to generate the rents that are necessary for corruption. This expertise would include a greater ability to dodge audits, avoid meritocratic staff reviews, and hide information from the outside world, consistent with *L3a* above.

In my empirical analysis, I choose data generated by an activity that all agencies must do (procurement) from a country (Russia) where all agencies must play by the same rules. This is less susceptible to the selection problems affecting the Hunt and Recanatini studies, and I purposely operationalize my key independent variable, information asymmetry, as a measure of agency choices over what to procure.[37] In this sense, my study characterizes an agency by its choices over what to buy and how. When operationalizing my theory I explain how underlying differences in agencies are generated, by their different product needs, and how this variation in needs is measured by their overall purchasing choices.

My focus on the relationship between bureaucrats within an agency to explain my puzzle is informed by an extensive literature that adapts the model of the firm developed in transaction cost economics and industrial organization.[38] In political science this literature has contributed significantly to our understanding of the core tension of any bureaucracy: the *principal-agent problem.* Put simply, when an agent is empowered by a principal to carry out some task, any difference between them

---

[37]Furthermore, there is little evidence that internal or external auditing are robust in any Russian agencies, which is a natural consequence of weak political accountability for government agencies' actions. Given the significant incentives for auditors to earn from corruption themselves, data of the kind used by Recanatini is unlikely to be reliable even if it existed. Meritocratic recruitment and promotion is also uncommon in Russia: on Russian-language search engines the most common auto-complete option for the question "how do I enter state service?" is "without connections".

[38]Formative work in their respective fields is Williamson (1981) and Tirole (1994).

in preferences and information imposes two types of *agency costs* on the principal. If their preferences diverge, an agent will pursue a different course than would the principal, changing the outcome; in the course of carrying out the task an agent gains information not available to the principal, which can be strategically withheld.[39]

Despite these problems, principal-agent relationships are useful for individuals seeking an outcome beyond the limits of their own productivity, for two reasons: by delegating to one or more agents, principals multiply the available labor; by hiring agents with different preferences and observing how their actions diverge, principals gain valuable information that helps them make better decisions.[40] When benefits exceed agency costs, delegation makes the principal better off. As it also results in compensation for agents, delegation 'grows the pie' of total benefits generated by the agency's business. Over time, and under certain conditions, agencies develop highly specialized expertise that makes production of these benefits more efficient. When agency expertise becomes high enough, politicians may even be willing to tolerate significant preference divergence and information asymmetry in their agencies, if the political gains from more informed policy outweigh the loss of partisan control.[41]

When applied to corruption, principal-agent models have typically assumed the principal (whether supervising bureaucrat or politician) is maximizing social welfare. This makes the model simpler, reducing the principal's problem to two matters of institutional design: how to give the agent incentives to avoid corruption (e.g. efficiency wages), and how to monitor the agent's compliance with these incentives (e.g. random audits). This simplification creates problems, however, when confronting the puzzle of clean agencies in corrupt states. First, it pushes the explanation for

---

[39]This problem is so fundamental that even if the principal could clone herself as an agent, the mere fact of delegation can lead to an outcome that differs from the principal's preferences.

[40]These motivations are the core of the framework developed by Huber and Shipan (2002) to explain variation in delegation of policy-making authority across democracies.

[41]This story of agency development is exemplified in the qualitative literature by Carpenter (2000), and in the formal study of bureaucracies by Gailmard and Patty (2007). Common to both stories are increasing grants of discretion by politicians to agencies as expertise develops.

variation back on the agency's features, because the welfare-maximizing principal's preferences do not vary by agency.[42] Furthermore, where national-level corruption is extensive, as in most states, it is implausible to assume that competition disciplines politicians to the extent they can be approximated as welfare maximizers; a better starting points is that they maximize rents from office.[43] The same is true for any bureaucrat supervising others, as stated succinctly by Olken and Pande (2012, p.498): "the very individuals tasked with monitoring and enforcing punishments may themselves be corruptible, so increasing monitoring may simply increase transfers from low-level officials to auditors".

This potential corruptibility of those responsible for monitoring corruption is a crucial building block of my theory. In a corrupt state, clean agencies must be doing *something* unusual to deviate from the national average, but existing theory cannot explain why some agencies would make drastically better choices over the features (e.g. auditing) that cause such deviation. Any attempt to explain clean agencies based on the argument 'they chose less corruption-prone features' will meet the same difficulties I noted in the Hunt and Recanatini studies.[44] Instead of trying to explain choices over such features, I take the institutional design of agencies as given, and minimize the threats to inference caused by this trade-off in three ways: my model is agnostic to the kind of bureaucratic decision being made; it still incorporates simple exogenous variation in agency features related to corruption, such as auditing;[45] and my data are generated by an activity that all agencies must engage in to survive, procurement of goods and services. In Russia's procurement system in the period

---

[42]This does not invalidate the effect of audits, meritocracy and openness on corruption outlined above; instead it reinforces the need to go beyond the Hunt and Recanatini studies to explain why agencies adopt these effective tools to varying degrees.

[43]This intuition can be motivated by the 'selectorate' theory of Bueno De Mesquita et al. (2005).

[44]This is a case of the difficult trade-off between explaining institutional equilibrium and equilibrium institutions identified by Shepsle (1989).

[45]Unfortunately, I cannot reliably observe variation in external auditing of Russian agencies, as the main audit institution examines agencies on a schedule. Furthermore, in equilibrium it is the threat of audits the disciplines bureaucrats; if an audit finds corruption it has already failed in one sense, though may act as a deterrent.

under study, bureaucrats had significant discretion individually, but within a set of common and highly constraining set of institutional design choices.

The model I develop below argues that the 'unusual' quality of clean agencies in corrupt states is a *failure to maintain robust collusion* necessary to extract rents through corruption. Though my theory is not dynamic,[46] it treats expertise in corruption like any other kind of expertise an agency might develop. In corrupt states, where bureaucrats' wages are universally poor, expertise in creating rents and extracting them through corruption is a significant source of benefits. This means I can use the same tools used to explain variation in corruption within principal-agent relationships (mainly in economics), and the development of agent expertise (mostly in political science), to approach my puzzle. I reserve further discussion of specific principal-agent tools from the literature to the section below where I develop my own model. For now, I generalize the following hypotheses:

L3b: Agency expertise increases in the specialization of agents relative to principals

L3c: Agency corruption is increasing in the expertise of its bureaucrats

In my theory, the interaction of these two mechanisms can lead to a breakdown of collusion between bureaucrats, producing a clean agency even when opportunities for corruption are extensive. While I treat expertise in corruption as existing work treats policy-making expertise, it differs in one crucial respect: the spoils of corruption must be kept secret to be retained, whereas the whole point of policy-making expertise is to change *public* policy. This changes how the senior bureaucrat assesses the (more informed) junior's report of what was earned from corruption; it is not simply the opening move in a standard bargaining game. If the junior bureaucrat under-reports the true extent of illicit benefits, in an attempt to deceive the senior, the underlying

---

[46]Also, my data do not span sufficient time to test a dynamic theory.

34

risk to the latter is unchanged despite her lower expected reward. She will be held responsible for all corruption, should the manager learn of it. I now turn to the literature that explains the manager's decision over punishment.

2.3.4 How do politicians react to corruption in their agencies?

Another relationship that might explain why corruption varies across agencies is between the agency and the political leaders who sit above it. Whether they create or inherit their set of agencies, politicians in every state are responsible to some degree for the actions of bureaucrats. I discuss this relationship only very briefly, as its effect on cross-agency variation depends on the extent of political accountability for agency decisions. In the corrupt states of my puzzle, this accountability is low by definition. Nevertheless, some discussion of this literature is necessary to motivate the objectives of the minimally-specified 'manager' who is an important actor in my model.[47]

All politicians must ensure their survival in office, and subject to achieving this aim, maximize their rents from holding office. Politicians care about survival first because it is a precondition of future control over rents, control that challengers covet. For democratic politicians these challengers are both external, in the opposition, as well as internal to their own party. The less democratic a leader, the lower the threat of an opposition challenger rising to prominence, but the higher the risk of replacement from within the elite (Svolik, 2008). In short, all politicians must worry about being bumped off in more or less damaging ways, and often by their 'friends'. Politicians fend off these challengers, regardless of their source, by maintaining a sufficiently large support coalition.[48] This requires governing politicians share the rents garnered

---

[47]The manager is thought of as a politician here to separate his incentives from bureaucrats', who implement policy. I use the term 'manager' to emphasize his role in the game as automaton arbiter of the agency's corruption threshold, rather than a strategic actor. This has the additional benefit of allowing the model to work at any level of a nested set of agencies; if senior and junior are respectively the second-lowest and lowest bureaucrats in the hierarchy, the senior's immediate superior is unlikely to be a politician anyway. Nevertheless, any superior is concerned about corruption below to at least some degree, as they must report up the chain themselves.

[48]What qualifies as 'sufficiently large' varies by regime type.

from the economy with some of the people who matter in choosing the government, commonly termed the selectorate (Bueno De Mesquita et al., 2005). In democracies, the selectorate is usually considered to be the electorate, and in non-democracies it is made up of members of the political elite. The problem politicians face is how to give away the minimum amount of rents while still retaining office.[49]

Political leaders use the bureaucracy to reduce the severity of this problem. They garner and distribute coalition-maintaining rents using agencies for the two reasons stated above: delegating power to intervene in the economy to many agents increases the available labor; separating this labor in to specialized agencies increases their efficiency. For states in a low-corruption equilibrium, these rents usually take the form of policy changes implemented by expert agents, and corruption is the exception. The reverse is the case in corrupt states, where support coalitions are nurtured through discretionary corruption rather than lasting policy concessions.

This brief discussion of political survival, extended to include the bureaucracy as a tool for creating and distributing rents, illuminates an important constraint on corruption even in corrupt states: every agency must do *something* that resembles its mandate, otherwise it has no reason to exist. In a two-bureaucrat agency with a given budget, the agents' ideal spending on actual policy implementation is zero because this leaves the entire budget for extraction. However, agencies in high-corruption contexts always deliver *some* observable output, even if significant energy is spent trying to extract as much as possible. This is because an agency delivering too little of its mandate becomes an *inefficient* way to capture rents, and may even strangle the sector of the economy it regulates (Block, 1977). If all the two bureaucrats do is argue over how to divide the entire budget between their pockets, even their weakly-incentivized manager will prefer to disband the agency and take the money directly.[50]

---

[49]This problem is strategic, rather than decision-theoretic, because the behavior of supporters and potential challengers depends on the leader's actions and vice-versa.

[50]This leads to an disturbing implication for bureaucratic politics in non-democratic settings, well-supported by case studies of states like Russia (Dawisha, 2014). Those running the state create and

This generates a hypothesis that I cannot test, but assume on solid grounds to be true:

L4a: Politicians' tolerance of agency corruption is never absolute[51]

Due to weaker political accountability for bureaucratic decisions in corrupt states, this literature does not provide an explanation of variation in agency-level corruption. Clean agencies might be the outcome of 'abnormally strong' political accountability, but it is not clear from existing theory what could generate that variation. Robust competition from credible oppositions is the source of this accountability (Grzymala-Busse, 2007), but in this class of explanations the state apparatus is more efficient *overall* when its political masters face a credible threat of replacement. The variation being explained in these studies is choice over broader political institutions, not the design or behavior of individual agencies.[52] What this literature does provide, however, is further support for the feasibility constraint for agencies specified in *L4a*: they must do 'just enough' policy implementation to justify their existence, because managers cannot tolerate 100% corruption.

2.3.5 How do citizens and firms react to corruption that benefits others?

I began this review by asking why citizens and firms demand corruption, and found this varies according to the prevalence of rents administered or created by an

---

expand agencies not only to 'do more' in terms of policy implementation (whether good or bad); they expand the state so that it can *spend more*, because spending agency budgets is the best way to extract rents. The best place to pursue corruption is right alongside 'justifiable spending' on undeniably necessary goals. Nobody can argue Russia does not need roads, and building more roads grows the 'fat' available in the budget for extraction by bureaucrats. To paraphrase the playwright Plautus: "you must spend money, if you wish to [t]ake money" (Plautus, 1992).

[51]I summarize the implications derived from the literature review in Table 2.2.

[52]An earlier version of this project attempted to explain agency-level variation in corruption as a function of variation in accountability of a political leader for agency decisions. This relied on a notion of political risk: in agencies where the risk to political survival of a corruption scandal is high, less corruption is tolerated. Tolerance was explained by the costs of scandal, for example an IT failure due to corruption in a nuclear regulatory agency is more risky to the leader than the exact same failure in the Ministry for Tourism. I abandoned this direction due to the convolution of the theory and the fact that 'political risk' is even more difficult to observe than corruption.

No.	Hypothesis	Testing
L1a	Rents are larger the more government intervenes in the economy	Not tested, foundational to literature
L1b	Demand for corruption increases in the size of rents	Not tested, common across theories
L1c	Agencies intervening more in the economy generate larger rents	Not tested, demonstrated by Hunt (2005)
L1d	Demand for agency corruption increases in the rents it generates	Not tested, demonstrated by Hunt (2005) & Recanatini, Prati and Tabellini (2005)
L1e	Bureaucrats' private returns on corruption increase monotonically with the opportunities for corruption generated by their agency's administration of rents	Modified by my theory, **tested** at agency level in Chapter 3
L2a	Paying a bureaucrat a higher income decreases the probability she engages in corruption for her personal benefit	Not tested, no bureaucrat-level data
L2b	Paying an agency's bureaucrats a higher income decreases the probability of corruption for personal benefit in that agency	**Limited test** in Chapter 4, for subset of agencies
L2c	Procurement market features that aid collusion between bidders increase the returns to corruption for bureaucrats by decreasing competition between bidders	**Tested** at purchase and agency level in Chapter 4
L2d	Agencies facing procurement markets with lower levels of competition are at greater risk of corruption	**Tested** at agency level in Chapter 4
L3a	Agency corruption is lower the more it adopts random auditing, meritocratic staffing policy, and openness about its actions	Not tested; all Russian agencies must play by same rules in single procurement system
L3b	Agency expertise increases in the specialization of agents relative to principals	Not tested, common in literature
L3c	Agency corruption is increasing in the expertise of its bureaucrats	Made conditional by my theory, **tested** at agency level in Chapter 3, and at purchase and regional levels in Chapter 4
L4a	Politicians' tolerance of agency corruption is never absolute	Not tested, but a safe assumption based on literature

Table 2.2: Summary of hypotheses generated by literature review

agency. Along with demand for corruption, these very same factors can also create demands *against* corruption from the same actors, if they lose from others' success in gaining access to rents.[53]

As my theory relies on only the weakest form of agency accountability to politicians (i.e. 'though shalt not extract everything'), it follows that the literature on how voters and firms pressure politicians over corruption is not relevant to my puzzle. Mechanisms of political accountability at the core of major political economy models (Ferejohn, 1986; Myerson, 1993; Persson, Tabellini and Brocas, 2000), in which competition dissipates politcians' (and in turn bureaucrats') rents, are riddled with weaknesses in corrupt states. Moreover, even if these mechanisms are operative, the "dominant view in the literature, that citizens want to vote out corrupt politicians and that political institutions that empower citizens to do so will lead to lower corruption" (Ichino, Faller and Glynn, 2013), is severely undercut by the evidence. Voters respond negatively to corruption under very narrow conditions: when they perceive it meets bureaucrats' greed rather than voters' need (Bauhr, 2016); when the state of the economy is poor, not when conditions are good (Klašnja and Tucker, 2013); when corrupt politicians are not 'their guy' (Winters and Weitz-Shapiro, 2013). These results rely heavily on political competition from credible oppositions, which is generally absent in corrupt states, and certainly in Russia during 2011-2015.[54]

2.4 Russia: A motivating case

In 2012, President Putin's United Russia party was hit by a series of corruption scandals. Enterprising members of the newly-emboldened opposition dug around in

---

[53]It is safe to assume these actors do not complain about corruption that benefits them, nor do they care particularly about dead-weight losses due to corruption.

[54]Despite its lack of implications for agency-level corruption, when testing my explanation against the main implications of national-level explanations, using Russia's regions as 'stand-ins' for states, the dominance of the governing party is an important variable. As with corruption, the low level of political competition at the national level obscures significant variation across regions.

the property registers of foreign states, and came up with damning evidence of unde-clared property belonging to elected officials, including the chair of the parliamentary ethics committee. Most of these officials, after initially protesting their innocence, resigned their mandate as MPs and moved out of politics, in to the bureaucracy. There, they were able to continue extracting rents but without the annoyance of even minimal oversight of state activity. Corruption in Russia, at the highest levels and for mind-bending amounts, is no obstacle to career progress for the well-connected. In 2013, then Minister of Defence Anatoliy Serdyukov was accused of participating in a scheme to illegally sell the land underneath a military base to a holding company partially under his control, and lease it back to his agency. He left his post, but faced no charges, even as his co-conspirators were convicted.[55]

Though more Russian citizens can survive today without paying bribes for critical services, the 'grand corruption' at the top, so crucial to maintaining the system of rent-seeking and risk-sharing that supports the regime, continues unabated. Today's level of extraction is an open secret with deep roots. Karen Dawisha (2014) outlines the following scheme involving associates of Vladimir Putin in 1990s St Petersburg, centering on his one-time neighbor Vladimir Smirnov:

> "The Nevskiy building stood empty for fifteen years, and yet [...] between October 1997 and July 2000 documents reproduced by *Novaya gazeta* show that 63.83 million Deutschmarks ($35 million) was transferred to [the construction company] Znamenskaya in twenty payments, over Smirnov's signature. This was self-service in the extreme in that he was signing for SPAG [a real estate holding company incorporated in Germany and co-funded by St Petersburg's City Hall] as a member of the board, authorized by Putin to vote the city's shares, and giving money to Znamenskaya,

[55]For details see: *New probe launched into Defense Ministry fraud scandal* (`http://rapsinews.com/news/20121206/265672434.html`) and *Russians Look Askance at Anticorruption Drive Even as New Scandals Arise* (`http://www.nytimes.com/2012/11/18/world/europe/russia-looks-askance-at-corruption-drive.html`)

which he headed, for a project that was not being built."

The energy and creativity required to formulate such an arrangement is far from the common image of the docile paper-pusher. Nor is it Kafka-esque, constructed simply to confuse the citizen in to submission to a distant and undiscoverable will; on the contrary the beneficiaries and methods are clear.[56]

In contrast to the 'wild west' schemes described above, popular in the 1990s, and brazenly open in their exchange of funds for *inaction*, in Russia today much of this competition for rents has moved to a more controlled arena: public procurement. The setting for my empirical tests, it accounts for approximately 30% of total government spending in any given year, making it a lucrative source of rents for those with the resources to form a company, bid, and win, whether honestly or in collusion with bureaucrats running the process. Despite procurement auctions following a prescribed format (albeit with some variation in procedure, which I use in the next chapter), there has traditionally been no oversight *in advance* of what agencies buy, and whether this is justified by some standard external to the agency.[57]

Contributing to the 'open secret' argument, despite an environment of endemic corruption, relatively independent and outspoken anti-corruption institutions exist in Russia. The Audit Chamber, the chief agency overseeing the bureaucracy and reporting to the parliament, has typically been headed by individuals who might even

---

[56]Corruption schemes this intricate vary significantly due to local factors. As a result, humor from contexts where bad governance is the norm translates poorly, as the following Russian jokes demonstrate: 'In the city-wide corruption championship, the winner was someone's nephew.'; 'How do you fight corruption in a country where even the envelopes are the size of banknotes?' This particular scheme relied on particular legal and political conditions during the late Soviet and early post-Soviet periods, during which some city governments held a monopoly on licenses for imports and experts. Later, in Moscow, Putin closed such loopholes as he centralized power.

[57]This is perhaps why the same sidewalk in Moscow can be torn up year after year: first to replace paving stones with asphalt, then to replace asphalt with concrete, then to replace concrete with different paving stones. In my experience walking Moscow's streets in 2011-13, the job is poorly done each time, and one winter is more than enough to create the need for further spending on another renovation. Most importantly for those running the process, the budget is spent every year, though extracting a cut is not automatic: work must be done on the sidewalk for the money to flow. This is why I give the manager in my simple model of an agency a tolerance for corruption below 100%; the money must be spent and *some* observable action taken for the scheme to work.

qualify for such a job in a low-corruption state. Sergei Stepyashin, who served from 2000-2013, and his predecessors were widely acknowledged to be impartial judges of the use of government finances, and Audit Chamber heads have not shied away from drawing attention to the overall size of Russia's corruption problem, nor to specific cases. Even Stepyashin's successor, Tatyana Golikova, who was previously the Minister of Health and Social Development and is married to the current Minister of Industry and Trade, has not slowed the pace of declarations that agencies' approach to corruption needs to improve. A 2012 Audit Chamber report estimated total losses due to corruption in procurement that year at US$32 billion, and received widespread media coverage in the business and popular press. Such an estimate, which is likely conservative, represents 17% of procurement spending and 7% of all government spending. It is not an ignorable sum, and points to a serious problem in Russian governance, yet it is highly unlikely that the Kremlin did not see this report before it was published, or even encourage its production.

Specifically in the realm of procurement, investigations of agencies are frequent. In 2013 the Audit Chamber reported that 70% of the 458 state contracts valued over US$30 million violated regulations designed to ensure competition.[58] A 2015 investigation in to contracts for the new Moscow Central Ring Road found many contracts were awarded to bidders who submitted a higher price, but succeeded nevertheless, based on dubious additional criteria such as 'use of nanotechnology'.[59] Several times a year, a highly-placed official in one of the many anti-corruption agencies within law enforcement is arrested after being entrapped taking a bribe to ignore corruption. Competition to occupy peak anti-corruption positions is intense, and such arrests often signal a struggle between factions of the security services to maintain oversight

---

[58]Source: RAPSI News, *Most major Russian government contracts breach rules - regulator*, 15 August 2013, http://www.rapsinews.com/news/20130815/268562284.html

[59]Source: The Moscow Times, *Russian Road Agency Investigated for Holding 'Suspicious' Tenders*, 6 July 2015, https://themoscowtimes.com/articles/russian-road-agency-investigated-for-holding-suspicious-tenders-47930

over the others.[60]

Taken together, these anecdotes begin to offer an explanation for surprisingly outspoken anti-corruption authorities in states such as Russia: the regime *needs* them, not to reduce corruption, nor to bolster incredible claims to legitimacy, but to create common knowledge between political leadership and bureaucrats of the appropriate level of 'takings' without writing it down.[61] Delegating this grand whistle-blowing power to a relatively independent voice increases the perceived accuracy of the signal. The ultimate goal of such institutions is to place some limits on corruption, while assuring bureaucrats that anti-corruption action is not totally arbitrary. Relatively impartial anti-corruption institutions encourage bureaucrats to believe that if they 'keep their noses clean' by predating within certain prescribed bounds, they will have little to fear. This further motivates why, in my model, the manager tolerates at least some corruption by bureaucrats, but never 100%, and punishes bureaucrats with non-zero probability.

2.5 A model of agency-level corruption

In this section I tie key hypotheses generated from the literature review to my own insights, using a simple formal model to structure the discussion and generate empirical implications that encompass my own and existing explanations. These inform the research design and conservative tests of my theory in Chapter 3, and the more ambitious and detailed tests in Chapter 4.

---

[60]In 2014, the head of the Interior Ministry's anti-corruption division was arrested in a sting operation relating to bribes paid for stopping investigations in to procurement fraud: `http://www.rbc.ru/politics/01/09/2014/946145.shtml`. In 2015, the head auditor of procurement at the Ministry of Defense was detained on suspicion of heading an extensive network of procurement corruption: `http://newsru.com/russia/22dec2015/zakazotkat.html`

[61]Codifying this level, however informally, would potentially reduce the discretion enjoyed by politicians in such systems to use anti-corruption as a tool to manage the elite.

2.5.1 What the model demonstrates

For an agency to be highly corrupt, there must be *robust collusion* between its bureaucrats. 'Collusion' here implies cooperation to extract state resources above the level preferred by some other actor; 'robust' emphasizes the fragility of that cooperation in the face of a temptation to defect and extract *from each other*. The preconditions for this collusion vary across agencies based on the degree of information asymmetry between levels of the agency hierarchy; sometimes this is high enough to produce low-corruption agencies even when average corruption is high in a given state. In the model these agencies are *not* generated by variation in the propensity of individuals for corruption, which would have clean agencies created by well-meaning bureaucrats. Instead, corruption fails where bureaucrats cannot trust each other enough to cooperate and maximize the joint returns to corruption, because one anticipates losing out when these returns are divided between them.

2.5.2 Who are the actors?

I follow most theories of agency corruption and focus on the actions taken by two actors. A subordinate can deviate from some prescribed procedure in order to create an opportunity for self-enrichment, by receiving a share of the rents generated by the deviation. A superior decides whether to impose sanctions for that deviation, if she detects it. In the canonical example in the literature, the subordinate grants a permit to an unqualified applicant in exchange for a bribe, while the superior looks the other way (and potentially asks for a cut).[62] I call these actors the *senior* and *junior* bureaucrat and their interaction defines the agency at the core of my model.

This interaction takes place in a context where the agency's budget, bureaucrats' wages, and the returns to whistle-blowing (also conceptualized as the political lead-

---

[62]Deviation from a prescribed procedure does not imply that discretion is never warranted, simply that the bounds of discretion be well-specified and not stretched beyond breaking to create rents. I discuss my preferred definitions of these terms at the beginning of the next chapter.

ership's tolerance for corruption) are already set. These factors help shape the interaction by shifting the incentives of my two strategic actors, and for the sake of exposition I bundle them in to a non-strategic actor: the *manager*.

This simple structure allows the logic of the model to apply to any level of a real agency. At the very top, the senior and junior roles would be assumed by the responsible minister/secretary and the head civil servant respectively, while the manager is a prime minister or president. At the very bottom of an agency hierarchy the lowly public-facing desk clerk is the junior and their immediate supervisor the senior, with the manager one level further above. This trade-off in favor of abstraction allows for more freedom in research design, as I elaborate below.

2.5.3 What are they trying to achieve?

Both bureaucrats want to maximize their stream of benefits from holding office. There are two sources of benefit: a fixed wage set by the manager[63] and a variable share of the rents extracted through corruption, which is determined by their interaction. The structure of their objective functions is therefore identical, what distinguishes them is that the senior is paid more and moves first. The tension *within* each actor is that pursuit of the illicit stream (rents extracted) increases the likelihood they lose their licit stream (their wage), although this problem is worse for the senior as the responsible agent. The tension *between* the actors follows from having to divide the rents extracted between them: one's gain is the other's loss.

The manager's ideal level of corruption in the agency is always lower than the bureaucrats',[64] though it is never zero, due to decreasing marginal returns of efforts to lower corruption.[65] This means she accepts there will always be a part of the

---

[63]By assumption the senior earns more than the junior.

[64]Following Block (1977) I assume the manager is "dependent on the maintenance of some reasonable level of economic activity" in the sector regulated by the agency, which in turn implies the bureaucrats spend "some reasonable level" of the budget on their core job. The impetus is that association with a sudden drop in activity increases the likelihood the manager is replaced.

[65]Monitoring and enforcement is always more costly than some level of corruption.

agency's budget diverted to the pockets of the bureaucrats. A more sophisticated model could endogenize the manager's punishment strategy; here I simply assume she incentivizes whistle-blowing to some degree, as a way to learn about the most egregious cases.[66]

2.5.4 What actions are available to each?

Outside the model, the manager sets a budget for the agency, the wages, and an incentive for whistle-blowing.[67] These are taken as given within the model itself, an extensive-form game between senior and junior bureaucrat.

The agency's budget (π) is a proxy for an alternative explanation of variation in corruption across agencies from the comparative politics literature: larger agencies have more *capacity* to control corruption. Even though total funds siphoned off in large agencies might be higher, this explanation predicts the proportion siphoned off to be lower than in a small agency. Introducing this as a parameter allows me to examine its effects on the incentives for robust corruption. As a duo the bureaucrats would prefer to extract the entire budget through corruption and receive respective shares of the budget (b_s, b_j) equaling 1, but if they fail the remainder of the budget b_m is spent on the manager's priorities.

The senior and junior bureaucrats' wages $(w_s > w_j)$ capture another important argument for variation in corruption: low official compensation encourages bureaucrats to seek alternative sources of income available due to their position. Conversely, the efficient wage hypothesis contends high compensation creates a future income stream that encourages bureaucrats to restrain themselves. Wages often vary considerably across agencies within the same state, and so are a plausible alternative explanation.

---

[66]In states with low levels of corruption, the returns to whistle-blowing are typically very high, in order to overcome bureaucrats' reluctance to risk job tenure.

[67]This incentive can subsume many national-level explanations of corruption that do not allow variation by agency, for example the limits to rent-extraction in high-corruption contexts motivated by the roving vs stationary bandit logic of Olson (1993). The manager knows that if the bureaucrats extract the entire budget for themselves this kills the 'golden goose'.

The returns to whistle-blowing (k) are derived from the concept of *kompromat*, a Russian portmanteau of 'compromising material'. If revealed, *kompromat* damages the target's reputation to such an extent he must leave office.[68] The senior bureaucrat, by virtue of hierarchy, can fire the junior bureaucrat at will. Under some equilibria of the model, the latter has no leverage and therefore no recourse. During their work, however, the junior bureaucrat has observed the senior's choices about corruption. If the manager places sufficiently high value on obtaining this information, the junior bureaucrat can benefit from revealing it to punish the senior. This parameter does not vary by agency, and can also be motivated as the return to a bureaucrat for ensuring adherence to the manager's tolerance for corruption across all agencies. The parameters π, w_s, w_j, k are taken as given in the model illustrated in Figure 2.1. The two bureaucrats must decide how to maximize their respective streams of benefits from wages and corruption. The central tension is that their interests coincide with respect to growing the pie, but conflict when it comes to dividing it.

The senior bureaucrat first chooses whether to grant the junior bureaucrat discretion to engage in the extraction of rents through corruption.[69] Then the junior bureaucrat decides whether to extract or not. If not, then regardless of whether extraction was allowed or forbidden, the game ends with each bureaucrat obtaining their respective wage ($w_s; w_j$).[70]

In the event the junior went against orders, and extracted corruption rents when

---

[68]Understanding *kompromat* hinges on distinguishing 'open secrets' from 'exposed secrets'. In Russia, for example, it is no secret that senior officials live well beyond the means implied by their salaries. The plausible justifications for this are limited to corruption and prior wealth, and it is an open secret that corruption is usually the reason. Exposure of specific instances of corruption, tied to particular individuals, has a very different effect because they create a 'fall guy'. Open secrets only deaden citizen expectations of good behavior, whereas exposed secrets become a focal point for discontent and a trigger for swift anti-corruption.

[69]Although this project focuses on high-corruption contexts, it is important the model is general enough to also generate an equilibrium with low corruption.

[70]These two outcomes are observationally equivalent, and occur only when the returns to whistle-blowing (k) are very high. The equivalence is less problematic in substantive terms than it is within the model. As corruption can never be stamped out entirely, a more comprehensively modeled manager would accept the risk that sometimes her senior bureaucrats allow extraction (and that she is blind to this), as long as extraction does not occur in equilibrium.
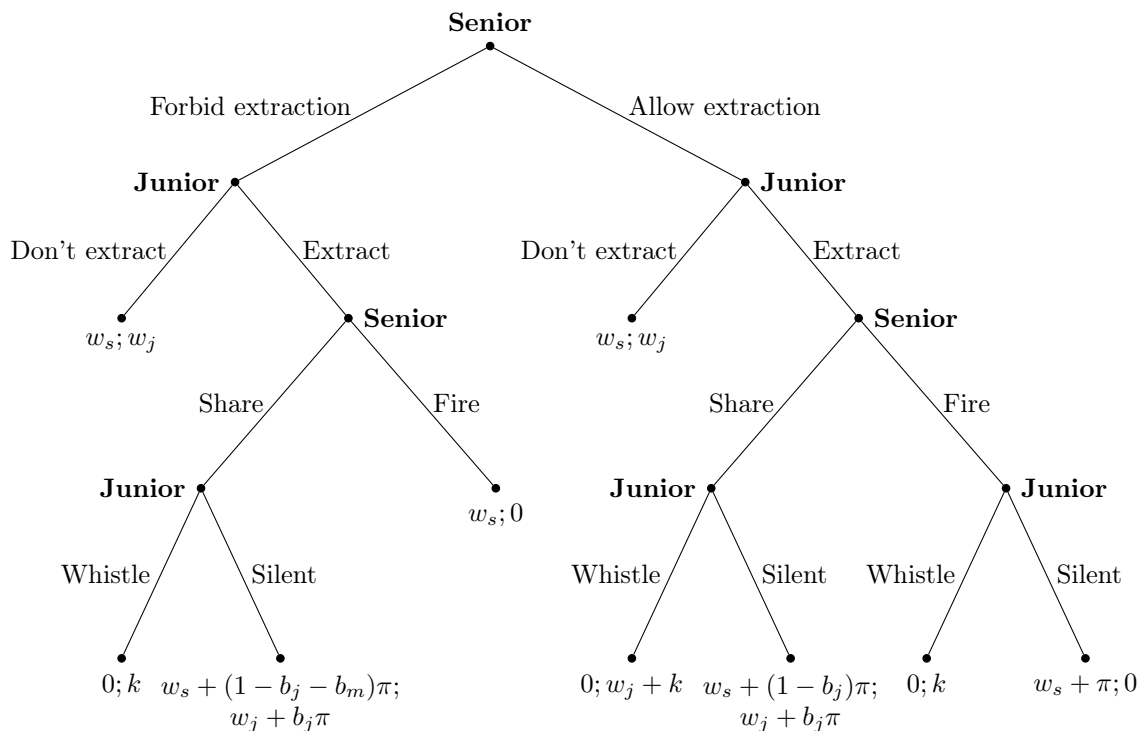
Figure 2.1: Extensive-form game between senior and junior bureaucrat

it was forbidden, the senior faces a choice: fire the junior and forgo the rent, or give
the junior a share b_j of π and retain $(1 - b_j - b_m)\pi$.[71] By going against his own
prior directive forbidding extraction, the senior creates *kompromat*, which the junior
can use to blow the whistle should his final share of π be too low. Furthermore,
because they did not both actively 'work for' extraction I assume the total share
extracted $(b_s + b_j)$ is less than 1, making $b_m > 0$ (so the agency provides some policy
implementation).

In the other major branch, the senior allows extraction and the junior takes up
the offer. They are now working together and can extract the entire agency budget
to divide between them, so $b_m = 0$.[72] The senior faces a similar decision: share the

---

[71]The agency retains part of its unextracted budget, b_m, because the bureaucrats did not coordi-
nate on (Allow; Extract) as the first two moves of the game. This penalizes the senior bureaucrat
for not authorizing corruption up front that could be of later benefit.

[72]In a real agency, absolute theft would be exceedingly difficult as it collapses the justification for
the agency to exist. The reduction undertaken here is less extreme if we think of π not as the entire
agency budget but rather its discretionary component. The manager would still prefer it be spent

corruption rents extracted, or fire the junior and take the whole lot. In this case, firing is not 'clean' because the junior bureaucrat was obeying (rather than disobeying) orders, so for any $k > 0$ the whistle is blown and the senior loses everything.[73] This forces the senior to share just enough so the junior is better off accepting the proposed split than taking their wage plus the returns to whistle-blowing.[74] This is the 'robust collusion' equilibrium outcome that describes the highest-corruption agencies.

2.5.5 What outcomes are possible?

Suppose the two bureaucrats reach the robust collusion equilibrium with the following actions along the equilibrium path: Allow extraction; Extract; Share; Silent. This breaks down if the benefit to the junior bureaucrat of whistle-blowing (their wage w_j plus the return to whistle-blowing k) is greater than the benefit of staying silent (their wage w_j plus the share of π offered by the senior). This occurs when $b_j < \frac{k}{\pi}$ so the minimum share necessary for robust collusion is, quite intuitively, increasing in the returns to whistle-blowing and decreasing in the agency's budget.

By the logic of this simple version of the model, agency-level corruption (the proportion of π extracted by bureaucrats) varies bluntly by equilibrium. When robust collusion is maintained as above, the whole budget is extracted. When this fails, either nothing is extracted, or when the bureaucrats do not both choose extraction, corruption is less effective, yielding some $b_m > 0$. Which outcome obtains is a function of mostly exogenous parameters capturing existing explanations: the agency's budget, agency wages, and the manager's incentivization of whistle-blowing.

---

on improving policy delivery, but it is all theoretically at risk of extraction.

[73]The junior retains w_j as a 'reward', this also be thought of as an expected future wage.

[74]As the junior was following the senior's lead, the wage is not forfeited.

2.5.6 What determines variation across agencies?

This simple model demonstrates the intuition behind the interaction I study, but is too abstract to push far beyond existing explanations. I complicate the model here using a narrative, rather than extending the model itself.[75] After the senior declares a corruption regime, the junior proposes a share of the agency budget to extract (rather than simply extracting all of it). The senior then agrees and proposes a split of the spoils, or fires the junior, as before. In the event there is still a job to do, the junior carries it out and reports back on what was extracted, but may have 'skimmed' off the top beforehand. The senior observes this report imperfectly, and either accepts it and distributes the agreed shares of the reported spoils, or fires the junior (here the junior retains whatever he hid from the report). If fired at any point, the junior can blow the whistle as before. Finally, the manager observes total extraction with some probability and if above her threshold, fires both bureaucrats. Just as before, the level of corruption in an agency is determined by the success or failure of its bureaucrats to *overcome their conflict over shares, and instead cooperate to 'grow the pie'*.

Under this logic (which simply makes some choices continuous rather than binary, and introduces one source of uncertainty) whether the senior and junior bureaucrat can safely reach *robust collusion* and maximize extraction, depends on the leverage each holds in the bargaining situation. This in turn depends more sharply than before on the *information available to each about the opportunities for corruption.* When the senior bureaucrat is less-informed the junior bureaucrat can more easily 'play the system' and defect, avoiding the senior's 'tax' to obtain greater absolute benefits, even as the total amount extracted from the agency budget shrinks *ceteris paribus*. When they are equally-informed this gaming is not possible and both are incentivized to cooperate, maximizing total extraction and sharing the spoils equally.

---

[75]The model above achieves the most important task of this chapter: to show that low-corruption agencies can arise even in high-corruption settings, using my smaller set of assumptions compared to existing explanations. I leave a full formalization of the logic below to future work.

To understand the implications of the model in a concrete case, recall the example of Vladimir Smirnov's elaborate property investment scam quoted in Section 2.4 above: many officials colluded in order to sustain such collective gains through corruption. Any one of several linchpins could have defected in pursuit of a greater share for themselves, collapsing the returns of the scheme as a whole. As Dawisha (2014) chronicles throughout her book, a crucial factor in the success of these schemes is the careful sharing of information between the participants. In my explanation, which generalizes this observation, information sharing balances out asymmetries between actors, reducing the incentive to defect.

In the narrative extension of the model, the potential for information asymmetry is added by making the senior bureaucrat uncertain about whether the size of corruption spoils revealed by the junior bureaucrat is the 'whole truth', or whether some has already been skimmed off the top. This uncertainty can be motivated by several findings in the literature on principal-agent relationships, the most important of which is the informational advantage granted to agents by expertise.[76] Agencies vary in the degree of expertise necessary to effectively implement policy, because the spheres of activity they regulate differ. I simply adapt the same idea to corruption: the same expertise in the actual work of the agency will also benefit a bureaucrat seeking to extract the agency budget through corruption. A nuclear regulatory agency, for example, requires a large proportion of its workforce to have deep technical skills in order to avert disaster. This means the senior bureaucrats, even if themselves specialists, cannot be expert at the jobs of all their subordinates. A tax agency, on the other hand, can get by with small teams of specialists in specific policy subdomains, but much of the work of the agency is purely administrative. Here the day-to-day work undertaken by the junior bureaucrat requires no special expertise beyond what their boss already possesses. When the senior bureaucrat more easily

---

[76]Often modeled as special knowledge about the 'state of the world' held by one actor.

understands what the junior did to implement policy (or extract rents via corruption), it is more difficult for the junior to pull policy in their preferred direction (or hide extraction).[77]

Any real agency is obviously more complex than what I model here, but the fact the senior bureaucrat is responsible to another official above her implies several of these relationships could be nested inside one another. This models top-level agencies (e.g. ministries) as a chain of bureaucrats, each responsible for spending decisions over successively smaller pots of money on successively narrower domains within the agency's remit.[78] I discuss in the next chapter how this flexibility helps use the full variation in the data to test my theory.

2.5.7 How does this within-state story resolve the puzzle?

When the national-level corruption regime is strict (represented as high returns to whistle-blowing) the two bureaucrats never have an incentive to cooperate in pursuit of rents. This is why low-corruption states have only 'clean' agencies. To stretch the model, in those contexts the incentive to defect grows faster (in agency budget, and in the proportion extracted) than either bureaucrat's share: the opportunity cost of silence is so great that neither can ever steal enough to pay the other to keep quiet. When this underlying control mechanism is absent, the mechanisms of my model take over, making sustained cooperation feasible, but never guaranteed. This produces some clean agencies and some that are far better off; what drives this difference is the structure of information inside the agency. When the junior bureaucrat has significantly better information about the opportunities for corruption (also referred to above as *expertise*) he can effectively steal from the senior under the

---

[77]In a formalized version of this narrative, information advantage would be represented by the senior bureaucrat's uncertainty as to whether the junior's report of corruption spoils is genuine.

[78]For example: the Minister/Secretary sets the budget for each division of the agency in consultation with division heads; the division heads do the same with respect to their branches; and so forth. The common kernel of all these recursive relationships is one bureaucrat responsible for spending a particular budget and accountable to another bureaucrat for how it is spent.

latter's nose. Recognizing this risk, in equilibrium the senior bureaucrat fails to authorize corruption, reducing the cooperative gains to corruption, and causing the junior's anticipated returns from unauthorized corruption below his licit wage.[79]

2.5.8 Making the leap from theory to empirics

The theory I have built in this chapter is general, and could plausibly apply to any context where bureaucrats in a hierarchy a) make a decision, b) have the opportunity to depart from the bounds of discretion granted over the decision, and c) have a mechanism to profit from that departure. In the next chapter I advance a research design that applies the theory to a specific context, agency procurement, which satisfies these conditions.

The logic of a well-functioning procurement market, described in the literature review earlier, is simple. Anything that enhances competition will more likely decrease than increase corruption; these factors were described in the table following the literature review. Two in particular, *L1e* and *L3c*, are modified by the model and narrative outlined above, and I take these to the data in Chapters 3 and 4. *L1e* hypothesized that corruption opportunities generate actual corruption, whereas I showed conditions under which this will not be possible: when the junior bureaucrat must extract rents from an activity requiring greater expertise, the senior bureaucrat will not support corruption in fear of missing out on rents while risking punishment.[80] The *L3c* provides the link between the individual level in *L1e* and the agency level.[81]

Combining these, the core implication of my theory is that bureaucrats' likelihood of engaging in corruption is increasing only when they can generate opportunities to extract *and* they are not too reliant on each others' expertise to take advantage of

---

[79]Though it stretches in to normative territory, I sometimes narrate this below as the senior bureaucrat 'not trusting' the junior enough to cooperate.

[80]In procurement, this corresponds to less-purchased products for which greater investigation is required regardless of corruption intentions.

[81]I re-emphasize this logic in detail when operationalizing the theory in Chapters 3 and 4.

these opportunities. Conversely, the likelihood of corruption decreases when either opportunities to extract fade away, or too much expertise is required to take advantage of these opportunities. For example, imagine Agency A's mission generates significant opportunities for corruption, but also requires a greater degree of expertise at junior levels to achieve this mission. Here, the absence of a level informational playing field means bureaucrats cannot cooperate to extract, despite ample opportunities, and the agency is 'clean'. Now imagine Agency B's mission does not require such expertise to be carried out effectively, but also does not generate as extensive corruption opportunities. Here, the gains from cooperation disappear not due to information asymmetry but from lack of opportunities, and the agency is also 'clean'. Even in corrupt states, corruption only succeeds in agencies where opportunities are high enough, *and* the information asymmetry described above is low enough. I carry this core implication forward to the remaining empirical chapters as an interactive hypothesis: the extensiveness of opportunities for procurement corruption changes how the information asymmetry between bureaucrats (implied by an agency's purchasing choices) affects corruption, and vice-versa.

CHAPTER III

Testing the theory while holding location fixed: Moscow-based agencies

"We estimate losses from corruption in public procurement amounted to US$30 billion in 2012." - Audit Chamber of the Russian Federation[1]

In the previous chapter I explained how corrupt systems sometimes produce clean agencies, which are puzzling given existing explanations. My theory builds on principal-agent models in which corruption is the outcome of successful collusion between bureaucrats. It differs in that it produces both corrupt and non-corrupt agencies without assuming either bureaucrat is concerned about social welfare. Clean agencies in corrupt contexts are simply the product of a *failure to coordinate* on the risky collusion necessary for corruption, caused by uncertainty in the bargaining situation between senior and junior bureaucrat. When uncertainty is mitigated, bureaucrats can engage in robust collusion and their agency is corrupt.

In my explanation this uncertainty is a consequence of the degree of expertise required to run the agency and generate opportunities for corrupt extraction in the process. As in canonical principal-agent models, the more expertise needed the greater the informational advantage of the junior bureaucrat over senior. Absent corruption

---

[1]This amounts to approximately 7% of total government expenditures, and 17% of procurement expenditures for 2012.

as an income stream, this leads to greater delegation to the better-informed actor: the principal cedes some decision-making power to learn more about the state of the world and 'splits the difference' with the agent. My model complicates this by making the senior bureaucrat uncertain about the junior bureaucrat's report on the extent of illicit takings from corruption. This uncertainty increases in the junior's informational advantage, to a point where no report would be credible, leading the senior to defect from collusion and withdraw support for corruption altogether.[2] The senior passes up the opportunity because safe income (salary) is preferable to the risk of having enabled corruption without sharing enough of its benefits.

In this chapter I explain my empirical strategy for testing these implications, which is complicated by the fact corruption is not directly observable beyond individual cases.[3] I begin by discussing the threats to inference presented by the theory I have built, and how to guard against them. These threats suggest a single-country, many-agency design is most appropriate. I then explain how Russia meets these conditions, and describe the procurement process that generates my dataset. With the assistance of related literature using similar data, I then propose and validate a series of proxy measures of corruption. Finally, I translate the implications of my theory in to an empirical model, which I estimate using data from agencies' spending in Moscow only. This holds location and market fixed, allowing me to examine the full set of goods and services procured by each agency. In the following chapter I build up a comprehensive analysis of individual purchases, and how agencies vary making them.

---

[2]Formally, in the model, the senior bureaucrat is choosing whether to delegate approval to engage in corruption to the junior. Delegation is necessary for what I call cooperation to generate the collusive equilibrium: without it neither actor can effectively extract rents. When the collusive equilibrium breaks down, delegated approval to extract rents via corruption is withdrawn.

[3]Individual cases of corruption are generated by a process more complex than the simple decisions I model in my theory. Individual corruption cases may be generated by political motivations when no underlying corruption has occurred; likewise where underlying corruption is undoubtedly endemic, only a fraction of those instances are revealed as individual cases. This phenomenon was documented in Hungary: in a comprehensive review of court documents, the *Corruption Research Center Budapest* found no court orders relating to public procurement from 2009-2014 (Corruption Research Center Budapest, 2015), despite extensive press coverage of corruption. These possibilities muddy the link from revealed corruption back to actual underlying corruption.

3.1 Research design

The puzzle I am trying to explain is the presence of low-corruption agencies in states where existing theory predicts all agencies should be high-corruption. As discussed in the previous chapter, this puzzle exists because explanations of corruption have not yet delved far enough inside the state, and because suitable data have not been available until recently. I tackle both problems in this dissertation.

In this section I describe how I test my theory using a research design focused on very fine-grained records of bureaucrats' behavior in Russia. Though focusing on a single country avoids many potential confounds it creates a new threat: variation across agencies on unobserved characteristics. One is the 'mission' of an agency, i.e what it was set up (and has evolved) to do within the state. Apart from belonging to the United States government, are the Department of State and Social Security Administration comparable cases? Here my procedural definition of corruption helps. Defining it as 'misuse of public office for private gain' would only exacerbate unobservable heterogeneity as the opportunities for misuse vary widely by agency. As I adopt a procedural definition (deviation from prescribed procedures when making a decision) a set of decisions common to all agencies would reduce the risk of comparing apples and oranges. This set of decisions, and prescriptions for how to make them, is the procurement of goods and services by agencies. Every paper-pusher needs paper, and every decision to purchase paper is an opportunity for corruption.

3.1.1 Empirical implications I will test

Most theories of corruption (including mine) predict that corruption opportunities are correlated with actual corruption: the more the door to the safe is left ajar, the more likely are bureaucrats to take advantage. This dissertation was motivated, however, by the inadequacy of this unconditional hypothesis. If it were true, all agencies in Russia would be as corrupt as each other. Instead, I advanced a conditional

version of this hypothesis: the conversion of corruption opportunities in to illicit gains depends on successful collusion between bureaucrats, which in turn depends on the information asymmetry generated by the specificity of the agency's purchases.

My theory predicts that increasing the information advantage of the junior bureaucrat (greater purchase specificity) decreases the likelihood of corruption, all else equal. This negative effect should be strongest when corruption opportunities are relatively scarce (high median auction efficiency); as corruption opportunities become more plentiful this effect should weaken, because the potential reward to the senior bureaucrat eventually overcomes the risk of punishment.

3.1.2 Balancing within- and between-state variation

The ideal research design to test an explanation of this puzzle would use measures of corruption that are equally comparable within and between states. The unit of analysis would be the individual agency nested within a state, allowing empirical separation of variation in corruption due to national-level factors from variation due to agency-level factors. This is the 'horse-race' set up in the introduction, when I demonstrated the curious cases of clean sectors in corrupt countries.

While the *Global Corruption Barometer* (GCB) data were suitable for motivating the puzzle, they are not fit for the purpose of testing my theory for several reasons. They measure citizen's reported experiences with corruption, which only partially capture the decisions of bureaucrats.[4] Also, to aid comparability across states, the GCB uses 'sectors' rather than agencies, which further separates the measurement from the agency in my theory. Finally, the sector groupings aggregate agencies together, obscuring most of the factors my (and other) theories claim cause corruption. Due to these shortcomings I set aside the GCB data until I return to exploratory

---

[4]The question generating the reported bribery rates is about payment of bribes for services, rather than requests for bribes. Whether a bribe is paid by a respondent within the last 12 months is the outcome of a long chain of decisions, only one of which maps to my theory.

analysis in the conclusion. For now, I require another source of data fit for testing the core mechanisms of my theory against alternatives. No other potential source improves on the GCB by satisfying the ideal of both within- and between-state variation across comparable units. This leaves a choice among second-best options that should be informed by a conceptual trade-off: to sacrifice either within- or between-state variation in pursuit of the other.[5]

I have argued from the outset that clean agencies in corrupt states are puzzling, because this within-state variation is not predicted by extant theory. My literature review showed how the dominance of cross-national explanations of corruption, and the between-state variation necessary to test them, has obscured this puzzle. This makes the optimal trade-off clear: as my theoretical contribution is specifically to go inside the state, the best option is to maximize within-state variation (and in turn the comparability of observed units), at the expense of between-state variation. Despite my trade-off in favor of within-state variation, between-state variation still matters for my choice of which state to go inside. As demonstrated earlier, there is no puzzle in the existence of clean agencies in clean states, so I need a setting where corruption is common in order to observe both *corrupt-as-expected* and *clean-as-unexpected* agencies within the same system.

Russia meets these criteria well. In national-level measures of corruption such as the *Corruption Perceptions Index* it is consistently ranked in the lower half of states. As demonstrated by the quote opening this chapter, corruption is common, extensive, and recognized as such by citizens and officials. And as I show below, there is significant within-state variation in proxy measures of agency corruption. Corruption in Russia is widespread but far from total: it has many clean agencies.

---

[5]For further discussion of this trade-off, and an application to measuring state capacity at the agency level in Latin America that inspired my approach here, see Gingerich (2013).

3.1.3 Selecting data to measure within-state variation

In setting aside the GCB data above, I alluded to criteria that would be satisfied by within-state corruption data suitable for testing my theory. First, data should measure the decisions of bureaucrats as directly as possible, and these decisions must generate opportunities for corruption. Data about revealed or experienced corruption fail these criteria because they mix together the decisions of bureaucrats inside my model with decisions by actors outside it. Investigations and prosecutions are the result of many bureaucrats' decisions, and reported bribery rates capture responses to a corrupt decision, not the decision to ask for a bribe itself.[6]

Second, data should be at the agency level (or even within-agency) so the context of bureaucrats' decisions is consistent with my theory and alternative explanations. Data comparing sectors rather than agencies, branches of government (e.g. executive vs legislature), or political institutions (e.g. political parties) offer within-state variation but are not close enough to the decisions under study.

Finally, the decisions measured should be comparable across agencies; this requires abstraction away from the specific mission of each agency. Each agency is set up for a different purpose, which clearly confounds explanations of *any* difference between agencies. Even if it were possible to measure bureaucrat's decisions about passports in one agency, and compare wait times and bribe requests to another agency issuing construction permits, the context and decisions are too different for useful comparison.[7]

Public procurement data meet all three criteria. Public procurement is the process by which agencies purchase goods and services from external market actors using

---

[6]An example is the GCB data discussed above. Such sources may nevertheless be useful as corroborating evidence, once an initial decision is identified and measured.

[7]Sources such as the World Bank's *Doing Business* report provide within-state variation across geographic units, but ask firms about their experience with specific decisions made by particular agencies. This provides a comprehensive view of business interactions with government as a whole, not with each agency.

government funds, under a framework of regulations specified for government as a whole. Taking the criteria in reverse: *all government agencies purchase goods and services* regardless of their mission, and many of these items will be common to all agencies; government funds are usually appropriated and spent at the agency level, because *agencies are the dominant organizational form within governments*; in public procurement *bureaucrats approach the market seeking an outcome* (which may include extraction of funds through corruption) and their decisions are usually separable from those of potential suppliers choosing their respective strategies.

There are important caveats to the argument above. Public procurement is not the *only* way agencies spend government funds: they disburse funds directly to citizens and firms as transfers, purchase employee time through salaries or wages, and buy goods and services outside a procurement process.[8] Politicians and bureaucrats may even directly steal government funds from agency budgets, a type of corruption I leave aside here to focus on opportunities for extraction that arise from the 'normal business' of agencies. Despite these caveats, the key advantage of procurement as a source of data remains: all agencies do it to some degree, making it a plausible source of comparable measures of bureaucrats' decisions *in a context where corruption is possible.* These caveats are further mitigated for a given state if public procurement: a) is centralized, so that all agencies use the same system governed by the same rules; b) has a low bar for inclusion, so that not only large purchases (highways) but also small (printers) are governed by the system, and; c) bureaucrats have significant discretion (within a common framework) over how to purchase goods and services, so that different patterns of decision-making can be inferred reliably. All else equal, the more centralized, inclusive, and discretionary is a state's public procurement system, the more useful are its data for testing explanations of my puzzle.

---

[8]This is particularly true for agencies with a security or military mission, where a significant proportion of the total budget may be obscured or secret.

3.1.4 Limits of public procurement data for studying corruption

The objection that procurement data can never overcome is that they do not directly measure underlying corruption. As argued above, we should be wary of any data claiming to directly measure corruption: as a social process in which secrecy is crucial, what can be directly observed is unlikely to be an informative revelation of the underlying phenomenon. When choosing among the second-best options, all we need is that a proxy measure increases as probability of corruption increases.[9] Based on my theory and the criteria developed above, procurement is the safest area of bureaucratic activity to look for such proxies, which I develop and validate below.

3.1.5 Choosing useful public procurement data

Public procurement has historically been a manual process: agencies publishing notifications in an official gazette, which are read by potential suppliers who submit competing bids. The winning bidder then concludes a contract with the agency and provides the goods or services. Any manual system is unlikely to meet all the criteria above (centralized, inclusive, discretionary), because of resource-intensiveness of the process. This constraint has typically led agencies to use public procurement for only the largest purchases, and a separate internal procedure for smaller purchases. When purchases run through the centralized procurement system are larger, they are more likely to be idiosyncratic to the mission of the particular agency, reducing comparability across agencies of any derived measures. Government must also live with lower visibility of how the procurement budget is spent.

The advent of modern information systems, and a means to connect them (the Internet), led to the rapid adoption of electronic procurement systems by governments from the late 1990s.[10] Initially these were simply a faster way to distribute tender

---

[9]Escresa and Picci (2015) stated this concept succinctly: "A valid index in the lexicon of this article means one that increases along with the probability that a transaction is corrupt."

[10]One stated aim of electronic procurement has been to reduce corruption, primarily through

notifications to the market, but over time evolved to run the entire process described above (notification through to contract) within one system. Though most states have moved at least some of the procurement process to a centralized electronic system, few have gone all the way. The United States, for example, operates a centralized listing service for all federal agencies[11] and provides an option for agencies to accept bids through the same system. Agencies can, however, choose to run the bidding process through their own system. Once contract awards are made (through whatever system), a separate unified system publishes those data.[12]

Several European states' public procurement systems have gone further than data aggregation on the notification and contract ends of the process. The Czech,[13] Hungarian,[14] and Slovakian[15] governments all run centralized public procurement systems that cover the entire procurement process: notification of intent to purchase, submission of bids, choice of winner, and contract registration. The problem with using data from these states is that clean agencies in EU members states are not surprising. While the extent of corruption across agencies within those states varies,[16] it is within a narrower band than the ideal setting for my puzzle: states with extensive corruption and unexpectedly clean agencies. As mentioned above, Russia is a good candidate state given the puzzle. Its clean agencies are certainly unexpectedly so, given how it scores on national-level corruption measures. Russia also operates one of the most centralized, inclusive, and discretionary procurement systems in the world. Furthermore, it is a significant source of government expenditures: the total procurement budget for 2012 in Russia was approximately US$188 billion, out of a total federal budget of approximately US$450 billion.

---

improving the transparency of government purchases (Olken and Pande, 2012).

[11]The US *Federal Business Opportunities* portal is located at `https://www.fbo.gov/`

[12]The US *Federal Procurement Data System - Next Generation* is at `https://www.fpds.gov/`

[13]Located at: `https://www.vestnikverejnychzakazek.cz/`

[14]Located at: `http://www.kozbeszerzes.hu/`

[15]Located at: `http://www.uvo.gov.sk/`

[16]Discussions with authors of research using these data confirm variation across agencies, though none have built a research agenda around it. They expect agencies to be clean by default.

3.1.6 Russian public procurement as a source of within-state variation

Russia is a suitable case for several reasons: it spends a lot on procurement, losses due to corruption are well-recognized, and below I demonstrate significant variation across agencies in the behavior of bureaucrats responsible for procurement.

The centralization of Russian public procurement culminated with implementation of an end-to-end online system in 2011. Prior to this, each of the 85+ subjects of the Russian Federation (regions, republics, and other sub-national entities) ran their own public procurement process (some online, some offline),[17] under a patchwork of regional and federal regulations, including Federal Law No. 94-FZ "On State and Municipal Procurement of Goods, Works and Services", passed in 2005.[18] From 1 January 2011, all new tender notifications by all public entities[19] were listed on a new portal at `www.zakupki.gov.ru` and Federal Law 94 was further amended to override the heterogeneous regional regulations. My data begin from this point, and continue until 2015, when the procurement regime was changed again.[20]

A key feature of the Russian system is its inclusiveness. In many EU states, agencies are required to use the public procurement system (rather than their own methods) only for relatively expensive purchases.[21] EU-wide rules for mandatory use of the EU procurement framework apply only at even higher thresholds, when

---

[17]These systems varied widely in their coverage, accessibility, and depth of detail. Many only covered a small subset of all procurement, for tenders above relatively high values. There was wide variation in the extent to which they made information easily accessible (eg plain text vs scanned hard-copy documents), and in the level of detail about procurement. One justification for centralization was a reduction in the burden on potential suppliers of monitoring dozens of portals.

[18]An English translation is available from the Federal Antimonopoly Service at `http://en.fas.gov.ru/documents/documentdetails.html?id=13920`

[19]The definition is broad, and includes all federal and regional government agencies, as well as state-funded institutions such as hospital and universities.

[20]Federal Law 94 was replaced with Federal Law 44 on 1 January 2014, but adoption by agencies was not immediate. The Economic Development Minister at the time, Andrei Belousov, stated confidently that the new law would almost rule out kickbacks, see: `https://themoscowtimes.com/articles/minister-says-tender-bill-to-eliminate-kickbacks-22940`

[21]In Hungary, for example, the threshold is €10,000. In Germany it varies by sub-national jurisdiction but is generally much higher.

contracts are for hundreds of thousands of euros.[22] In stark contrast, Russian Federal Law 94 exempted from its framework only those tenders considered 'cash transactions' by the Central Bank, a threshold held at 100,000 rubles since 2007. This amounts to US$3,300 at the average exchange rate across 2011-2015 of 30 rubles to the dollar, well below the threshold for other states. Due to this low threshold at which the public procurement framework starts to bind agencies, we can observe a much higher proportion of purchasing activity than is common in other states.

Beyond this high degree of centralization and inclusiveness, the Russian public procurement system from 2011-2015 is useful for the discretion it affords bureaucrats, within a unified framework. On one hand, the low bar for inclusion under the rules of Federal Law 94 disincentivizes gaming to avoid oversight altogether. This limits the kind of discretion bureaucrats have in many other systems to carefully split one logical purchase in to several smaller tenders, each falling below a certain threshold.[23] On the other hand, having brought almost all agency purchasing under its remit, Federal Law 94 then permits significant freedom to bureaucrats caught in its grip. I explain the bounds of this discretion below.

3.1.7 Corruption in Russian public procurement

I have argued my theory would be best tested using within-state variation from a corrupt state, that procurement data is the best source of that variation, and that these requirements make Russia a good candidate. Corruption is widespread, touches every agency, and could be part of almost any interaction with a government official. Public procurement is no exception, as these examples show.

A Reuters investigation of public procurement in the health sector from 2006-13 found "Russian state hospitals frequently pay two or three times more than hospitals

---

[22]Thresholds are published at `http://europa.eu/youreurope/business/public-tenders/rules-procedures/index_en.htm`

[23]Procurement of a long highway could, for example, be split in to several small purchases, each falling under less scrutiny than the whole.

in the West for the same equipment".[24] Of course some difference in price is to be expected for any export of expensive equipment, but the investigation also found a bewildering array of intermediary companies involved in these transactions, for which a leading hypothesis must be to hide the many 'cuts' added between manufacture and final delivery. Given the demand for this modern equipment and the availability of state budget funds to purchase it, the relatively small group of health officials making such purchasing decisions face a problem, even if they are not corrupt: they need to obtain it but various intermediaries will insert themselves in to the process. This could be the head of the hospital, or a customs official, or a medical expert who helps choose between the alternatives. In order to conclude the purchase, all of these potential veto players must be compensated, which requires creation of a wedge between the actual market price and the final sale price. As per the definition adopted in Chapter 2, *bureaucrats create this wedge by deviating from prescribed procedure*. If the equipment manufacturer wants to make a sale, they will be forced to accept this 'padding' by intermediaries, and enabled by the purchasing bureaucrat. This is not an isolated example, and such violations are clearly big business:

> In his presentation the auditor also drew attention to the large number of violations detected in public procurement. "During the first nine months of this year, according to the results of operational and expert monitoring, the Audit Chamber found about 400 violations of procurement legislation (both procedural and financial), totalling about 112 billion rubles [US$2 billion]." - Audit Chamber press release, 25 November 2014 [25]

Further evidence of the volume of rents presented by procurement is the frequent arrest of anti-corruption officials. In a system where corruption is endemic, control

---

[24]Source: The Moscow Times, *Russian Hospitals Paying Over the Odds for Western Medical Equipment*, 19 December 2014, http://www.themoscowtimes.com/business/article/russian-hospitals-paying-over-the-odds-for-western-medical-equipment/513693.html

[25]Available at: http://audit.gov.ru/press_center/news/19640

over which investigations proceed and which falter is lucrative. In 2014, the head of the Interior Ministry's anti-corruption division was arrested in a sting operation relating to bribes paid for stopping investigations in to procurement fraud.[26] In 2015, the head auditor of procurement at the Ministry of Defense was detained on suspicion of heading an extensive network of procurement corruption.[27] Finally, this quote emphasizes the sheer scale of diversion of state funds in Russia from crucial needs, which must end up somewhere:

> "Despite receiving $1.6 trillion from oil and gas exports from 2000 to 2011, Russia was not able to build a single interstate highway during this time. There is still no interstate highway linking Moscow to the Far East; in contrast, China, another top-down authoritarian regime, has built 4,360 miles of modern highways annually for the last ten years - equivalent to three times around the circumference of the earth." (Dawisha, 2014, p.314)

I provide more examples of corruption specifically in procurement in the next chapter, to help explain my more detailed approaches to corruption detection.

3.2   Data

In this section I describe the bureaucratic process that generates my dataset, introduce the data themselves, and illustrate their plausibility as measures of the key agency characteristics necessary to test my theory. In the following section I apply these measures to agencies' spending within Moscow.

---

[26]Source: http://www.rbc.ru/politics/01/09/2014/946145.shtml
[27]Source: http://newsru.com/russia/22dec2015/zakazotkat.html

3.2.1 Details of the Russian public procurement system 2011-2015

Russian officials can choose from three distinct procedures for making a purchase, outlined in table 3.1. In **electronic auctions** they decide what to buy and a maximum (starting) price, then a reverse auction proceeds and the lowest bidder automatically wins.[28] Though intended to minimize bureaucrats' discretion, the exclusion of quality standards has led some to speculate the opposite (Balsevich and Podkolzina, 2014). A supplier may win and provide a technically compliant (but not fit-for-purpose) good, sharing the profits with the organizer. **Open tenders** clearly allow more discretion to bureaucrats, as they introduce non-price criteria in to the decision over which bidder is granted the contract. This is useful in more complex purchases where quality might be important in assessing bids. The introduction of alternative criteria also creates a mechanism to tilt the playing field towards a particular supplier, an opportunity bureaucrats might choose to monetize. This opportunity is somewhat limited by the openness of tendering: any firm may submit a tender without passing any pre-qualification check, which is common in other jurisdictions.[29]

In a **request for quotes**[31] bureaucrats announce their requirements, and choose among the proposals by selecting the lowest-price bid that meets their criteria. This allows well-connected firms influence over the way the requirements are phrased, which in turn gives leverage to bureaucrats before notifying the market. Nevertheless, the only formal criterion is price, so the procedure provides less discretion than the open tender with multiple criteria. Due to the extensive risk of corruption under

---

[28]The Government publishes a list of goods and services 'for which there is a functioning market and which can only be compared by their prices' (Article 10, Part 4, Federal Law No. 94 of the Russian Federation 2011); electronic auction is mandatory for purchasing these commodities.

[29]The ability to use pre-qualification to whittle down the field of eligible suppliers, by introducing idiosyncratic criteria, provided one of the justifications for Federal Law 94's banning of this technique altogether. This had several unintended consequences: truly unqualified firms that find another way to affect the relevant bureaucrat's decision can avoid scrutiny; those bureaucrats committed to clean procurement are deprived a key tool for guaranteeing the quality of specialized goods and services. For further commentary see Shamrin and Yakovlev (2009)

[31]Sometimes translated as 'sealed-bid auction' because the bids are submitted in sealed envelopes, and by the rules should all be opened at once.

Procedure	Discretion	Conditions of use
Electronic auction	Ambiguous	Mandatory for certain goods and services specified by the Government
Open tender	High	Recommended for technically complex projects with important criteria besides price; mandatory for purchases above 3,000,000 (US$100,000)
Request for quotes	Ambiguous	Not allowed for purchases above 500,000 rubles (US$16,000)[30]
Single-supplier	Highest	Only allowed when goods produced by a monopoly (natural or otherwise)

Table 3.1: Summary of procedures available under Federal Law 94

such procedures, its use is limited to purchases below 500,000 rubles (US$16,000).[32] Finally, the **single-supplier** procedure offers maximum discretion to bureaucrats. There is no competitive procedure and the purchasing agency contracts directly with its preferred supplier; only the conclusion of the contract is notified to the market through the online portal. This procedure is only allowed in cases where the necessary goods and services are available from one supplier, either a natural monopoly or exclusive producer, conditions which have obtained for steadily more markets since 2000.[33] This procedure maximizes the risk of corruption, because official oversight is post-facto, and other firms can only protest when it is effectively too late.

In the process of purchasing goods and services via the public procurement system, bureaucrats make the following choices that I can recover from the database:

- Procedure: open electronic auction; open tender; request for quotes; single-supplier (subject to the rules outlined above)

- Product: a description of the good/service, and a standardized product code

---

[32] Article 42, Part 2, Federal Law No. 94 of the Russian Federation (2011). The risk of corruption arises from the shorter time-frames available under this procedure: while efficient for small purchases, this also gives an advantage to any supplier with advance knowledge.

[33] The standard economic justification, that it is inefficient to ask a market with one supplier to bid against itself, is less robust under Putin than in the 1990s. Since coming to power, a hallmark of Putin's political and economic strategy has been the re-monopolization of certain markets.

- Quantity: how much of the good/service is required (and a unit price)

- Maximum price: the most the agency will pay for the specified goods, which becomes the starting price in the reverse auction if that procedure is chosen (otherwise it is just a guide to suppliers)

- Revisions: number of revisions to the notification and contract

- Bidding window: time over which bids are accepted

- Final price: the price at which the contract is concluded

Data from these decision points generate measures of the concepts in my theory. For example, I capture the informational advantage of the junior bureaucrat using a measure of the uniqueness of the product 'basket' purchased by a given agency. Bunching of maximum prices around thresholds where rules change, e.g. where *request for quotes* procedures are not allowed, show attempts to game the system that are correlated with corruption. Likewise, when the difference between initial (maximum) price and final (winning) price is smaller, the opportunity for corruption is greater.

3.2.2 Obtaining and processing Russian public procurement data

That procurement in Russia was centralized to a single, federally-run portal, is not surprising given Vladimir Putin's well-documented drive to centralize political power generally.[34] Nor is it surprising that even while centralizing and bringing under control much more purchasing activity by setting a low threshold, bureaucrats were still afforded significant discretion in their choices. Puzzling, however, is the stark openness of the system and its data, which has caused considerable, albeit informal, debate among Russia experts. The data I describe below were not scraped semi-legally off a website, but downloaded en masse, following official instructions posted

---

[34]The centralizing decisions made early in his presidency, which set the stage for the entire post-Yeltsin period, are well-documented in (Dawisha, 2014)

by the maintainers of the portal at `www.zakupki.gov.ru`. As well as bulk data, any individual purchase through the system can be discovered and examined by all.[35]

The publicly available data from the portal at `www.zakupki.gov.ru` is downloadable from a back-end FTP connection.[36] The total payload is about 70 gigabytes, and contains tens of thousands of zipped XML files describing the separate stages of the procurement process (e.g. notification, contract), grouped by the region of the purchasing agency. XML files are only loosely structured, and generating tabular data useful for measurement requires parsing each file for the relevant content.[37]

I wrote a large program in the **R** language that processes the zip files, exactly as downloaded, in to the data I present below. This means every choice made during construction of my dataset can be audited, reproduced, and altered. The program is hosted online and includes instructions to download the raw procurement data, as well as other data sources such as the *Global Corruption Barometer*.[38]

As with any data from an administrative system, the user must make consequential decisions about which data to exclude for quality reasons, how to deal with duplicates, and the appropriate standard for matching data on various parts of the

---

[35]The openness of the system is certainly unusual given the general inaccessibility of government information in Russia. Ad hoc theories discussed on the sidelines of a conference I attended, during field research in Moscow, ranged from "best to hide rampant corruption in plain sight" to "openness helps cartels of suppliers monitor each others' compliance with bid-rigging arrangements". The theory I favor is that the process was made more transparent not to minimize corruption, but to maximize the ability of the federal center to monitor 'who gets what' out of procurement; this generates potentially compromising information that can be selectively used against particular actors inside and outside government, as necessary *for other purposes*. The more paperwork required to purchase goods and services (in the name of transparency), the more such information is generated, and the stronger the lever. These unintended consequences, and the general null effect on corruption levels, were noted already in the period between Federal Law 94's initial introduction in 2005 and the centralization to one portal in 2011, see Yakovlev (2010).

[36]Instructions in Russian, downloaded from the official site, are part of the software program I wrote, described below.

[37]The author needed to learn several new technical skills to transform the downloaded data in to a usable format *en masse*, but in principle every individual procurement in the system can be found by browsing or searching. Any moderately-skilled computer programmer would have little difficulty transforming the data as I eventually did, although she would need to know Russian and the details of the system to make sense of the data.

[38]Available in full at `https://github.com/shaunmcgirr/shaun-mcgirr-dissertation`. Interested parties can download the same data and change any one of my data quality decisions.

underlying process to each other. These decisions were relatively straightforward and the resulting losses inconsequential; I document these in **Appendix A.**

3.2.3   Accuracy of Russian public procurement data

This section motivates and provides context for the empirical test to follow. This is important because the data used to test my theory are generated by bureaucrats themselves, going about the business of procurement. For my statistical analyses to be credible, we must first establish these data are a defensible representation of public procurement, and in particular corruption. The general measurement approach I employ is to ask, for a given decision, how a compliant and efficient outcome would be represented in the data. Agencies where decisions frequently deviate from this norm are more likely to be corrupt than those where such deviations are rare. This is because deviation from prescribed procedure is necessary to create rents; there is no other mechanism by which the agency can overpay (Fazekas, Tóth and King, 2016).

An obvious question is whether these data represent the procurement process they claim to. If a system of record requires, for example, extensive manual copying of information from elsewhere, inaccuracies inevitably occur. These can be innocent, in the case of onerous data entry procedures or poorly trained staff, or malicious in the case of willful mis-entering of data for the purpose of obfuscation. Neither threat is likely to be significant in the case of Russian procurement data, because the web portal controlling the procurement process limits the choices of the bureaucrats using it. For example, though the full description of goods to be purchased is a free-text field, the user must also choose a product code and defined quantity from the relevant state standard; my analysis uses the latter.[39] The choice of procedure is likewise limited to the options allowed by the web-based system, and cannot be changed mid-purchase. Cryptographic digital signatures are required to run a purchase through the system,

---

[39]This is the *All-Russian Classification of Economic Activities, Production and Services (OKDP).*

and these are tied to specific organizations, identified by another state-administered scheme.[40] Finally, when announcing the conclusion of a contract, the user must select which notification (i.e. announcement of intended purchase) it relates to. This is only allowed to be missing in the case of unannounced notifications, under the single-supplier (highest discretion) procedure outlined above. These features of the system mean the wide discretion available to bureaucrats under the law is reflected in the data, without generating opportunities for wholesale manipulation.

3.2.4   Face validity of Russian public procurement data

Even though these data are an accurate record of bureaucrats' use of the underlying web-based procurement system, I must still satisfy two concerns: do the data represent the decisions under study, and can they detect deviation from specified procedure in those decisions? There is evidence that both can be answered affirmatively. Figure 3.1 shows how many of the 550,000 purchases notified in Moscow from 2011-2015 were listed at various maximum (initial) prices.

Recall that *request for quotes* (green) is not allowed for purchases above 500,000 rubles. If bureaucrats either ignored the rules completely, or played by them, we would not expect the large concentration of purchases right against this threshold. This shows the data capture sensitive procurement decisions of bureaucrats as they balance compliance with the rules against opportunities for corruption: they know purchases beyond this threshold may offer less discretion in the awarding process. If an agency's procurement activity is highly discontinuous at this threshold, it is more likely to be a site of corruption than an agency where this gaming is not observed.

A similar pattern is evident for *open electronic auctions*, where the bunching of listings right near 3,000,000 rubles is unexpected. Under this procedure, suppliers attend an online auction at a specified time, and submit competing bids. The lowest

---

[40]The *Taxpayer Identification Number* is unique to each individual or organizational taxpayer, see: `https://www.nalog.ru/eng/exchinf/inn/`
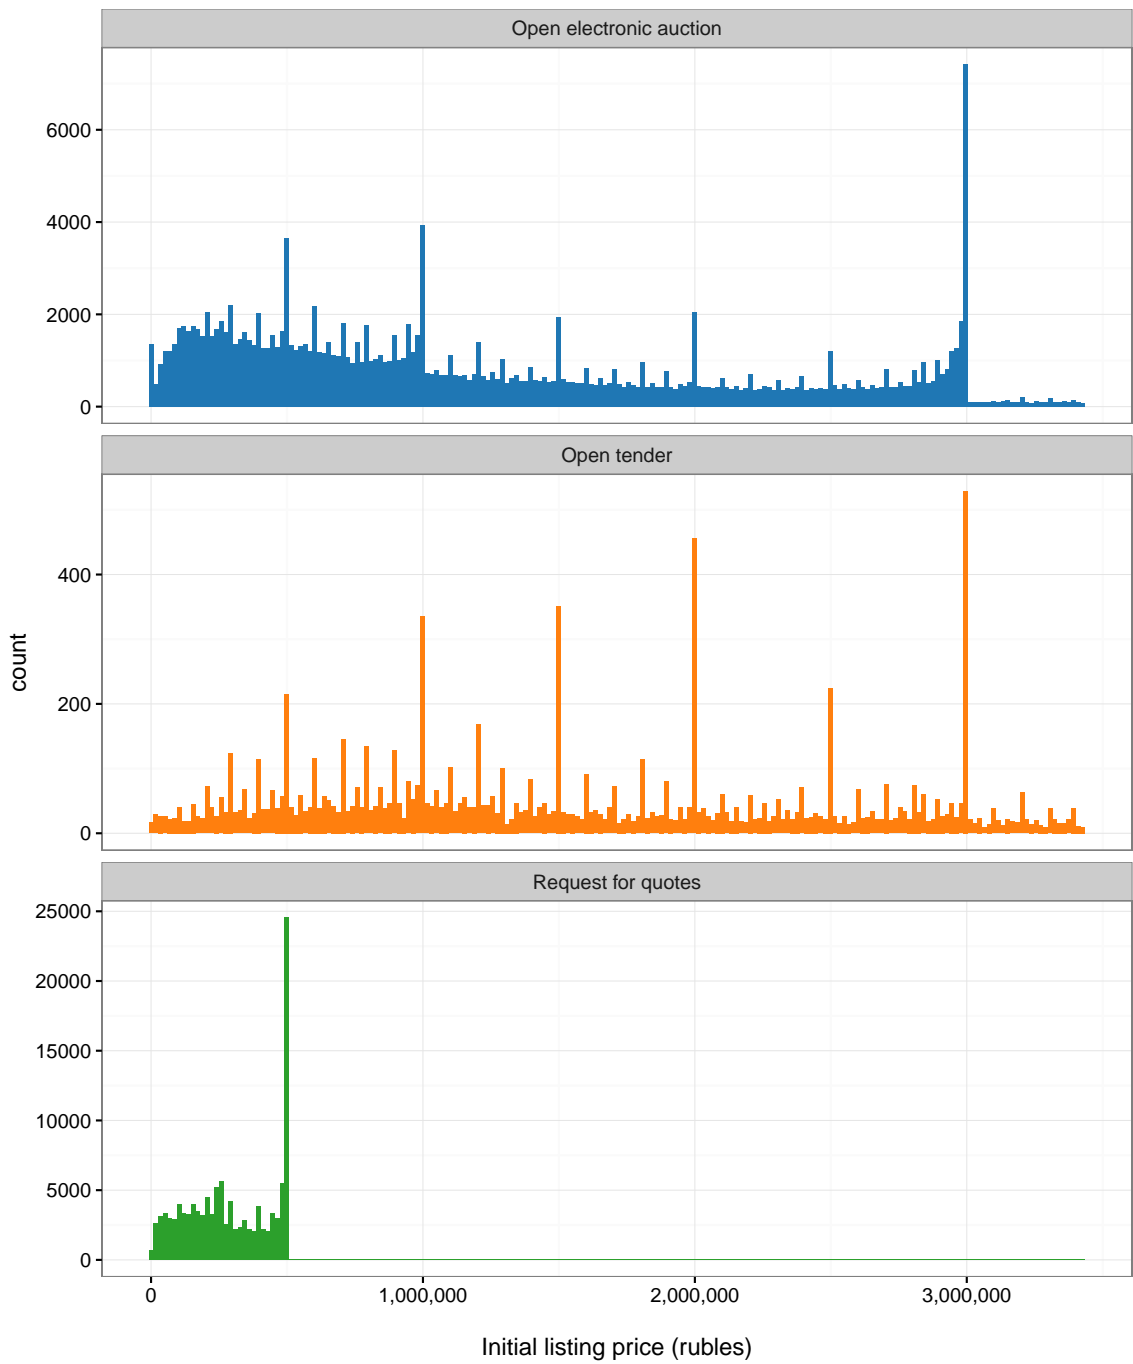
Figure 3.1: Frequency of purchases listed at each price level, by procedure. Note the bunching of purchases using the 'request for quotes' procedure near the 500,000 ruble threshold, above which it cannot be used.

price wins automatically, which reduces discretion of bureaucrats during bidding; any corruption must be organized in advance. Nothing in the law or regulations indicates additional scrutiny beyond this threshold, but the ability to bring a purchase to a close quickly under this procedure, by late notification of the details of the auction, is curtailed.[41] Figure 3.1 shows both that bureaucrats using the system respond to the rules built in to it, and that they game these rules consistently by manipulating the maximum (ie starting) price of auctions. The true distribution of the purchase price of items is simply highly unlikely to show such 'bunching' around these points of significant changes in incentives.

3.3   Measurement

So far I have demonstrated that these procurement data can detect responses of bureaucrats to the rules constraining their choices. Simply by examining the choice of procedure, I have demonstrated two patterns that should not exist if procurement is free from corruption: 'bunching' of initial prices near thresholds where rules change, and over-reliance by some agencies on particular procedures. These patterns in the data are not, however, robust evidence of corruption. As discussed above, direct measurements of corruption with the coverage needed to resolve my puzzle do not exist. To find *better* evidence of corruption than has previously been available I now turn to more detailed measures of purchasing activity possible with these data. The measurement philosophy does not change: the best I can hope for is that an increase in corruption is *highly likely* with an increase in my measures.

My measures are derived from two literatures, one explaining the efficiency of auctions, the other measuring corruption in Russia with similar data on a smaller

---

[41]Specifically, Article 41.5.2 of Federal Law 94 provides that as long as the starting price is less than three million rubles, the purchaser can delay open announcement of the auction details until just seven days before it proceeds. This obviously makes it more challenging for all bidders to arrange to bid. For auctions starting above this threshold, fifteen days notice is required.

Stage	Actions taken
1	Agency chooses what goods and services to purchase
2	Agency decides which procedure to utilize (open electronic auction, open tender, or request for quotes) under the constraints specified above
3	Agency notifies market of purchase, procedure, initial (maximum) price
4	Bidding proceeds according to the procedure, until the specified deadline
5	Contract is awarded to winner at a published final price
6	Goods delivered or services rendered

Table 3.2: Sequence of actions taken by agencies purchasing goods and services under Russian Federal Law 94

scale. I supplement these with additional measures of irregularities in the specific procurement process under study. For clarity through my discussion below, Table 3.2 shows the procurement process from start to finish.

3.3.1 Measuring opportunities for corruption with auction efficiency

Regardless of the procedures chosen by Russian bureaucrats using their public procurement system, we can apply existing knowledge from the analysis of auctions to understand how their behavior deviates from efficient outcomes.[42] Inefficiency is neither necessary or sufficient for corruption, but does create opportunities for corruption. For example, if an agency consistently purchases a particular good above the market price, at least one actor (the supplier) always benefits. From there, the jump to corruption is much shorter than the case without such 'padding'. Note that my study departs from previous work studying Russian procurement (Balsevich and Podkolzina, 2014; Yakovlev, 2012) in that *I do not consider 'padding' to indicate corruption per se*; under the definitions I adopted in Chapter 2 it creates a rent that may or may not be extracted by bureaucrats. For this reason, I will consistently refer to 'padding' that reduces auction efficiency as creating *corruption opportunities*, and rely on other measures to detect corruption itself.[43]

---

[42]For example, see Compte, Lambert-Mogiliansky and Verdier (2005), Lambert-Mogiliansky and Sonin (2006), and Menezes and Monteiro (2006).

[43]In the next chapter I outline in detail a case where padding clearly occurred, and very likely as a prelude to corruption. During the procurement procedure for a new hospital, the maximum

In order to create 'padding', the bureaucrat organizing the purchase must first set the initial (maximum) price above the lowest (market) price. The greater this difference, the more padding is available to extract as private gain. Padding represents the *opportunity for corruption*, and I develop an approach to measuring it here. Yet cultivating opportunities is only half the task: having created padding, the bureaucrat must prevent natural market competition from driving that initial price down, otherwise no padding remains to be extracted.[44]

I provided evidence above that at least some bureaucrats in Russia are concerned about restricting competition: they 'bunch' purchases below certain price thresholds, beyond which their ability to influence who wins would be restricted. As procurement corruption relies on goods and services being exchanged for money (unlike direct theft of funds) then, by extension, the same bureaucrats should also try to set an initial price above the market price. Without first creating padding, attempts to steer the gains to a particular bidder are futile, as there is nothing to extract. This implies bureaucrats seeking to make corruption easier will try to set initial (maximum) prices well above the cheapest price available in the market, and then restrict competition to ensure the final purchase price is little different. Bureaucrats not seeking to make corruption easier will tend to do neither: building padding in to the contract without intention to extract it generates risk without any potential reward.[45]

Unfortunately, the data available do not directly observe the full process by which bureaucrats set the initial (maximum) price.[46] This would allow a simple test of the

---

price increased five-fold as the purchasing agency continued revising the requirements to include more non-hospital-related equipment. Two bidders complained and were disqualified, leaving just two bidders.

[44]Recall that in this system, bids descend from an initial maximum price. How much of this gain accrues to the organizing bureaucrat vs. the winning supplier is a function of their respective outside options. The sharing of this private gain between the colluding bureaucrat and supplier(s) is the subject of much analysis, which I leave aside here.

[45]That a bureaucrat does neither is consistent with my explanation, but does not distinguish it from others.

[46]Under the *open tender* procedure, the agency must obtain two quotes and set its initial (maximum) price as the lower of these. The justification is this prevents abuse of the price criterion, as non-price criteria are introduced to the decision under this procedure. As discussed above, the

implication developed above, that restriction of competition and 'padding' of initial prices are positively correlated. A second-best option is to adapt the concept of *auction efficiency* to provide information about the opportunities for corruption.[47] In a conventional auction, efficiency means the seller gains the highest price from potential buyers. In a reverse auction the efficient outcome for the buyer is paying the lowest possible price. In the context of agency procurement, efficiency minimizes the opportunity for corruption, by making 'padding' less feasible. Theoretically then, auction efficiency is increasing (and opportunity for corruption decreasing) in the difference between the initial and final price, as long as that initial price is set exogenously.

As well as being unobservable in these data, the process by which bureaucrats set the initial (maximum) price is endogenous to the rest of the procurement process. Bureaucrats do not maximize agency efficiency, but rather their stream of licit and illicit benefits, so may manipulate the initial price precisely to generate 'padding'. This disqualifies auction efficiency from validity as a proxy for corruption, but as I demonstrate below the concept is still useful as a measure of corruption opportunities. I develop measures of corruption itself in the following subsection.

Suppose one bureaucrat in an agency purchases two identical goods, using the same procedure on the same day. If the difference between initial and final price is identical for both purchases, we cannot infer anything about opportunities for corruption from the gap alone. If, however, that gap *changes* between purchases, this could provide useful information in some cases, which I summarize in Table 3.3.

Despite the difficulties interpreting the corruption opportunities generated by changes in auction efficiency, a pattern emerges: a smaller difference between ini-

---

underlying logic of the reforms producing the data in this period was that the corruption to be minimized was predominantly manipulation of bids after the fact. This obviously ignored other avenues for manipulating prices and pre-selecting winners, in turn justifying further reforms in 2015.

[47]Though only one of the procedures available in Russia uses the word auction (*open electronic auction*, a multiple-round, open-bid, reverse auction) the *request for quotes* procedure is a single-round, sealed-bid, reverse auction. The *open tender* procedure includes non-price criteria so cannot be properly considered an auction. Regardless, the *open tender* is used infrequently as shown by the lower frequencies of the orange histogram in Figure 3.1.

| Initial prices | Change in observed difference between initial and final price for two otherwise identical purchases | |
	Smaller in second purchase	*Larger in second purchase*
Same initial price	Agency pays **more** for second purchase (more opportunity for corruption)	Agency pays less for second purchase (less opportunity for corruption)
Initial price higher in second purchase	Agency pays **more** for second purchase (more opportunity for corruption)	Agency pays **more**, same or less for second purchase
Initial price lower in second purchase	Agency pays **more** or same for second purchase	Agency pays less for second purchase (less opportunity for corruption)

Table 3.3: Linking changes in auction efficiency to opportunities for corruption

tial (maximum) and final (winning) price is more commonly associated with greater opportunity for corruption. Table 3.3 shows that only in one logical case is it even possible, and not guaranteed, that a larger difference is associated with a greater opportunity.[48] For all the cases where auction efficiency is lower for the second purchase (left column), it is always possible for the agency to pay more, and in two of the three this is the only inference possible.

Generating these six logical cases required strong assumptions in the underlying example: an agency buying the same good (to fix the suppliers), using the same procedure (to fix the rules), on the same day (to fix the price). These assumptions gave clarity over the interpretation of differences in auction efficiency, with a smaller gap between initial (maximum) and final (winning) prices more likely to create corruption opportunities than a larger gap. As outlined above, only in exceptional cases would

---

[48]This is when the initial price is set higher for the second purchase but the auction efficiency is also higher. In some case the higher auction efficiency 'overcomes' the effect of the increased initial price, but it is at least equally likely that the increase of initial price is greater than the increase in auction efficiency, leaving the agency to pay more. In the other ambiguous case, the initial price is lower for the second purchase, but so is the difference between initial and final price. Here the agency would only pay less for the second purchase if the change in initial price was greater than the reduction in observed auction efficiency. This could occur if the bureaucrat re-calibrated the initial price to market conditions after the first purchase, and new suppliers entered to further drive down the final price; this is less likely than the same bidders remaining (in which the final price is unchanged from the first purchase), which in turn is less likely than some bidders dropping out upon observing the lower initial price for the second purchase (resulting in unchanged or higher price).

a larger gap (i.e. greater auction efficiency) be associated with greater opportunity. In general, padding is more likely to have occurred the less the price changes across the course of the auction. If these assumptions can be relaxed safely, the concept of auction efficiency measures corruption opportunities more generally, and not just in pairs of otherwise identical purchases. Recall the test for validity of a proxy: when the underlying phenomenon increases, it must be more likely that this measure increases than decreases.

First, relax the temporal assumption, that the two purchases occur on the same day, which was made to ensure the lowest price would not change between purchases. Instead allow them to occur a week or month apart, but hold the initial price fixed for the sake of illustration.[49] If, for example, a supplier that failed to win the first auction managed to cut its costs further, it could make a lower bid in the second auction, generating a larger observed difference between initial and final price. Conversely, if the winning supplier from the first auction went out of business in the week or month between auctions, the final price would be higher, resulting in a smaller observed gap between initial and final price. In these examples, the change in auction efficiency is related only to the elapsed time and resulting changes in the composition of suppliers in the market.

An alternative explanation is available, however, for the higher final price in the second example. Instead of going out of business, the previous winner was somehow excluded from the second auction, allowing the agency to pay an above-market price, which creates a rent that in turn makes corruption more likely. Importantly, such a corruption-driven alternative explanation is not possible in the first example, in which the difference between initial and final price is greater in the second auction. There, regardless of the time elapsed, the lower bid by an existing supplier (or entry

---

[49] As described above, the greater likelihood of corruption when auction efficiency is lower (smaller difference between initial and final price) generalizes from the case where the two purchases have the same initial price, to the cases in Table 3.3 where the initial price differs.

of a new supplier) cannot be the result of collusion to restrict competition; the lower price can only be the outcome of natural competition. This implies that regardless of the time elapsed between purchases of the same good using the same procedure, lower auction efficiency can be associated with either corruption or natural changes in market composition, whereas higher auction efficiency is only associated with the latter. If the underlying opportunity for corruption increases between two auctions, therefore, it is more likely that auction efficiency decreases than increases. Put simply, a smaller gap between the initial (maximum) and final (winning) price is more likely to be associated with increased corruption opportunities than is a larger gap.

Similar logic applies when relaxing the assumptions that the two purchases use the same procedure and are of the same good. Given a particular good, and a sufficiently narrow time period between purchases, if an agency's purchase using procedure A results in greater auction efficiency than a purchase using procedure B, the opportunity for corruption under procedure A is less than under procedure B. And given a procedure and narrow time period, if an agency's purchase of good X concludes with greater auction efficiency than its purchase of good Y, the opportunity for corruption (i.e. padding) is smaller in the purchase of good X than of good Y, all else equal. While the logic developed above is obviously clearest in a comparison between two otherwise identical auctions, the overall message is simple: *on average a bureaucrat at risk of engaging in corruption will prefer the agency pay more, rather than less, for any given purchase, as this increases the opportunity for corruption.*

In a real agency, unlike in these examples, all else is not equal, and many other factors likely affect auction efficiency. Unlike my two-purchase examples, however, real agencies purchase hundreds or thousands of goods across the 2011-2015 period covered in my data. Any persistent difference in the gap between initial and final prices across agency (or good, or procedure, or time) is less likely an artifact of these other factors the more purchases observed. Rather, if two agencies differ markedly in

auction efficiency across many purchases, and my assumptions and validation above hold, we can cautiously infer one generates greater corruption opportunities than the other. In empirical tests I operationalize auction efficiency as a measure of corruption opportunities by calculating the *median percentage difference between initial (maximum) and final (winning) price* across each agency's purchases.[50] I can also break this down further by class of good, procedure, and by month. As the final price is (almost) always lower than the initial price, when I present auction efficiency without rescaling it runs from 0 to -50 percentage points: -50 signifies the agency's median purchase on this measure concluded at *half* the initial price. 0 denotes an agency's median purchase on this measure concluded at the *same* price at which it started. If my logic above is sound, the latter agency presents greater opportunity for corruption.

Figure 3.2 plots this measure for each agency's purchases in Moscow, against its total spending (log scale) in Moscow, for the period 2011-2015.[51] An interesting pattern emerges from this graph. There is only a weak correlation between spending and auction efficiency: agencies spending relatively less (and making fewer purchases) demonstrate both high and low auction efficiency, whereas agencies spending a lot tend to experience low auction efficiency. If auction efficiency is a valid measure of corruption opportunities, this suggests that extensive spending on procurement is sufficient, but not necessary, to generate large-scale corruption opportunities. This is consistent with the predictions of my model: beyond a certain agency budget threshold, the senior and junior bureaucrat always agree to collude to maximize corruption

---

[50]I use the median because I am interested in inferences about agencies' behavior in a 'typical purchase'. Mean auction efficiency would effectively pool purchases that differ on other dimensions; interpretation of this quantity is unclear. Furthermore, for most agencies the distribution of auction efficiency is highly skewed, with most falling between 0 and -10 percentage points.

[51]The distribution of spending across agencies is highly skewed, even excluding agencies that made fewer than 100 purchases. In this period, the most profligate agency, the Ministry of Internal Affairs, spent 86,509,355,647 rubles through the procurement system within Moscow alone, equating to US$2,900,000,000 at the average exchange rate of the time. The lowest-spending federal ministry, the Ministry of Justice, spent only 881,896,426 rubles within Moscow, almost 100 times less. The smallest agency in the dataset, the "Moscow Center for Hygiene and Epidemiology, South-East Branch", spent 18,405,241 rubles, or 4,700 times less.
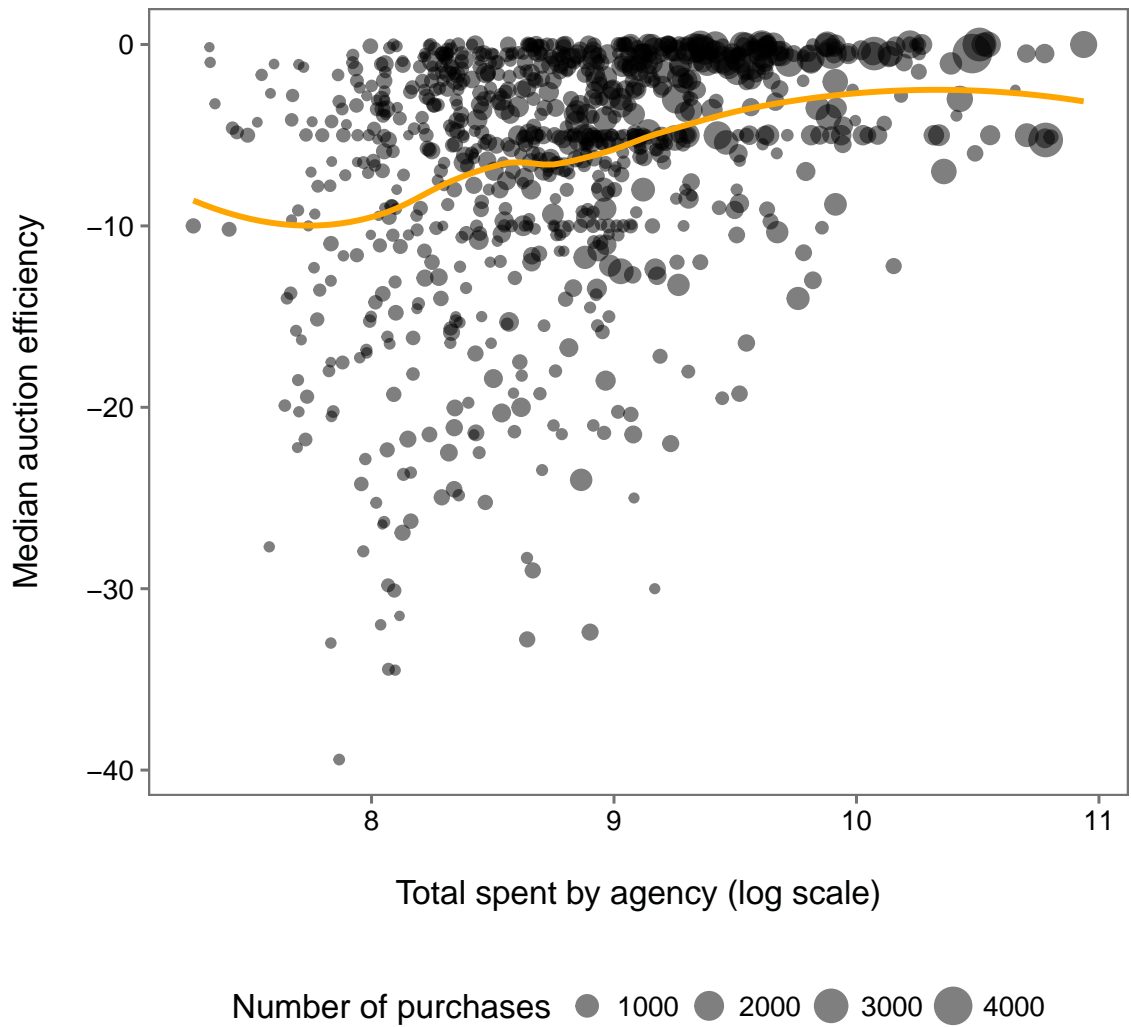
Figure 3.2: Corruption opportunities (as measured by auction efficiency), by agency. Orange line is a local smoother.

opportunities.[52] This pattern is also consistent with the way procurement corruption actually happens, through 'padding' the market price to create an extractable surplus. As discussed above, allowing the final (winning) price to drop considerably is counterproductive if one's goal is corruption. So far then, the model, measures and data agree in this graph.

There is, however a simpler explanation available for this pattern. Agencies that make more purchases (larger dots) and spend more (further right) might simply be better at choosing an initial (maximum) price, i.e one closer to the best price in the market. This is essentially a *state capacity* argument applied to procurement at the agency level; over time bureaucrats gain expertise through interactions with firms, affording them better information about the world outside the agency. This insulates the agency from external pressures and aids the transmission of politicians' preferences in to policy implementation. If this argument holds in the case of Russian public procurement, we should certainly expect the pattern observed in the top-right corner of Figure 3.2. These agencies would be staffed with experts making frequent purchases, accurately assessing the true market price for each good and avoiding attempts by suppliers to game the procedures. As a result, most of their auctions would finish close to the starting price, without generating opportunities for corruption. The problem for this argument is the mass of agencies in the top-left quadrant that spend far less, but still show low auction efficiency. A corollary of the 'capacity' explanation for the low auction efficiency of high-spending agencies, is that low-spending agencies should demonstrate much greater median auction efficiency. Bureaucrats in these agencies would set initial prices without expertise in the market; competition between suppliers should then cause prices to drop significantly in many auctions. While this outcome obtains for some lower-spending agencies (bottom-left quadrant), low auction efficiency is common across all levels of spending. If capacity

---

[52]Recall that this is just the first stage of their interaction, and does not imply corruption succeeds. Below this threshold, their decision to collude or not is conditional on other factors.

deadens opportunities for corruption, it cannot do so through the most plausible channel (frequent, larger purchases), which casts doubt on this alternative explanation as a whole.[53]

Another alternative to my interpretation of Figure 3.2 is that higher-spending agencies tend to purchase goods and services having fewer sellers. For example, there are fewer suppliers of fighter jets than there are of printers. The latter is a commodity good relevant to every agency and commercial business, while the former is highly specific to the mission of very few agencies. This would have two effects, and produce the relationship between auction efficiency and spending observed above: 1) goods purchased more rarely by agencies are likely to be more expensive simply for being 'exotic'; 2) fewer suppliers means weaker market pressure on prices. This argument suffers the same conceptual flaw, however, as the capacity argument. While it may account for the high-spending, low-efficiency agencies (if indeed they tend to purchase low-competition goods) it fails to explain why the low-spending agencies show both low and high auction efficiency. Furthermore, as Figure 3.3 shows, there is no relationship between spending and a measure of specificity of agency purchasing (a proxy for market competition), which I develop further below. If anything, lower-spending agencies are more likely to purchase 'unusual' baskets of goods than are higher-spending agencies. This runs counter to the argument that higher-spending agencies' low auction efficiency is driven by unusual purchases.

There is anecdotal evidence that the link I have drawn between auction efficiency and corruption opportunities exists in the minds of Russian bureaucrats. Ostrovnaya and Podkolzina (2015) examined drug procurement by St Petersburg's public health agency in 2008-10, a period when that agency was under investigation by the Federal

---

[53]Furthermore, in Russia the incentives to develop this kind of expert capacity are generally weak. As with any bureaucracy, agencies must spend their budget or risk losing it, and bureaucrats receive little formal training in the principles of efficient procurement. In this environment, starting the auction too low, such that many suppliers are put off entering, is risky as it requires starting the paperwork again if nobody bids.
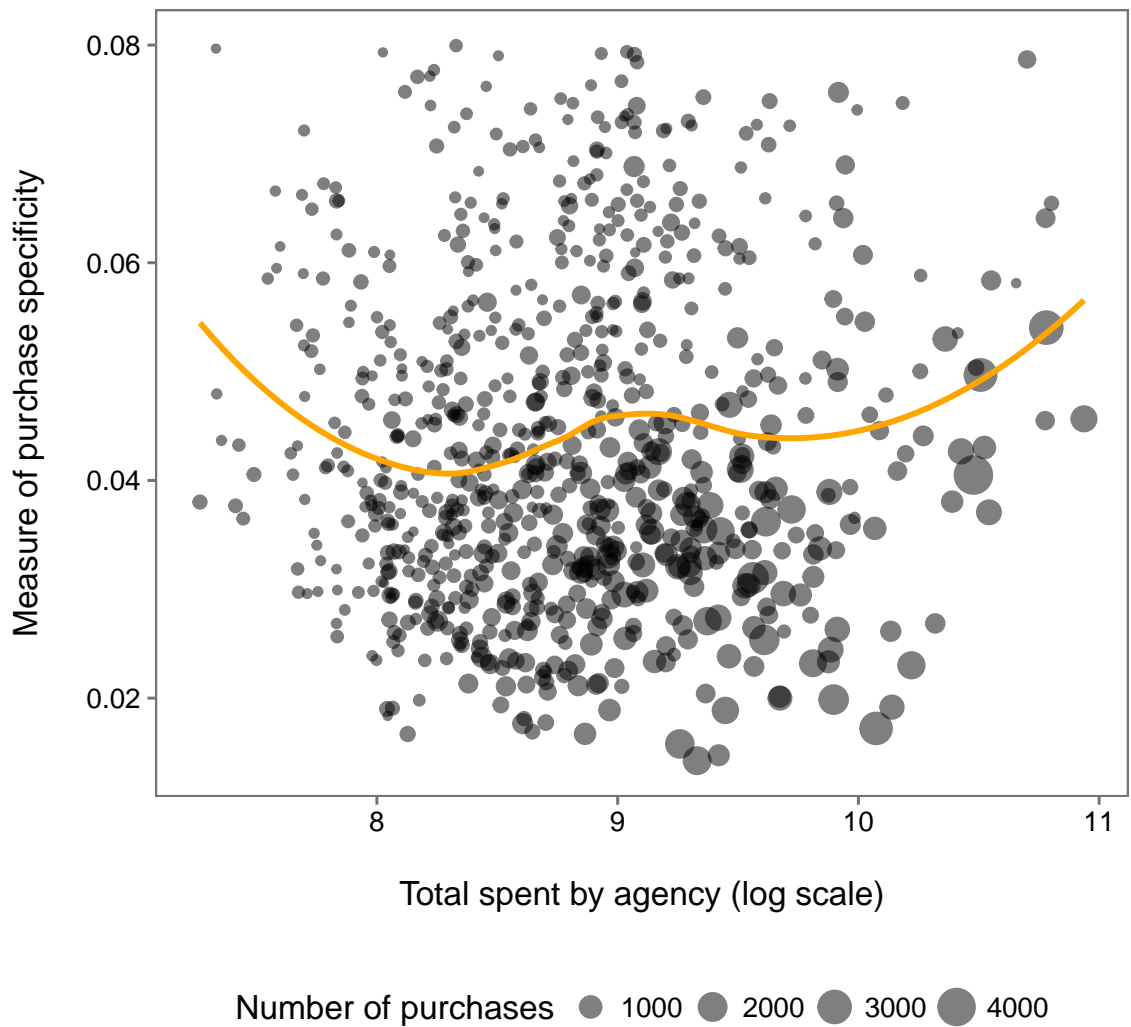
Figure 3.3: Purchase specificity (a measure of how unusual is an agency's procurement) by agency. Orange line is a local smoother.

Anti-Monopoly Service (FAS) for collusion with a so-called 'preferred supplier'. The alleged technique to generate a surplus (i.e. rent), to be shared by supplier and bureaucrats, was restriction of competition by disqualifying bidders who underbid the 'preferred' supplier. The local FAS took as evidence of law-breaking the fact most auctions were concluding with very few bidders, despite the agency purchasing generic drugs with many potential suppliers. After the FAS intervention, the agency's procurement involved more bidders, helping them avoid a court judgment; the new competitors, however, were fake. The number of bidders rose, without prices falling. This demonstrates the lengths to which bureaucrats will go to maintain the gap between initial and final prices, which manifests as low auction efficiency across many purchases.

One study, by Morozov and Podkolzina (2013), has found systematic evidence to support the case I have made above with respect to Russian procurement. The authors studied one market, highway construction, in one region of Russia, collecting very detailed data (from slightly before the period of my data) to study the effects of competition on auction efficiency.[54] They found that for purchases with auction efficiency between 0 and -10 (i.e. final price was 90-100% of initial price), indicators of competitive procurement, such as the number and prior experience of bidders, were not correlated with auction efficiency. For purchases where the final price was less than 90% of the initial price, indicators of market competition *were* correlated with the final price, suggesting two distinct populations of purchases: those subject to competition and those not.[55] In the former group the potential rent is dissipated by competition, which is consistent with my argument above: when we observe a large price change (high auction efficiency), that purchase is on average less likely to

---

[54]They call auction efficiency 'relative contract value', and construct it as the percentage of the initial price represented by the final price. Their approach is adapted from similar analysis of Ohio school milk auctions by Porter and Zona (1999).

[55]The 90% threshold in their paper is robust to manipulation, with the same pattern observed when their sample is split at 80%, and in to upper and lower quintiles.

generate rents that can be captured by suppliers or bureaucrats. Furthermore, as the market they study was also investigated by regional authorities, the authors compare their findings to a published schedule of meetings between alleged cartel members. Meetings tended to precede cases where the final (winning) price was identical or close to the initial (maximum) price; cases where auction efficiency was high were not preceded by meetings. This identifies how suppliers can benefit from low auction efficiency (the agency pays a higher price) and provides a mechanism (restriction of competition) that bureaucrats could attempt to sell.[56] It is important to note that low auction efficiency only creates opportunities for corruption: collusion between market participants can occur without any benefit to bureaucrats. We need other measures of bureaucrats' behavior to infer the extent of their involvement and benefit.

A final check of this measure of corruption opportunities is whether it is associated with any particular procedure. If so, it might indicate that auction efficiency is not a generalizable measure of corruption opportunities across all purchases, but instead a result of idiosyncratic differences in procedures. Figure 3.4, which plots hundreds of thousands of purchases in Moscow, shows that auction efficiency is not significantly different across procedure. The masses of the respective distributions are located near to each other, as indicated by the overlaid dots and the similar medians (solid horizontal lines across each box). It is worth nothing that *open electronic auction* procedure generates the most cases where the final price paid to the winning supplier is actually higher than the initial price set by the agency.

To conclude this subsection, there can be no 'smoking gun' evidence that auction efficiency is a valid measure of corruption opportunities, though the study by Morozov

---

[56]A concrete case illustrates how even a fake cartel helps maintain padding. In 2015 Novgorod's regional government announced a 'request for quotes' procedure to buy 300 copies of the book *Faces of Novgorod History*, for a maximum price of 225,000 rubles (US$3,500 at 2015 exchange rates). All 300 copies in existence were held by the publishing house *Veche*, which bid the lowest per-unit price by one ruble. The two competing bidders were later discovered to be shell companies owned by *Veche*, whose slightly higher bids were well-informed and gave the appearance of competition despite its absence. See: `http://transparency.org.ru/goszakupki/sgovor-na-torgakh-novgorodskoe-pravitelstvo-zakupaet-knigi-s-narusheniem-zakonodatelstva`

Auction efficiency by procedure,
for agencies in Moscow, 2011-2015
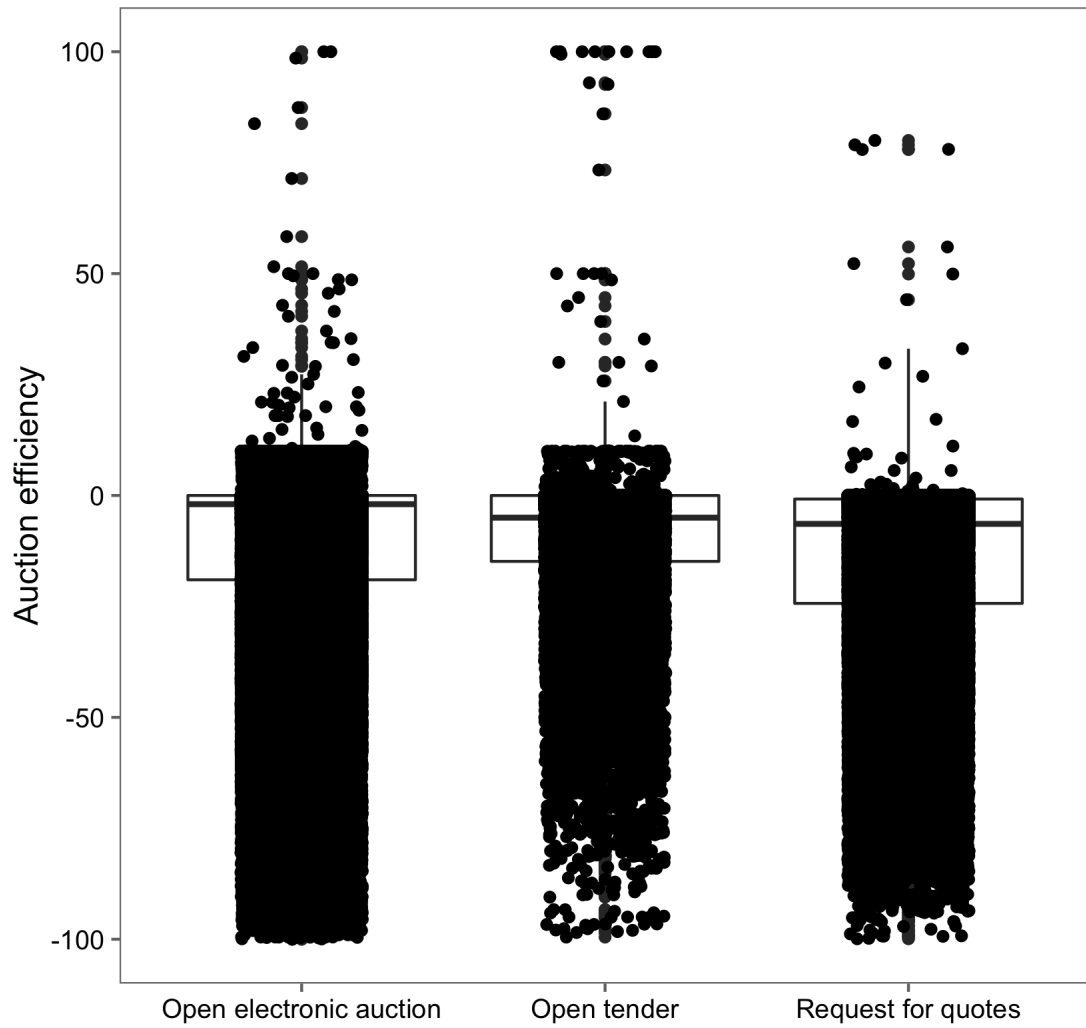
Figure 3.4: Box-plots of auction efficiency by the three choices of procedure. Distributions are similar, but note preponderance of cases where *open electronic auction* led to a price increase (efficiency greater than zero).

and Podkolzina (2013) discussed above is compelling. Instead, I have examined the most robust arguments that the patterns observed are driven by mechanisms unrelated to corruption, and found those explanations lacking. Furthermore, the measure is motivated directly by my theory: in the model 'padding' is necessary (but not sufficient) for actual corruption and auction efficiency is the best available measure of this padding. Across the logical cases describing the relationship between auction efficiency and opportunities for corruption, presented in Table 3.3, auction efficiency is more likely to be lower than higher the more bureaucrats engage in 'padding'. In addition to validating this measure as best as is possible, I also uncovered evidence that the total spent by an agency is an important variable in agency-level empirical tests, as predicted by my model. Now I proceed to measure corruption itself.

3.3.2 Measuring corruption itself with procedural 'red flags'

The previous subsection developed an approach to measuring the opportunities for corruption generated by agencies spending their procurement budget. Though auction efficiency was unsuitable as a proxy for actual corruption, I adapted it to infer the extent of 'padding' likely to have occurred across many purchases by the same agency. The creation of opportunities, however, is insufficient for bureaucrats to benefit from corruption. Having created a rent by committing the agency to pay more than the lowest market price for a good or service, they need a technology for extracting this rent for their benefit. Otherwise, it is completely captured by the winning supplier, who bore far less risk in its creation than the enterprising bureaucrat.

As discussed above, corruption is not directly observable at the level of an entire agency.[57] Instead, I must use the data available to generate proxy measures that are *more likely to increase than decrease* when underlying corruption increases. Here I rely on prior work measuring procurement corruption in other states, in particu-

---

[57]Reports of corruption by third parties can be motivated by many other factors, making them unreliable measures of underlying corruption.

lar Fazekas, Tóth and King (2016),[58] as well as practitioner-focused publications by OECD (2016), The World Bank (2010), and Kramer (2012). The former identify 'red flags' in Hungarian procurement that correlate with other measures of corruption; the latter describe generalizable features of procurement likely to be correlated with corruption. While electronic procurement systems are different in their details, they all follow the approximate sequence described in Table 3.2 above: agencies notify suppliers, who submit bids, and one bidder wins.[59] This prior work is consistent with my overall measurement approach, which derives conditions that should obtain if purchases are corruption-free; robust exceptions to these conditions indicate corruption is more likely to have occurred in a given agency. I call these exceptions *procedural red flags* and identify and validate them in the Russian procurement data below, restricting my attention to the 777 agencies that made at least 100 purchases in Moscow. Restriction to the largest procurement market reduces the probability that idiosyncratic effects of economic geography drive results, while ignoring smaller agencies results in measures less affected by any unusual purchases. The measures are summarized in Table 3.4 and described below.

Red flag 1: no price decrease. Perhaps the most intuitive 'red flag' measures result from procurement behavior and outcomes that simply should not be observed under the prevailing rules. For example, regardless of the specific procedure chosen by agencies in Russia, all are *reverse* auctions: only in very exceptional circumstances should the final (winning) price be higher than the initial (maximum) price. Furthermore, if the rules and market function as designed it should be rare for an auction to finish at exactly its starting price. This would only be expected when agencies buy highly unusual goods and services, which happens infrequently; the overwhelming ma-

---

[58]The cited article captures the results of a multi-year research agenda at the *Corruption Research Center Budapest*, the publications of which are available at http://www.crcb.eu/?cat=7

[59]Corruption could happen at any stage: bureaucrats can perturb the notification process, and influence who will bid, and then manipulate the process by which one bid is selected.

Measure	Definition	Rationale
No price decrease	Proportion of agency purchases where final (winning) price higher than initial (maximum)	Increases should never occur under the rules as all bids should be lower than initial (maximum) price
Single-supplier procedure	Proportion of agency purchases using single-supplier procedure	This represents use of the highest-discretion 'single supplier' procedure, under which the agency informs a contract has already been signed
Bunched initial prices	Proportion of agency purchases where initial price within 1% of rule thresholds	Prices should not be responsive to thresholds if bureaucrats not concerned with effect of rules on ability to extract
Dramatic price decrease	Proportion of agency purchases where final price is less than half initial price	Only possible when winning supplier has drastically lower costs or plans to provide a cheaper substitute
Winner concentration	Herfindahl-Hirschman index of purchase values by agency-supplier pair	Buying from fewer and/or larger suppliers helps build long-term collusive relationships that aid corruption

Table 3.4: Summary of corruption proxy measures

jority of purchases through this system should have multiple suppliers in Moscow.[60]

Violating this expected pattern (price decreases in most auctions) is the easiest way for bureaucrats to capture the rent generated by any procurement purchase. When we observed low auction efficiency (i.e. small price changes) above, we could not be sure who exactly benefited: bureaucrat or supplier. All we knew was that on average, opportunities for corruption are greater when auction efficiency is lower; the rent generated might be entirely captured by suppliers manipulating ill-informed bureaucrats. When, on the other hand, a procedure ends with the agency paying a *higher* price than the maximum notified, or *exactly the same* as the maximum, it is much more likely to be the result of foul play that must involve and benefit the bureaucrat. In the Russian case, as elsewhere, this is most easily achieved by disqualifying all but one bidder, who then automatically wins the auction even if they bid higher

---

[60]In Russia's remote regions, of course, there may be fewer suppliers due to market size. I address this comprehensively in the more detailed tests in Chapter 4.

than (or exactly) the initial (maximum) price.[61] A price increase can also be achieved by the agency simply ignoring the rules of the system and signing a contract for a price higher than the final auction price, though this is much more easily detected by other bidders, who are likely to complain. While I cannot observe the reason for a price increase (or no price change) in the Russian data, the logic just described leaves only one alternative explanation not involving foul play by bureaucrats: the agency set the initial price poorly. According to the law, however, in these circumstances the procedure should be canceled and re-run. For these reasons, *no price decrease* is the first of the 'red flag' measures I employ as a proxy for underlying corruption, at both the individual purchase level (Y/N) and at the agency level (proportion of all purchases). Figure 3.5 shows this red flag is only prevalent in agencies where corruption opportunities are extensive, as measured by auction efficiency.[62] Though these two measures are correlated by construction, as predicted by my theory, corruption opportunities are not sufficient for corruption. Many agencies with relatively poor auction efficiency (i.e. greater corruption opportunities) score low on this red flag.

Red flag 2: single-supplier. Obviously, one such flag is insufficient evidence for the link between opportunities and corruption. Another avenue for bureaucrats to take advantage of corruption opportunities is to select the highest-discretion procedure, *single supplier*, in which they simply notify the market of a signed contract without following a competitive process.[63] While there are justifications in Russian law, and in general procurement practice for this procedure,[64] it is widely recognized

---

[61]This generates an initial appearance of competition by the entry of many bidders, even though the outcome is rigged. A clear example of this comes from the 2015 investigation by the Federal Anti-Monopoly Service (FAS) of the road construction company Avtodor's contracts to build the Moscow Central Ring Road. The deputy head of the FAS alleged Avtodor reached "anti-competitive agreements" that allowed it to propose higher prices than other bidders but win on dubious non-price criteria like "use of nanotechnology". Source:https://themoscowtimes.com/articles/russian-road-agency-investigated-for-holding-suspicious-tenders-47930

[62]Rescaled here so 1 represents the agency with the lowest median auction efficiency.

[63]For this measure I restrict attention to purchases greater than 100,000 rubles, as below this there is no obligation to run a competitive procedure.

[64]If there is only one supplier of a good or service, and the agency is certain of this, it can be more efficient to proceed directly to negotiate directly with that supplier.

Proportion of purchases with no price decrease vs corruption opportunities,
by agencies in Moscow, 2011–2015
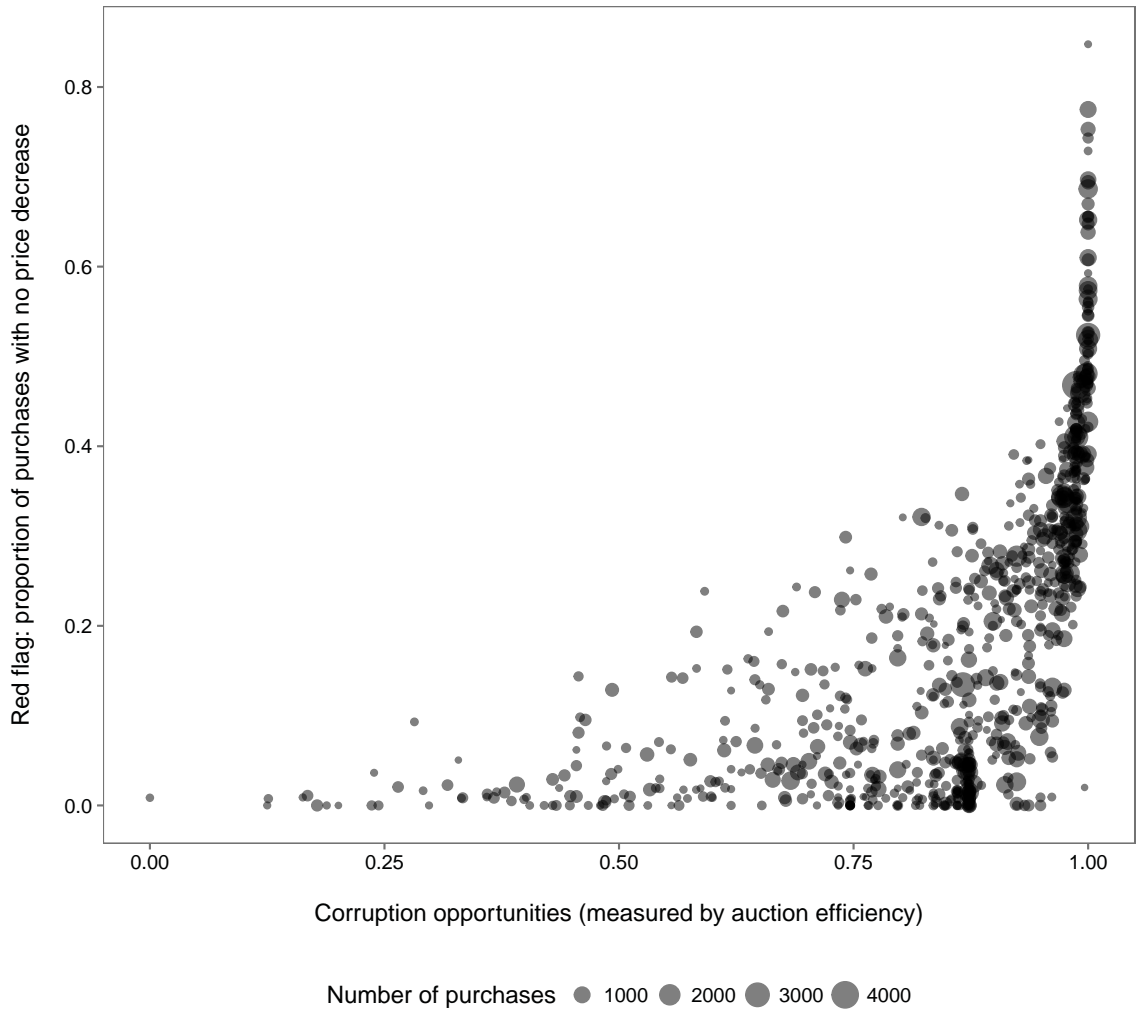
Figure 3.5: Corruption opportunities (i.e. low auction efficiency) are necessary but
not sufficient for a high proportion of an agency's purchases to display
the procedural 'red flag' of no price decrease.

by scholars of procurement to encourage corruption.[65] Only when internal agency checks and balances are robust, and accountability of politicians for those checks is high, should we expect the efficiency benefits to outweigh the costs borne through corruption. Based on this logic, the *proportion of agency purchases using the single-supplier procedure* is the second of my 'red flag' measures. Its empirical distribution is similar to the measure above, but less striking. Several agencies in Moscow run 7-10% of their purchases under this highest-discretion procedure, and all are above the average on my measure of corruption opportunities.

Red flag 3: bunching at thresholds. In Figure 3.1 above I used the 'bunching' of initial (maximum) prices near certain rule-based thresholds to demonstrate these data show the response of bureaucrats to constraints of the procurement system. Here I extend that logic to generate a measure correlated with underlying corruption. If bureaucrats are unconcerned with the effects of these thresholds on their ability to steer outcomes, then we should not observe such sharp discontinuities. Put differently, if initial prices are set only based on knowledge of the market, however incomplete, the frequency of purchases should decrease monotonically as the price increases. Instead, when pooling all agencies together in Figure 3.1, we see patterns that defy explanation under the logic of corruption-free procurement. Bureaucrats seem to be going out of their way to list purchases just below an initial price of 3,000,000 rubles in the case of *open electronic auctions*, and just under 500,000 rubles when using *request for quotes*. This is often a result of so-called 'contract splitting', in which a larger purchase is split in to several smaller purchases, to evade oversight.[66] As with all my proxy measures of corruption, I do not claim that in every single case this implies the bureaucrat derives some corrupt benefit. Yet the preponderance of initial prices that just 'beat' these thresholds (and the reduced discretion they represent) simply cannot

---

[65]In the Russian case, see Balsevich and Podkolzina (2014). In the Hungarian case see Fazekas, Tóth and King (2013) for a detailed overview of this and other techniques in the 'toolkit'

[66]Described in The World Bank (2010).

Proportion of purchases near thresholds vs corruption opportunites,
by agencies in Moscow, 2011−2015
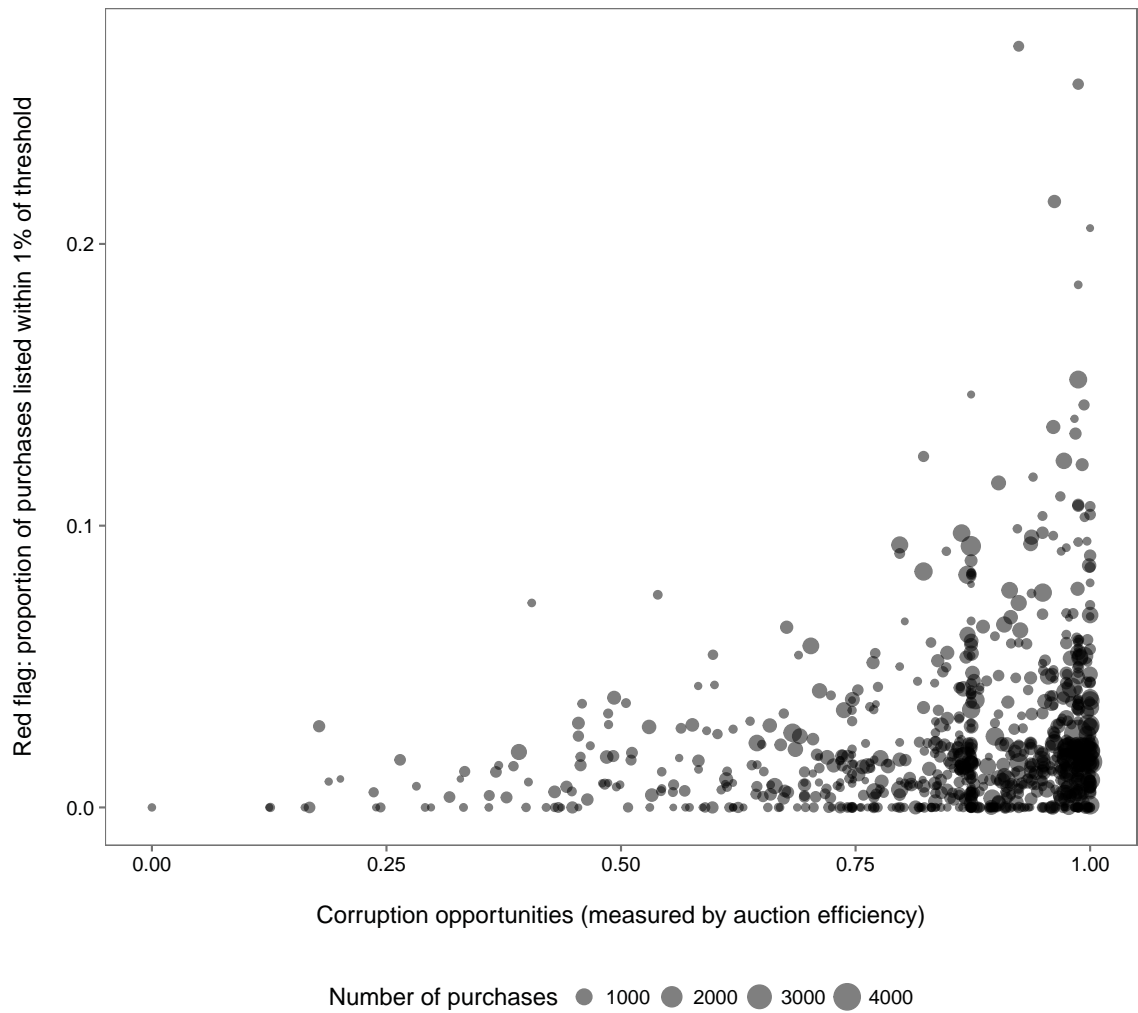
Figure 3.6: Corruption opportunities (i.e. low auction efficiency) are necessary but not sufficient for a higher proportion of an agency's purchases to be 'bunched' near thresholds where procurement rules change.

be accidental. The observed pattern is the result of intentional action by bureaucrats, and there is no plausible explanation of the motivation behind this action, other than to extract a portion of the rent created through procurement.[67] To measure the extent of an agency's exception to the corruption-free case, I calculate the proportion of its purchases in which the initial (maximum) price is within 1% of a threshold. Figure 3.6 shows the same pattern for this *bunching initial prices* red flag measure that we saw for the *no price decrease* measure above. Opportunities are necessary, but not sufficient, to score highly.

Red flag 4: dramatic price decrease. When calculating measures of auction efficiency, I excluded purchases where the price changed (between initial notification and winning bid) by more than 100%. I did not want extreme values to overly influence my characterization of each agency's typical purchase. This price change is useful not only to measure corruption opportunities through auction efficiency, as above, but can also proxy for actual corruption under specific conditions. According to many in the public procurement literature, the 'ideal' price change over a reverse auction is around -10%, so that the final (winning) price is roughly 90% of the initial (maximum) price.[68] When price changes are closer to zero, the most likely explanation is 'padding', which in turn creates opportunities for corruption as discussed above.[69] When instead the price decreases systematically by *too much*, different concerns emerge. Either the procuring agency was hopelessly misinformed about the state of the market, or its bureaucrats colluded with a supplier to win the auction but fulfill the contract with a drastically cheaper alternative. The former explanation

---

[67]If bureaucrats were clueless and innocent across the board, the distribution of prices around these thresholds would not be discontinuous. If bureaucrats were unfailingly competent social welfare maximizers, the same would be true.

[68]This is the threshold chosen by Morozov and Podkolzina (2013), based on an empirical finding in Ishii (2009) that cartel bidding in Japan rarely brought the price down by more than 5-10%. The exact number is less important than the idea that efficient auctions conclude at some level sufficiently below their initial price, such that padding is unlikely to be a problem.

[69]A highly expert agency might also achieve this outcome, but this relies on its bureaucrats having better knowledge of the market than bidders.

becomes less plausible over the course of many purchases.

Imagine an agency wishes to purchase one US ton of refined sugar, which has a commodity market price of $38,000.[70] If the agency lists this purchase with an initial (maximum) price of around $42,000, and suppliers bid down to the market price, auction efficiency for the purchase is approximately 10%. If the winner bids $34,000 then perhaps it is simply disposing of surplus stock at a lower profit margin. If, however, a supplier bids $21,000 and wins, something is clearly amiss: it should not be possible to provide this good at that price. If the procedure concludes and a contract is signed, it is highly unlikely the agency will receive exactly what it wished to purchase. Instead, the contract will go unfulfilled, or the agency will receive a substitute that meets at least some of the requirements specified in the procurement notice, but which can be provided at much lower cost. This logic is more clearly explained in the case of commodity goods, where quality standards are easier to define, but the threat of such behavior is actually greater the less commoditized the good or service: it gives suppliers more discretion to work around specifications.

This technique is known as product substitution, and the red flag indicator for it is a *dramatic price decrease*, which I define as a final (winning) price 50% or less than the initial (maximum) price.[71] This means the agency only pays half of what it expected, which more likely than not means it will not receive exactly what it expected. The link to corruption is that this should happen infrequently in those agencies where bureaucrats are closely monitoring the quality of purchased goods and services, rather than extracting rents. After being stung once they would learn the warning signs and tighten their specification. Conversely, if dramatic price decreases are common in an agency, bureaucrats running procurement are less likely to be concerned about

---

[70]At a price per pound of sugar (July 2016) of $19.05. Source: `http://www.indexmundi.com/commodities/?commodity=sugar`

[71]When calculating the *dramatic price decrease* red flag measure, I restrict attention to the *open electronic auction* procedure, where the logic outlined here applies cleanly. Bidders observe each others' bids so in the corruption-free case are more likely to offer an 'honest' price incrementally lower than the previous bid, rather than drastically lower in the case of likely corruption.

Proportion of purchases with dramatic price decrease vs corruption opportunities,
by agencies in Moscow, 2011–2015
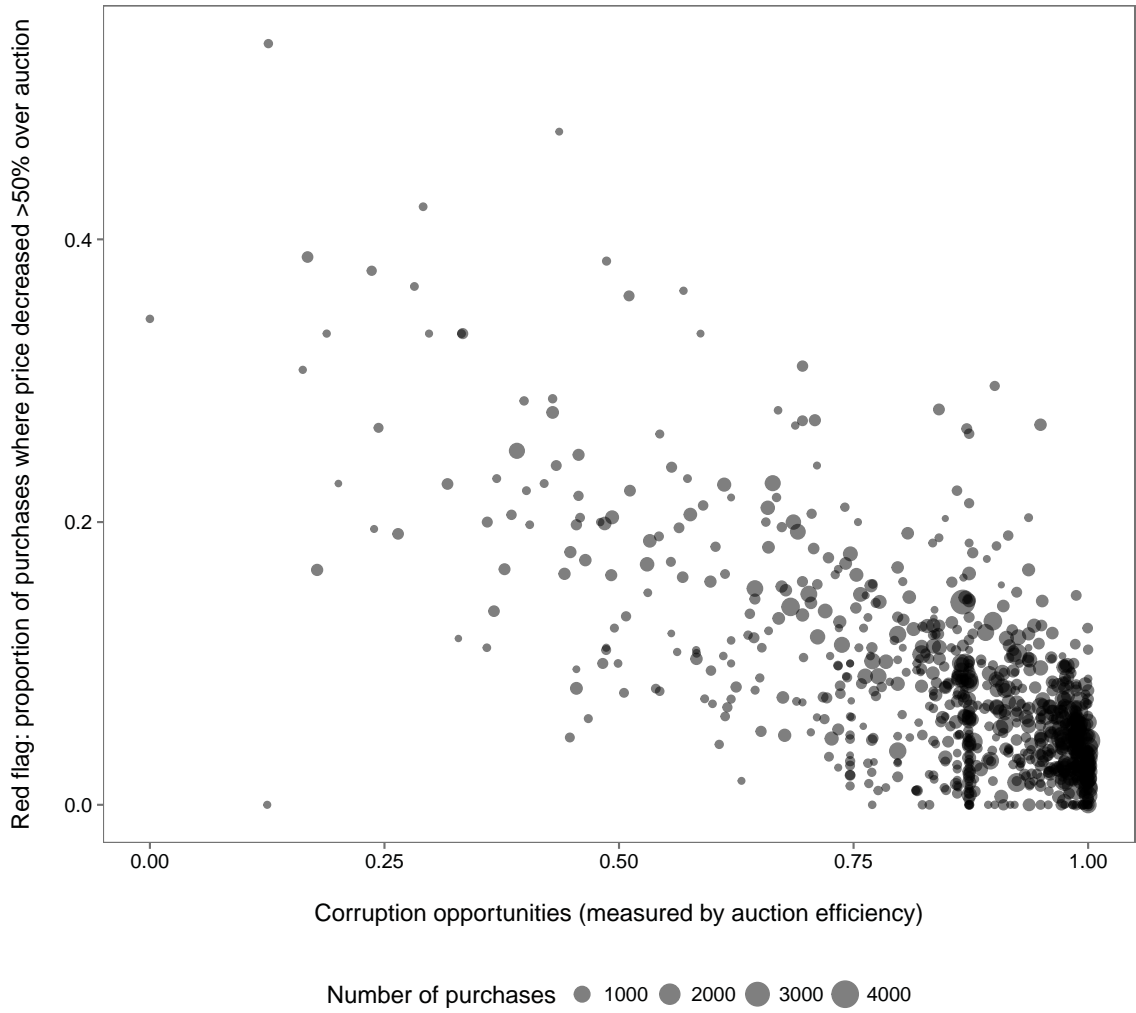
Figure 3.7: Despite the negative correlation, dramatic price decreases are still preva-
lent (up to 10%) when corruption opportunities are large.

paying for goods of dubious quality, or services not delivered at all. Over relatively few auctions this could be the result of incompetence or a poor understanding of the market; over many auctions it is unlikely to be possible without corruption, i.e. the bureaucrat's payment for colluding with the winner to 'look the other way'.

The distribution of this measure relative to corruption opportunities (Figure 3.7) differs from the measures above, for an obvious reason: they both rely on calculating the difference between initial (maximum) and final (winning) prices.[72] Nevertheless, the correlation is far from tight, as even the agencies where corruption opportunities are higher (i.e. lower median auction efficiency) vary in the proportion of purchases where the price decreases dramatically. Some experienced no dramatic price decrease across thousands of purchases; for others this occurred in 10% of purchases.

Red flag 5: winner concentration. The final red flag measure I develop here also correlates with collusion between bureaucrats organizing procurement and suppliers bidding. In a well-functioning procurement market without systematic corruption, and especially in a large market like Moscow, we would expect agencies to source their goods and services from many different suppliers.[73] This is simply a consequence of scale of competition: for the vast majority of products, no single firm will offer the best bid in all cases. If instead an agency purchases predominantly from few suppliers, this markedly increases the likelihood of collusion between supplier and bureaucrat, which in turn makes it more likely the bureaucrat derives illicit benefit from procurement decisions. The logic here is borrowed from analysis of competition in markets, particularly for antitrust enforcement.[74] Market concentration is correlated with corruption through two complementary mechanisms: a market with fewer players is more susceptible to cartel behavior; fewer competitors makes it easier for

---

[72]Corruption opportunities are measured with median auction efficiency, whereas this red flag is the proportion of purchases below a certain threshold (-50%) of auction efficiency.

[73]This may not be the case in more remote regions of Russia, where more unusual goods and services may have relatively few suppliers.

[74]This logic, and its measurement, is outlined in Kelly (1981).

bureaucrats to identify the most lucrative targets for collusion.[75]

Political science has already adapted the *Herfindahl-Hirschman Index* (HHI), used to measure market concentration/competition, to measure political competition.[76] I use the original formula, which computes a measure between 0 and 1 for each agency, based on the sum of the squared shares of its spending with each of its suppliers.[77] Higher values indicate fewer, larger suppliers receive most of the agency's procurement budget, lower values indicate the budget is spread among fewer, smaller suppliers. According to the logic outlined above, when underlying corruption (driven by collusion) increases, the HHI is more likely to increase than decrease.[78]

The distribution of this *winner concentration* measure vs corruption opportunities is similar to what we observed for the *no price decrease* and *bunching at thresholds* red flag measures developed above. A high score on the corruption opportunities measure is necessary, but not sufficient, for an agency to score highly on this red flag measure. A concern with this measure emerges, however, from the interaction between the market concentration of suppliers and the differing needs of agencies. Despite the fact all agencies engage in procurement, their different missions are still reflected in the basket of goods and services purchased. If agencies buying more unusual products tend to score higher on this concentration measure, I have simply revealed the obvious: there are fewer suppliers of products for which there is lower demand. Figure 3.8 shows that this is not a significant danger, though there is less

---

[75]They are complementary because it is easier for a bureaucrat to collude with one cartel, rather than with many individual suppliers; it is easier to form and maintain a cartel if decisions by the bureaucrat on the other side of the transaction can be purchased without facing competitors.

[76]The 'effective number of parties' introduced by Laakso and Taagepera (1979) is an inverse HHI.

[77]Formally, for agency a, its suppliers n through N and their shares $s \in (0, 1)$ of its spending:

$$HHI_a = \sum_{n=1}^{N} s_n^2$$

[78]As should be expected given the prevalence of corruption in Russia, survey data show that firms believe bribing is important to winning government business (Bashina, 2013). Unfortunately there is no comparable survey covering bureaucrats' views on the necessity of taking bribes.

Concentration of spending on winning suppliers vs purchase specificity,
by agencies in Moscow, 2011–2015
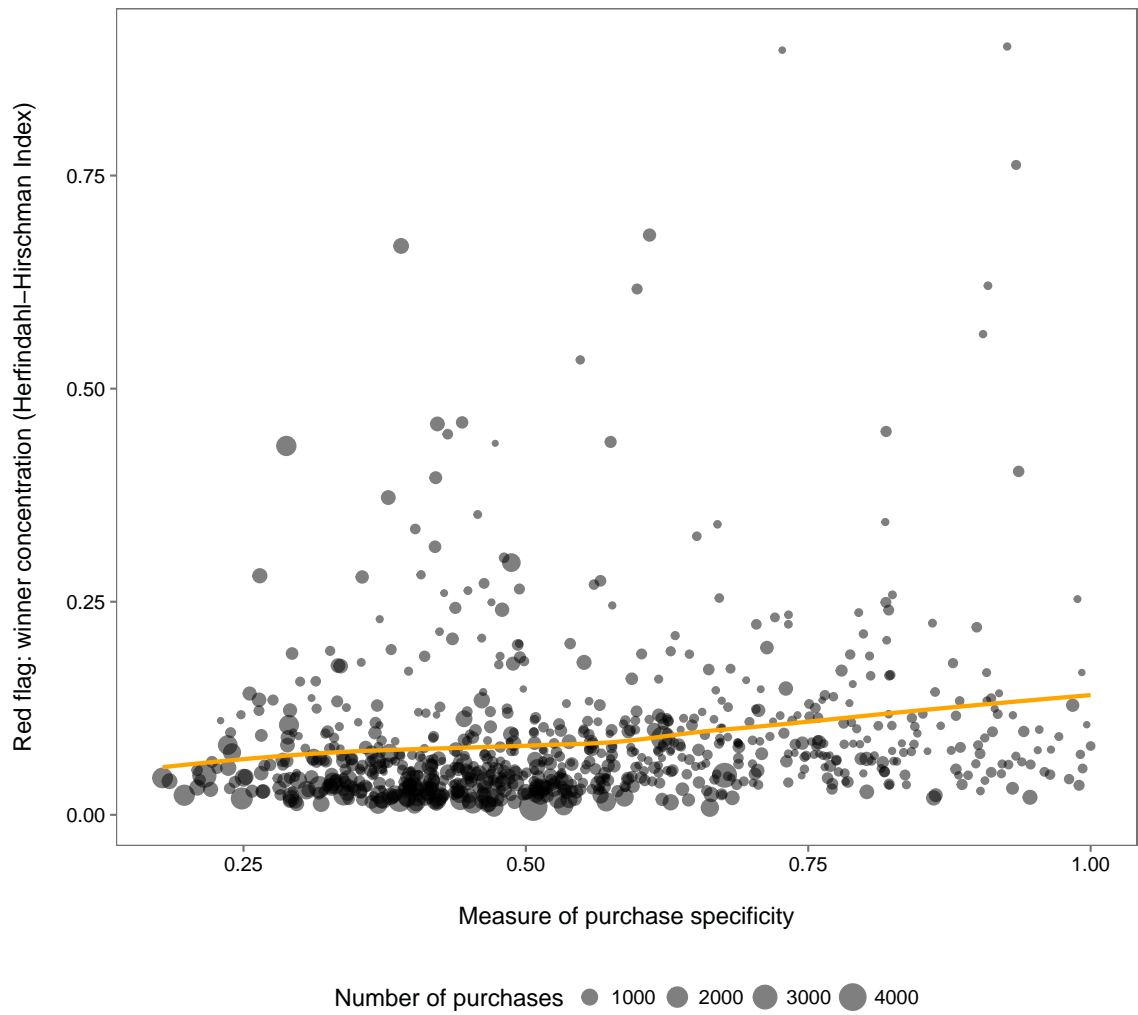
Figure 3.8: Concentration of each agency's spending on its winning suppliers is only
weakly correlated with the 'unusualness' of products it purchases.

variation in the HHI than the other red flag measures.

3.3.3 Comparison with existing measures

As discussed in the literature review, there have been few attempts to measure agency-level corruption in any state, for two reasons. First, corruption has been predominantly studied through the paradigm of bribing a bureaucrat for a permit, which makes agencies incomparable given their differing missions. Second, data on the scale I present here have only recently become available for the one activity all agencies have in common: procurement.

I searched for rankings of corruption in Russian agencies, and found only one that is suitable to compare with my measure.[79] The *National Rating of Procurement Transparency* (NRPZ)[80] is conducted annually using the same data I downloaded and parsed, but differs in several important respects. First, its focus is on transparency, which is related but separate to corruption, and is a much broader concept. This dissertation explains why bureaucrats in some agencies, but not others, systematically engage in less corruption than would be expected given the prevalence of corruption in their state. Increased transparency may make corruption more difficult, but my theory does not rely on it for explanatory power, nor make predictions about its effect. Second, the NRPZ methodology is aggressive where mine is cautious: it treats many types of behavior as indicative of transparency (or lack thereof), even when these have robust alternative explanations. For example, one of its criteria is *proportion of purchases conducted through open electronic auction*, with the justification that this procedure is believed to offer the greatest level of competition and transparency. This assertion is not tested, despite its sensitivity to the mix and value of products purchased by an agency.[81] In another example, the NRPZ methodology calculates

---

[79]The remainder are ad-hoc surveys of experts, usually focused on a subset of agencies, and most do not publish any methodology.

[80]Source: `http://nrpz.ru/`

[81]Recall that this procedure is mandatory for certain commodity goods.

mean auction efficiency by procedure, but states that both too high and too low values are correlated with corruption, without separating out the two mechanisms I described above. This means an agency actively working to both maximize padding, and encourage product substitution by suppliers, would appear as a model of clean procurement. Finally, the methodology includes the proportion of contracts awarded to small businesses, as a measure of the contribution of public procurement to the development of small business. I was also interested in a measure correlated with firm size, but rather as a measure of the degree of over-reliance on relatively few suppliers; the link between my logic and corruption is clearer than the link between small business development and transparency.

As a result of these differences in approach, we should not expect the NRPZ and my measures to be tightly correlated. Given that transparency is often viewed as an antidote to corruption, it would be useful validation if the measures agreed on the 'leaders' and 'laggards' among the agencies in Moscow. In the cases where an agency in the NRPZ matches an agency in my dataset,[82] I assign it to "above median" or "below median" on each measure. Only in one case do they agree strongly: *no price decrease* and the NRPZ transparency score agree on whether an agency is above or below their respective median in 72% of cases. The NRPZ and the *winner concentration* agree with above/below their respective median only half the time, while agreement with the remaining measures is around 35-40% only. It is difficult to draw any firm conclusion from these results. On one hand, it is comforting to see at least some level of agreement, rather than total disagreement. On the other hand, the underlying concepts and their implementation are quite different.

---

[82]Recall that for agency-level measures in this chapter I restrict attention to those with at least 100 purchases. I relax this constraint in my purchase-level tests in Chapter 4.

3.3.4 Identifying the units of analysis: agencies in procurement

My theory explains variation in agency corruption as a function of information asymmetry between bureaucrats, and I use a game between a senior and junior bureaucrat to develop its testable implications. As discussed previously these implications apply to senior-junior relationships at any level in a bureaucratic hierarchy. I can make this claim because my theory is built around the most important decisions over corruption in an agency: whether to engage in it, *and* how to manage its risks.[83] These decisions are ultimately made by individuals, in the context of a relationship with their immediate supervisor, who faces a similar decision with respect to their supervisor. What changes these decisions is the degree to which junior bureaucrats hold better information about the opportunities for corruption than senior bureaucrats.

The next subsection outlines how I operationalize information asymmetry in the Russian procurement data, as *purchase specificity*. This requires I first identify agencies in the data, and link these empirical units of analysis to the stylized agency in my theory. This is facilitated by the system that generates the data, which covers all procurement by all agencies (from childrens' hospital to Federal Ministry) from 2011-2015. In the system a *purchasing unit* links a set of purchases to a particular organizational identity.[84] While one purchasing unit is named *Ministry of Defense* this does not mean it is the sole buyer of goods and services for that entire Ministry. There are also purchasing units in the data named *Central directorate of medical procurement, Ministry of Defense* and *Ministry of Defense, Moscow Barracks No.6.* Ideally, every constituent purchasing unit of each top-level agency would be linked, allowing me to characterize corruption in the entire Ministry of Defense vs. the entire Ministry of Finance. Unfortunately, such a systematic hierarchy is not available in

---

[83]In the previous chapter I identified the missing link between the corruption and bureaucracy literatures as a failure to recognize the fragility of corruption. Even in corrupt states, corruption is risky for bureaucrats, especially for the senior bureaucrat who must authorize or forbid it.

[84]This identity is linked to a secure USB key that verifies the right of the operator to run purchases for a given unit.

these data. There are more than 4,000 purchasing units in Moscow alone and 110,000 across all regions, which negates the possibility of linking these manually.[85]

While this prevents me from calculating a comprehensive score for each top-level agency, this limitation has its advantages. First of all, collapsing so many diverse purchasing units in to one observation necessarily discards useful information about the link between purchase specificity and corruption. This could mean top-level agencies look little different based on these highly aggregated purchasing profiles, and that variation in corruption is so attenuated, that we learn nothing about my explanation or the alternatives. The second advantage of treating each individual purchasing unit as a separate agency is a much closer link between my theory and these data. The work in my theory is done by the accountability and delegation relationship between senior and junior bureaucrat, which operates as cleanly on a purchasing unit as it does on the theoretical agency; the detailed logic of my theory should apply regardless of the position of each purchasing unit within its respective top-level agency.[86] Finally, treating each purchasing unit as equivalent to an agency within the theory is less problematic given we can observe the same details for all of them: what they buy, how unusual or specific that purchase is, and how the auction proceeds.

3.3.5 Operationalizing information asymmetry

My explanation for why corruption is different across agencies is founded on the information available to a senior and junior bureaucrat while making procurement decisions. As in any principal-agent model, the senior bureaucrat can delegate to the junior when the latter has better information; this is a sustainable equilibrium when they both are better off having exchanged information for control. My theory

---

[85]This might be achieved programmatically, which I leave to future work.

[86]While the senior-junior relationships I describe can exist at any level of the bureaucratic hierarchy, I have no theory for how these levels interact, only that they follow the same general pattern. In this sense, leaving the many separate purchasing units of the Ministry of Defense disaggregated improves observability of this core relationship across the many levels of a top-level agency, which would be lost if the agency were treated as a homogeneous unit.

turns this standard treatment sideways to generate new implications, by requiring the bureaucrats to engage in risky collusion to earn more than their licit wage.

There are two risks in the model should the senior bureaucrat over-delegate. First and foremost, the senior might cede so much control in pursuit of successful collusion that the junior bureaucrat extracts substantial corrupt rents during procurement but hides the extent of this benefit, cheating the senior from a 'fair share'. This conflict takes place in the shadow of the second risk: if the junior bureaucrat extracts more than what the minister-in-charge will tolerate, both will be punished. The central tension is therefore between the senior's need to delegate to earn more, and the fear that doing so will generate risk without reward, and the optimal choice depends on the information advantage of the junior over senior bureaucrat. Up to a point, this advantage is profitable for both; past that point it creates more risk for the senior than the potential reward of having a subordinate expert in corruption. Unexpectedly clean agencies are thus the outcome of failed collusion, driven by the junior bureaucrat knowing more about the possibilities for extraction than the senior can accept.

I capture this information asymmetry by adapting the concept of asset specificity[87] in to a measure of *purchase specificity*. At the level of a single purchase, the most specific purchase is when an agency buys a good or service that is in no other agency's 'basket' of purchases; the least specific purchase is when the good or service is in every agency's 'basket'. Only one agency is likely to procure fighter jets (though not through this procurement system) which makes that purchase highly specific, whereas almost every agency will buy transportation services, making that purchase highly generic. Aggregating many purchases together, an agency's *purchase specificity* is higher the further its procurement spending across all goods and services deviates from the spending profile of the 'mean' agency. The mean agency's spending per product group is simply the proportion of all agencies' spending falling in the 61

---

[87]For an explanation and application of asset specificity see Clark and McGirr (2010).

product high-level product groupings in the dataset.[88] For each individual agency, I then compare the proportion of its budget spent on each product group to the mean proportion; agencies that deviate more across more product groups receive a higher purchase specificity score. For example, if there were only three products (A, B, C) that agencies could procure, and across all agencies the proportion spent on these was the vector (0.5, 0.3, 0.2), then an agency spending (0.48, 0.31, 0.26) has lower purchase specificity than an agency spending (0.2, 0.7, 0.1).[89]

The immediate implication of variation in purchase specificity is the number of potential suppliers: there will naturally be fewer for more specific goods and services. This is why I took pains above to check my measure of information asymmetry between senior and junior bureaucrat was not correlated at the agency level with my measure of *winner concentration* (Figure 3.8) or with total agency spending (Figure 3.3). Otherwise, I would risk measuring the same underlying concept on both sides of the equation, or confusing purchase specificity with agency capacity. Another approach altogether would be to define *a priori* more or less unusual products, but I have no grounds on which to make that determination. Instead I choose to generate a theoretically-informed measure directly from the data in each region of Russia; first in Moscow to validate my approach, and then separately across the other regions so the spending profile of the mean agency can change with respect to local conditions.

Having guarded against these risks, I can make the following link from purchase specificity back to the decisions described in my model. When an agency routinely buys more unusual products, as measured by purchase specificity, it is more likely to require specialist knowledge to successfully extract rents through corruption. In any agency, more specialized knowledge is held at lower levels rather than higher

---

[88]There is also a lower-level product coding available with 12,048 categories. This is too many to form the basis of a useful measure, as many have very little spending against them across all agencies. Aggregating these 12,408 products to 61 product groups most likely works against my theory by attenuating variation in my key independent variable.

[89]I calculate absolute deviations so that over- and under-spending in a product group relative to the mean agency do not cancel each other out.

levels of the hierarchy. This implies that in agencies where purchase specificity is higher, the junior bureaucrat has a greater informational advantage over the senior bureaucrat; in these agencies corruption is at risk of failing under the weight of the twin risks outlined above. When purchase specificity is low, little specialist knowledge is required, weakening the informational advantage and with it the senior's suspicion that the junior is hiding the true extent of the gain derived from their collusion. Corruption can flourish because nobody is 'too much' of an expert.

3.3.6 Endogeneity concerns

As with any non-experimental study there is a risk that the key explanatory variable just described (purchase specificity) is correlated with the error term in the equation estimating the dependent variable ('red flag' measures of corruption). The most likely sources of this problem are omitted variables correlated with both purchase specificity and corruption, and corruption reverse-causing purchase specificity. I address the first source fully in the following chapter, where my approach using purchase-level data provides fine-grained and theoretically-informed controls for plausible omitted variables.[90] These are only meaningful at the level of individual purchases where the logic of my theory operates; aggregation to agency level as a mean or median suppresses the very purchase-level variation that connects them to my theory. The number of bidders applying to an auction, for example, makes sense only in the context of a particular product market, and does not have a clear interpretation at the agency level across all its purchases.[91] One potential omitted variable, the amount spent, makes equal sense at the purchase and agency levels, but for different reasons. In empirical tests in this chapter, I control for total agency

---

[90]Length of auction, degree of favoritism towards suppliers, and number of bidders applying.

[91]This is because the market of suppliers is defined by the product purchased, and this definition is muddied by aggregating purchases. Furthermore, at the agency level the mean/median number of bidders is likely collinear with purchase specificity; as I show in Chapter 4 more specific (i.e. less common) products have fewer suppliers.

spending as the best available proxy for the agency-level state capacity argument, that higher-spending agencies are more able to resist corruption. Another mechanism, working in the opposite direction, is that corruption is worse where spending is greater; I control for this interpretation in Chapter 4.[92]

The second source of endogeneity, reverse causation, is plausible because much of the literature reviewed in Chapter 2 implies bureaucrats will adjust their behavior to maximize opportunities for corruption. In the context of procurement, this would mean bureaucrats choose what to purchase based on how much corruption they expect to be possible in the process. While this possibility can never be entirely dismissed, several features of my agency-level approach in this chapter, and my more detailed purchase-level approach in Chapter 4, mitigate this concern. In the more conservative tests here, aggregation to the agency level reduces the likelihood of reverse causation: while it is plausible that individuals' purchasing decisions *within* an agency are driven by the anticipated success of corruption, what a given agency can buy *overall* is constrained by its mission and the need to deliver at least some minimal level of policy implementation.[93] In the following chapter, the purchase-level omitted variables noted above account explicitly for the several channels that extant literature predicts would change bureaucrats' decisions of what to buy.[94] Including these in my fully-specified purchase-level models does not systematically change the estimates of my core variables of interest, building confidence that reverse-causation is not a threat to my inferences. Finally, this concern is further mitigated by careful avoidance of strong causal claims in my empirical tests. In the next chapter I use the metaphor of a 'recipe' to interpret whether corruption is produced under the conditions predicted by my theory: if my theory is correct, certain ingredients should be

---

[92]I always log-transform spending measures as they are heavily right-skewed.

[93]Expanding the intuition of my theory, if every bureaucrat in an agency tries to purchase only the most corruption-prone products, it will swiftly come to the attention of its political overseer. If the Treasury tries to purchase its own fighter jets, this too will become obvious.

[94]I outline these channels, and how I measure them, in Chapter 4.

observed together in purchases alongside successful corruption. I observe, at several levels of analysis, the 'mix' of ingredients predicted by my theory.

3.4 Preliminary test of the theory in Moscow

So far in this chapter I have developed a research design for studying variation in corruption at the agency level. I identified procurement as an activity shared across agencies, susceptible to corruption, and for which measures relevant to my theory can be plausibly defined. I defended Russia as my source of data then proposed and validated a measure of opportunities for corruption (auction efficiency), several proxies for actual corruption, and a measure of the key independent variable in my explanation. In this section I use these resources to test my theory's implications using purchases within Moscow. I begin by reiterating implications of the theory in detail, before estimating and interpreting minimally-specified statistical models.

The empirical strategy in this chapter is conservative by design. In the following chapter I use all available data from all agencies, and measure purchase specificity at the individual purchase level.[95] Here I restrict attention to agencies with 100 or more purchases over 2011-2015, measure purchase specificity at agency level using a more blunt classification of goods and services, and weight it by spending as described above. While this obscures part of my micro-level theory, it tests the feasibility of its core implications on a subset of the data with lower potential for spurious noise.[96]

3.4.1 Detailed empirical implications

Most theories of corruption (including mine) predict that corruption opportunities are correlated with actual corruption: the more the 'door to the safe' is left ajar, the

---

[95]I motivate this quantity of interest in the next chapter from the opposite perspective, as *product commonness*. It works the same way but is a better fit for the measurement approach in that chapter.

[96]In particular the aggregation of behavior across at least 100 (and up to 4,285) distinct purchases attenuates the effect of any exceedingly unusual procedures by a small agency that might perturb my measures. I validate my measures individually using purchase-level data in the next chapter.

more likely are bureaucrats to take advantage. This dissertation was motivated, however, by the inadequacy of this unconditional hypothesis. If it were true, *all* agencies in Russia would be as corrupt as each other. Instead, I advanced a conditional version of this hypothesis: the conversion of corruption opportunities in to illicit gains depends on successful collusion between bureaucrats, which in turn depends on the information asymmetry generated by the specificity of the agency's purchases.

My theory predicts that increasing the junior bureaucrat's information advantage (greater purchase specificity) decreases the likelihood of corruption, all else equal. This negative effect should be strongest when corruption opportunities are relatively scarce (high median auction efficiency); as opportunities become more plentiful (low median auction efficiency) this effect should weaken, as the potential reward to the senior bureaucrat eventually overcomes the risk of punishment.

3.4.2 Models

To test these implications I specify an OLS model suitable for testing conditional hypotheses, namely an interaction model as recommended by Brambor, Clark and Golder (2005). This form makes clear that the effect of my key independent variable, *purchase specificity*, depends on the opportunities for corruption:

$$
\begin{aligned}
CorruptionProxy = \beta_0 &+ \beta_1 PurchaseSpecificity + \beta_2 CorruptionOpportunities \\
&+ \beta_3 (PurchaseSpecificity \times CorruptionOpportunities) \\
&+ \beta_4 TotalSpent(log) + \epsilon
\end{aligned}
$$

In these agency-level empirical models I include the log (base 10) of the total spent by agency, as a control for the alternative capacity argument specified above, which stated agencies more experienced at procurement would engage in less corruption. According to my theoretical implications derived above, I expect the sign on β_1 to be

negative, and the signs on β_2, β_3 and β_4 to be positive. I estimate one model for each of the five 'red flag' measures of corruption separately, and present the interaction plot for each model below. All models are estimated on measures at the agency-level, in Moscow only, for agencies with at least 100 purchases (N = 775).

3.4.3 Results

I report results of the model above for the five red flag measures in Table 3.5. My theory's core implication is that the interaction between purchase specificity and corruption opportunities, as estimated by β_3 in the equation above, and reported as **Specificity x Opportunities** below, is positive and significantly different from zero. Coefficient estimates for three of the five red flag measures are as predicted.[97] Combined with my prediction that the unconditional effect of specificity (β_1) is negative, this implies that when corruption opportunities are limited, corruption is more frequent in agencies purchasing more specific goods.[98] This effect weakens the greater the opportunities for corruption, which is an intermediate prediction of my theory. Although obscured here by the extent of aggregation, and my restriction to agencies with at least 100 purchases, I demonstrate in Chapter 4 that in some regions of Russia (including Moscow) the interaction term is strong enough that empirical tests support all predictions of my theory at once. In addition to what we observe here, in those tests when corruption opportunities are prevalent (as opposed to limited), *corruption is less frequent in agencies purchasing more specific goods.* In other words, when opportunities for corruption are sufficiently high, corruption is more widespread the less specific (more common) the average product, where the information asymmetry between senior and junior bureaucrat is lower. Though not conclusive, due to their intercepts being too far below zero, the models estimated here and their marginal

---

[97]I explain below when interpreting the marginal effects plots, and in the next chapter in greater detail, why the price-driven red flags are less likely to generate the predicted pattern.

[98]'More specific goods' here are 'less common products' in the language of Chapter 4, where I observe the equivalent pattern in more detailed data.

effect plots below provide tentative evidence in favor of my theory.[99]

For further interpretation of these tests I focus on the patterns in marginal effects plots that illustrate each model's results for the core implication of my theory. The results of four of the five models are in line with my expectations. In these, the estimated effect of higher purchase specificity (greater information advantage for the junior bureaucrat) is negative when corruption opportunities are scarce. As this resource constraint is relaxed, the negative effect weakens. When corruption opportunities are at their greatest (median auction efficiency close to 0), changes in purchase specificity are estimated to have no effect. The *no price decrease* measure is substantively strongest, followed by *bunched initial prices* and *dramatic price decrease* measures. The *single supplier* measure is substantively very weak.

In the fifth model, where the dependent variable is *winner concentration* (Figure 3.13), the estimated coefficient is indistinguishable from zero for most of the range of the corruption opportunities measure. When opportunities are near their maximum, a higher purchase specificity scored is estimated to increase this measure of corruption. One explanation for the different results estimated with this 'red flag' measure as the dependent variable is its positive-/right-skew: its range is 0-1 and the highest-scoring agency is 0.9, but 75% of agencies score less than 0.1. This could indicate that the Herfindahl-Hirschman index does not capture well what I intended: favoritism towards particular suppliers that is likely correlated with corruption.

These tests represent a conservative initial approach to the variation in my data and what it can demonstrate about corruption at the agency level. In the next chapter I expand this approach significantly, using a variety of models and several levels of analysis (individual purchase, agency, region) to fully explore the evidence for and against my theory and its alternatives. This culminates in a mixed effects model that

---

[99]In Chapter 4 I take the same logic to the level of individual purchases, re-validating my red flag measures and building up a more complete picture of variation within and across agencies, and across regions.

Table 3.5: Preliminary tests of theory using five 'red flag' measures, large* Moscow agencies

	No price decrease	Single supplier	Bunching	Large price decrease	Winner concentration
	(1)	(2)	(3)	(4)	(5)
Purchase specificity	−0.298**	−0.042***	−0.270***	−0.121**	−0.115
	(0.121)	(0.014)	(0.069)	(0.050)	(0.084)
Corruption opportunities	0.407***	−0.025***	−0.176***	−0.339***	−0.135**
	(0.079)	(0.009)	(0.045)	(0.032)	(0.054)
Total spent (log)	0.033***	−0.002*	−0.019***	0.001	0.035***
	(0.008)	(0.001)	(0.004)	(0.003)	(0.005)
Specificity x Opportunities	0.228	0.048***	0.285***	0.085	0.246**
	(0.142)	(0.016)	(0.081)	(0.058)	(0.098)
Constant	−0.383***	0.041***	0.426***	0.388***	−0.156***
	(0.087)	(0.010)	(0.050)	(0.036)	(0.060)
N	775	775	775	769	775
R²	0.384	0.018	0.061	0.491	0.098

*p < .1; **p < .05; ***p < .01
Note: these tests are restricted to agencies with at least 100 purchases

115

Marginal effect of purchase specificity on 'no price decrease' corruption indicator,
for agencies in Moscow, 2011–2015
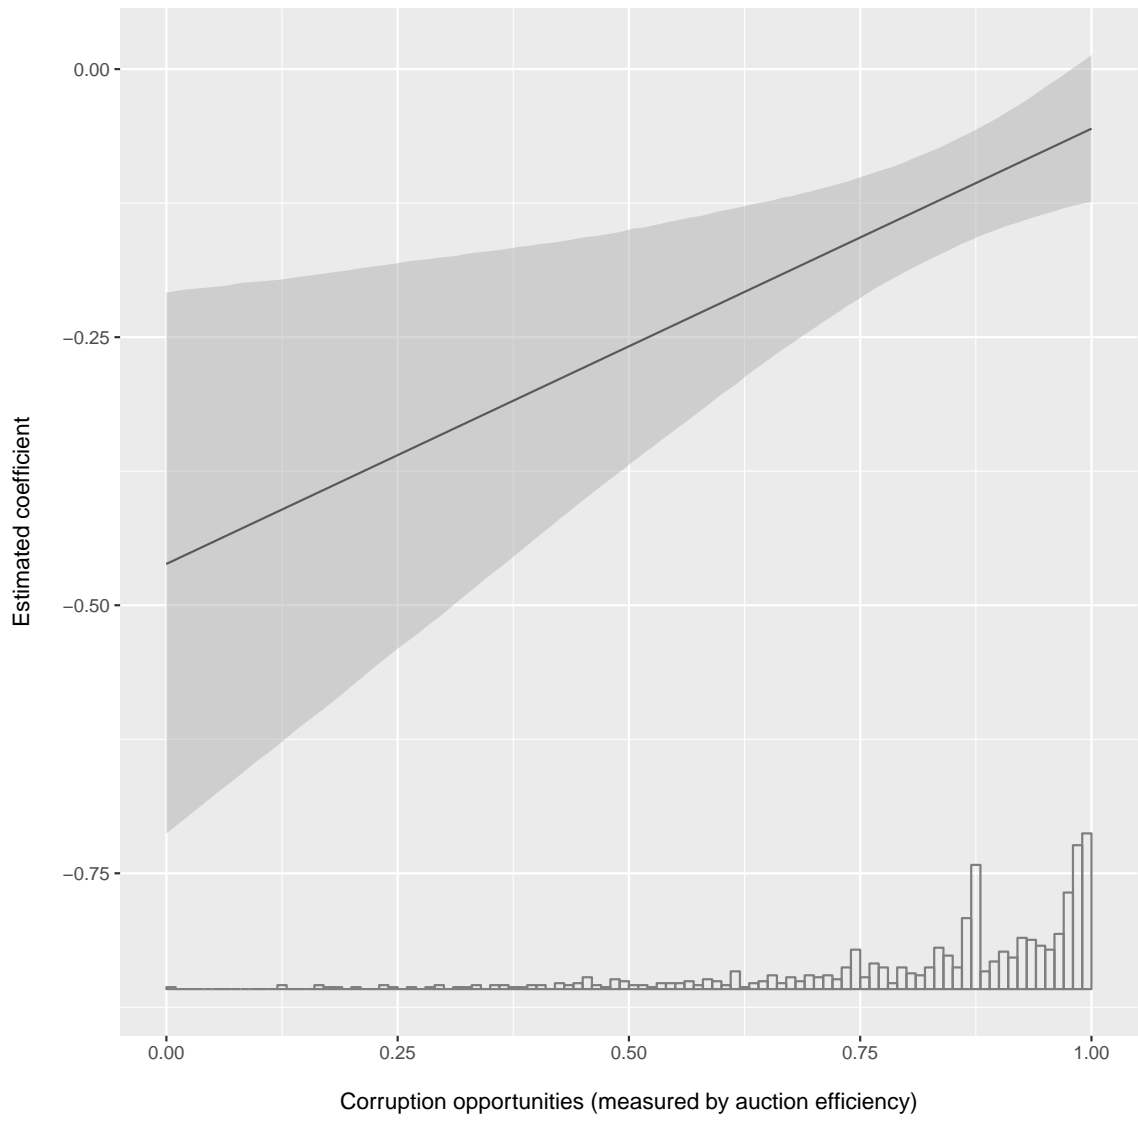
Figure 3.9: Marginal effect plot for model estimating red flag 1: no price decrease.

Figure 3.10: Marginal effect plot for model estimating red flag 2: single supplier.

Marginal effect of purchase specificity on 'bunched initial prices' corruption indicator, for agencies in Moscow, 2011–2015
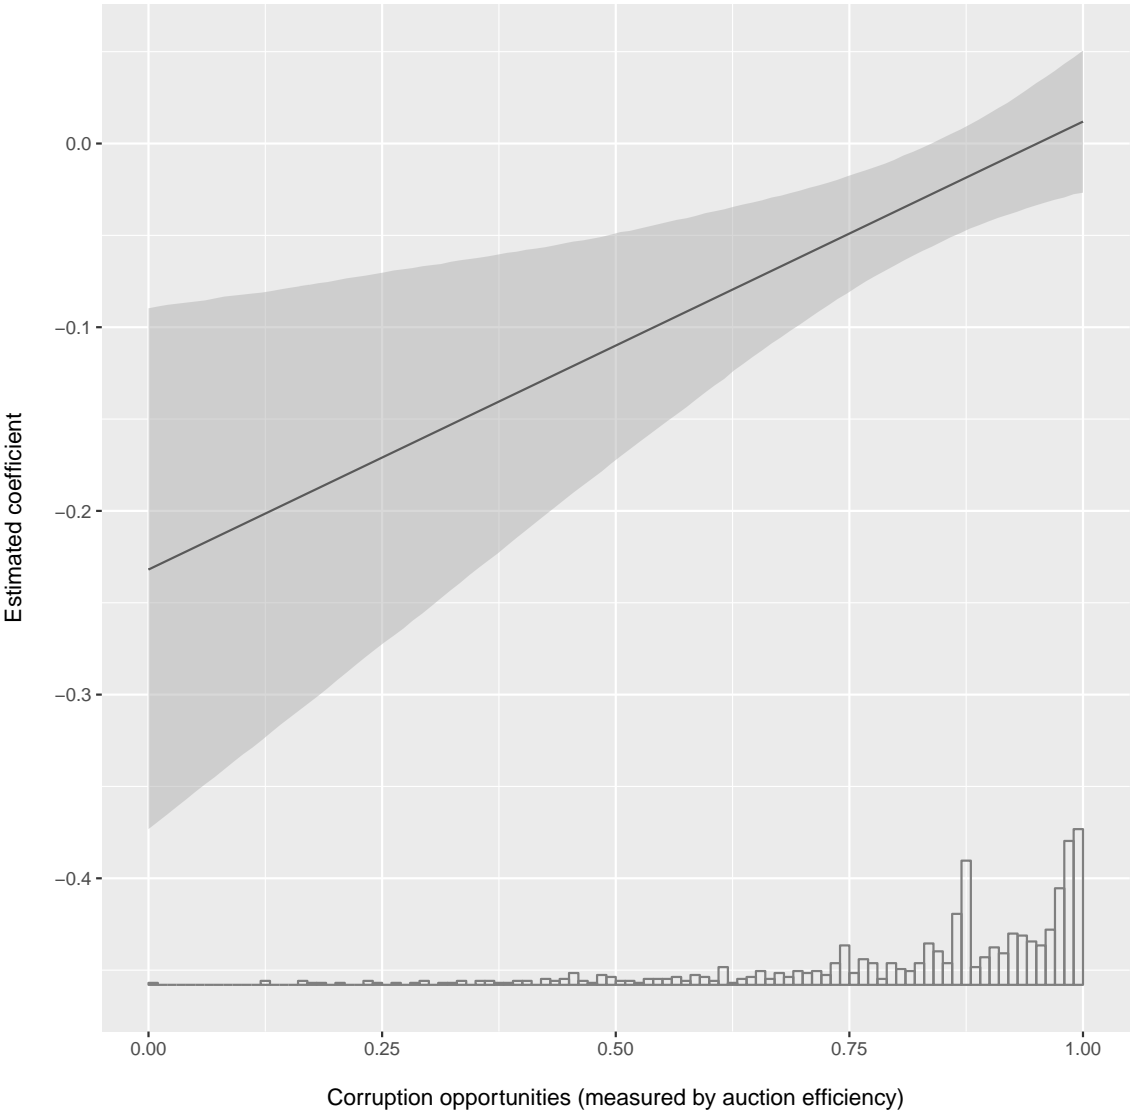
Figure 3.11: Marginal effect plot for model estimating red flag 3: bunched initial prices.

Marginal effect of purchase specificity on 'dramatic price decrease' corruption indicator, for agencies in Moscow, 2011–2015
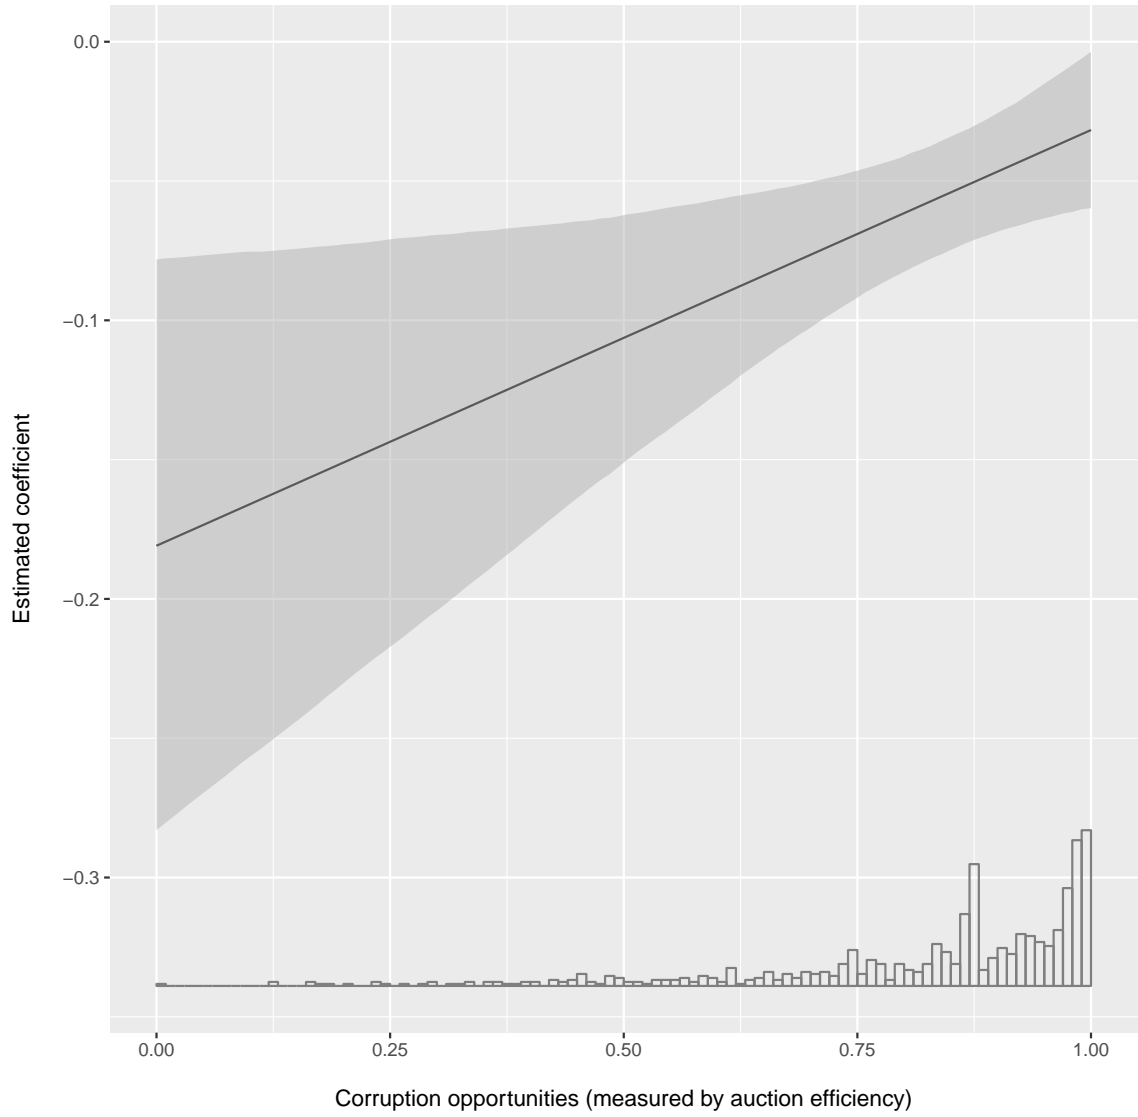
Figure 3.12: Marginal effect plot for model estimating red flag 4: dramatic price decrease.

Marginal effect of purchase specificity on 'winner concentration' corruption indicator, for agencies in Moscow, 2011–2015
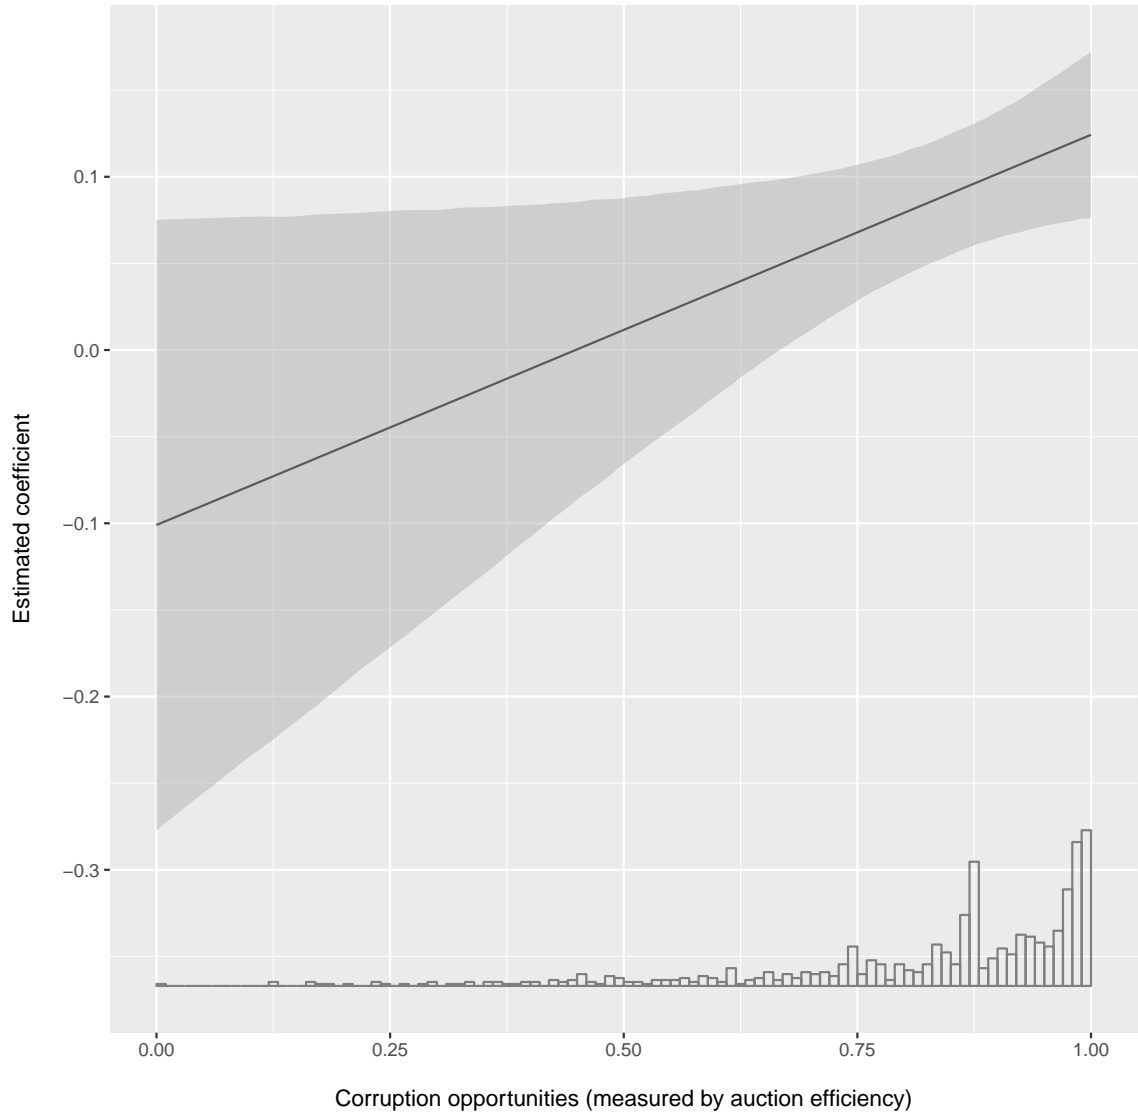
Figure 3.13: Marginal effect plot for model estimating red flag 5: winner concentration.

tests my agency-level theory within the context of regional variation, where Russia's regions stand in for states in cross-national theories of corruption.

3.5   Conclusion

In this chapter I developed an approach to measuring variation in agency corruption using data on an activity common to all agencies, procurement of goods and services. I applied this approach to the 550,000 procurement procedures engaged in by agencies in Moscow from 2011-2015, generating and validating proxy measures of corruption opportunities, and corruption itself. Aggregating these by agency, I found the core implications of my theory are supported by most of the evidence.

In agencies where opportunities for corruption are more scarce, a more unique basket of procured goods and services associated with lower likelihood of actual corruption according to my proxies. This is consistent with my explanation for curiously clean agencies in corrupt settings: when competition for opportunities is more pressing due to scarcity, and the information advantage of junior bureaucrats is greater due to higher purchase specificity, the collusion necessary for corruption to succeed is more difficult to maintain. Under such scarcity of opportunities, only agencies purchasing more generic goods and services can avoid the bad outcome from the perspective of bureaucrats, and turn opportunities in to extraction. Conversely, when corruption opportunities are plentiful this purchase specificity effect weakens. This is consistent with my explanation, that the conflict between bureaucrats is less pronounced when there are ample rents to go around.

These patterns were observed across several different proxies for actual corruption, all derived from the logic that exceptions to standard procedures are correlated with actual corruption. Having found robust evidence for the associations expected by my theory in just one area of Russia (Moscow), a strategy that helped validate my measures by controlling for important market-driven confounds, the next chapter

develops and validates similar measures for much more detailed purchase-level data. I include all agencies in my dataset, and test more comprehensive models.

CHAPTER IV

Testing the theory within and across Russia's diverse regions

"The fight against corruption must become truly national, and not the subject of political speculation, nor a field for populism, political exploitation, campaigning and throwaway solutions." - Putin2012.ru campaign website

In the previous chapter I developed a research design that can resolve the puzzle motivating this dissertation: why corrupt states have clean agencies. I introduced a unique dataset capable of executing this research design, and demonstrated the plausibility of my measures in a limited geographic setting, to avoid several confounds. In the process, I collapsed 555,000 purchases by agencies in Moscow across 2011-2015 into 800 observations of individual agencies' aggregate behavior. This loss of information[1] ensured unusual purchases did not skew my agency-level measures, but came at the cost of discarding useful variation across the full set of purchases. In this chapter I use progressively more and more of the available variation in the Russian procurement data to test further implications of my theory.

First, I show that purchases where a 'red flag' is present are systematically different to those without. This is important to establish, as I will later score purchases

---

[1]Of course, for the purposes of empirical modeling there are fewer than 550,000 independent observations in Moscow, as the individual purchases by an agency are related.

by their red flags when summarizing corruption at the agency and regional levels. Second, I show evidence that corruption varies by agency as predicted by my theory: agencies purchasing more common goods have more purchases with red flags than those purchasing less common goods. This relationship is robust to several empirical approaches. Finally, I use significant variation in economic and political context across Russia to test my agency-level theory against correlates of corruption that are prominent in cross-national studies. This is possible because the richest region (Tyumen) has GDP per capita (PPP) equivalent to Norway, while the poorest region (Ingushetia) is equivalent to Iraq; regions also vary politically. A hierarchical model nesting agencies inside regions allows me to estimate the shares of variance in corruption explained at each level. Using data from a single national procurement system mitigates some confounding heterogeneity that has plagued cross-national work.

4.1   Are purchases with 'red flags' any different?

In this section I validate my red flag measures against purchase-level data, within Moscow and within Russia's regions individually. I present evidence that in purchases with red flags, the relationships between variables that would represent an efficient market-driven outcome are weakened, absent, or reversed. This builds confidence that red flags detect behavior that differs from norms of clean procurement, clearing the way to aggregate these by agency to finally explain agency-level corruption.

Throughout this chapter, I contrast the logic of corruption with that of a 'normal' procurement market, which is very simple: the more common the product being purchased, the more difficult is corruption because more suppliers results in greater competitive pressures. This contrasts with the logic of my theory, which runs in the opposite direction: the more common the product the lesser the information asymmetry between bureaucrats, making corruption easier to maintain.

124

4.1.1 Groundwork: more common/generic goods have more suppliers

In Chapter 2 I identified a hypothesis prevalent in the literature: procurement officials facing more competitive markets will tend to run more efficient auctions that reduce their opportunities for corruption through 'padding'. Extending this, for my theory to work, it must be the case that more common goods have more suppliers.

At the level of individual purchases, I define *product commonness* as the probability a purchase is within a given product category (good or service), for a given set of comparison purchases (within a region, agency, or both).[2] It takes a low value when the product is bought infrequently ('less common' or 'specific' goods/services), and a high value when the product is purchased frequently ('more common' or 'generic' goods/services). For example, in Moscow the most commonly purchased product is *4560227: Landscaping and gardening* and it has the third-largest number of unique suppliers (2,030). The second-most purchased is *4560521: Construction - assembly work*, and has the largest number of unique suppliers. On the other hand, *1112850: Gas condensate* is bought very infrequently and has only one supplier.[3]

Of course, if a product is purchased only once in a given market, I can observe only one unique (winning) supplier, and so the relationship between product commonness and the number of suppliers is trivially true. Figure 4.1 shows the relationship holds beyond this edge case within Moscow, and the pattern holds across regions generally: the more frequently a given product is purchased, the more unique suppliers it has. For a given set of agencies, we can assume that local suppliers of a product are highly

---

[2]At the whole-agency level in Chapter 3, I defined the conceptual inverse of product commonness, *purchase specificity*, as an agency's *deviation from the average spending profile across products* for all comparison agencies, to account for intensity as well as frequency. The term 'specificity' made intuitive sense when discussing an agency's aggregate purchasing profile, while ignoring purchase level variation. In this chapter focused on purchase-level data it makes more intuitive sense to focus on products and how common it is for agencies to purchase them. To link the two concretely: an agency buying only very specific products, which by definition have low values on the product commonness measure, would have a high agency-level score for purchases specificity.

[3]These examples show the very fine detail of product classifications available in the Russian data. The pattern in Figure 4.1 also holds if products are aggregated to higher levels of classification.

Relationship between product commoness and
number of unique suppliers, Moscow agencies

Number of (unique) suppliers

Probability a purchase is of a given product
(higher value = more common/generic product)
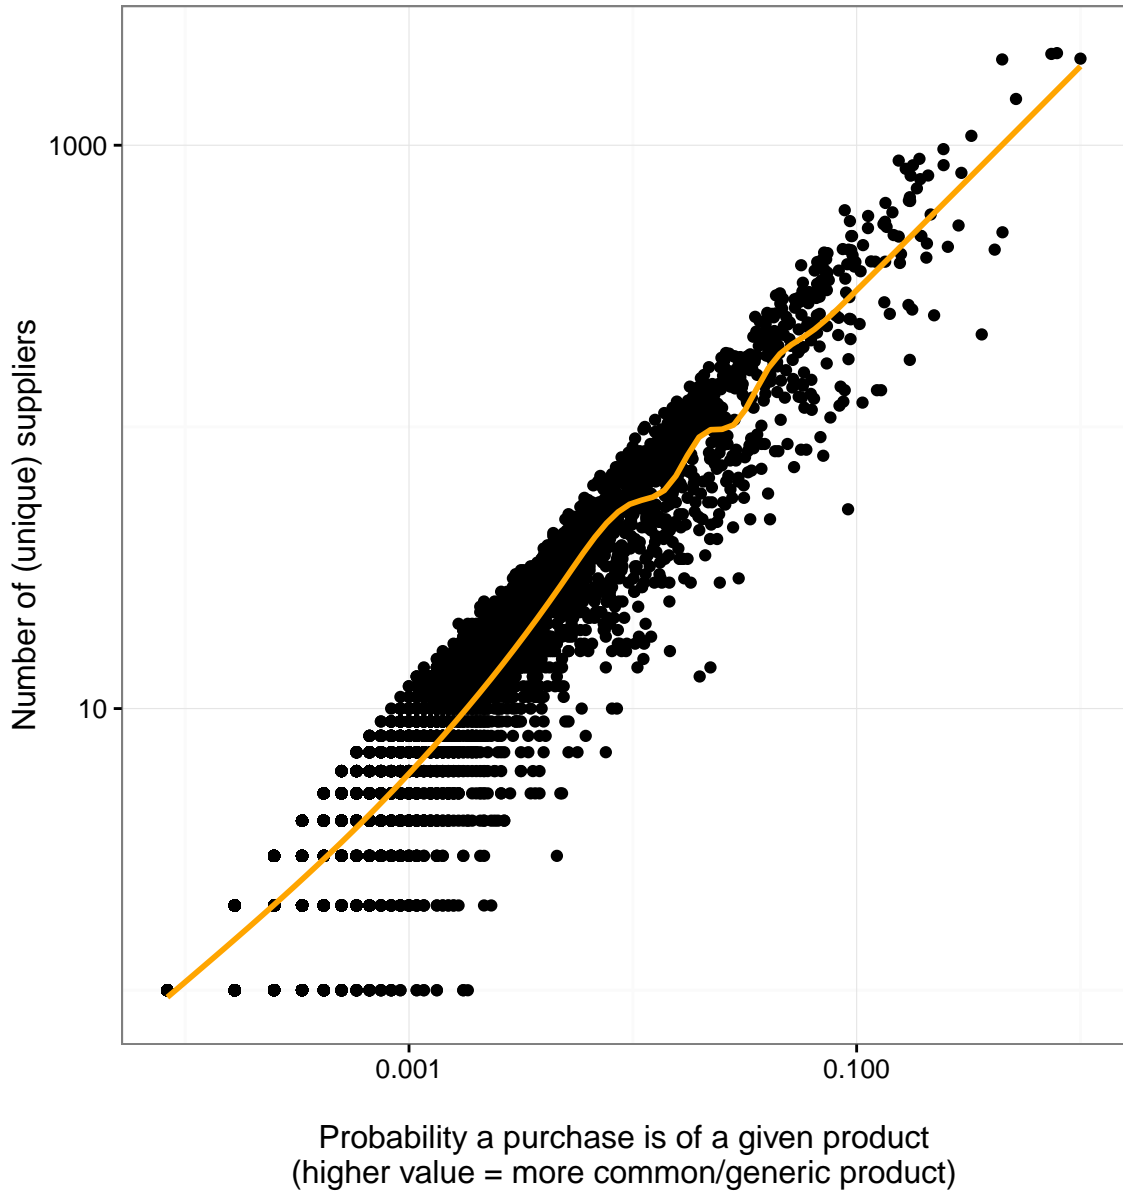
Figure 4.1: More commonly-purchased products have more unique suppliers.

likely to bid on other auctions of the same kind, especially because of the detailed product classification used in Russian procurement.[4] This evidence supports the hypothesized relationship upon which my theory relies: more common products have more competitive markets, less common products have less competitive markets.

This discussion of suppliers and product markets leads to a natural follow-up question: why do suppliers and their behavior not appear in my theory? Beyond the parsimony that comes from a sharp focus on exactly one interaction between two actors (senior and junior bureaucrat), I rely on the findings of the literature discussed in Chapter 2 to simplify my treatment of suppliers. While my theory of collusion between bureaucrats relies on their ability to collude with outside parties to a) systematically overpay for goods and services, and b) extract a portion of the resulting rents, when operationalized to procurement any variation in these factors across agencies is *already captured* by the notion of purchase specificity.[5] The most important effect of supplier behavior on corruption arises from whether they compete in saturated or sparse product markets; as just demonstrated this is strongly correlated with my measure of product commonness. Furthermore, in empirical tests below I control for the number of suppliers who apply to bid in each auction, which captures in sufficient detail the specific competitive environment in each purchase.

4.1.2 'Bunching' at thresholds where oversight increases

This red flag measure was first illustrated in Figure 3.1 in Chapter 3, which showed an unexpected number of purchases listed at an initial (maximum) price just below thresholds where the rules change.[6] Listing at an initial price beyond one of these

---

[4]Supporting this point, a number of third party services download daily the same data I did, and offer paid services including alerts when a relevant purchase is listed. Some also provide services to the agencies listing tenders, see for example: `http://www.zakupki.ru/art/static/8`

[5]Usually described here, at the purchase level, as product commonness.

[6]Specifically, the *request for quotes* procedure is forbidden for purchases with an initial price beyond 500,000 rubles (forcing choice of a more competitive procedure), while the *open electronic auction* procedure is more burdensome beyond 3,000,000 rubles.

thresholds reduces the discretion of bureaucrats as they are subject to additional constraints, providing an incentive to list just below. Later in Chapter 3 I found that the proportion of agency purchases with this feature was associated with my measure of information asymmetry between bureaucrats in the expected direction.[7] There is no feasible alternative explanation for the 'bunching' observed at agency level, besides bureaucrats manipulating the initial (maximum) price to maximize their discretion. When they buy one expensive item (e.g. car) they choose the maximum price; when they buy many smaller items (e.g. sugar) they can specify the quantity. More often than would be expected if the choice of initial price was driven only by market conditions, 'unit price times quantity' is close to these thresholds.

Despite this robust logic, 'bunching' is only observable across *many* purchases: any single purchase could be listed at an initial price near the threshold for perfectly innocent reasons. There are two defenses of this limitation. First, as all my corruption measures are proxies, a positive correlation between the underlying phenomenon and its observation is the standard, and is met by the evidence presented earlier.[8] A second defense that reinforces the first is my demonstration that the set of purchases with the red flag are systematically different from those without.

In the model[9] underlying Figure 4.2 the mean change in price between initial (maximum) and final (winning) price, when a purchase does not show this red flag and is of the least common product, is approximately -13%. For purchases far from suspicion of 'bunching' (left bar), moving through the sample from least common to most common product is associated with a further price decrease of around 0.9 percentage points (coefficient = -0.9). For this subset of purchases, the market-driven explanation holds: purchases of more common goods, for which there are more sup-

---

[7]On average across agencies in Moscow, those purchasing less common goods were less likely to engage in 'bunching'.

[8]If bureaucrats manipulate initial price, it is more likely the measure increases than decreases.

[9]A linear model with percentage price decrease as the dependent variable, the 'bunching' red flag interacted with product commonness, and a control for the size of purchase (N = 357,915).

Effect of 'bunching' red flag on association between product commonness and auction efficiency, Moscow agencies

AnyBunching1

Expected change in price decrease over course of auction, moving from least common to most common good

0 = Not listed within 5% of threshold; 1 = Listed within 5% of threshold
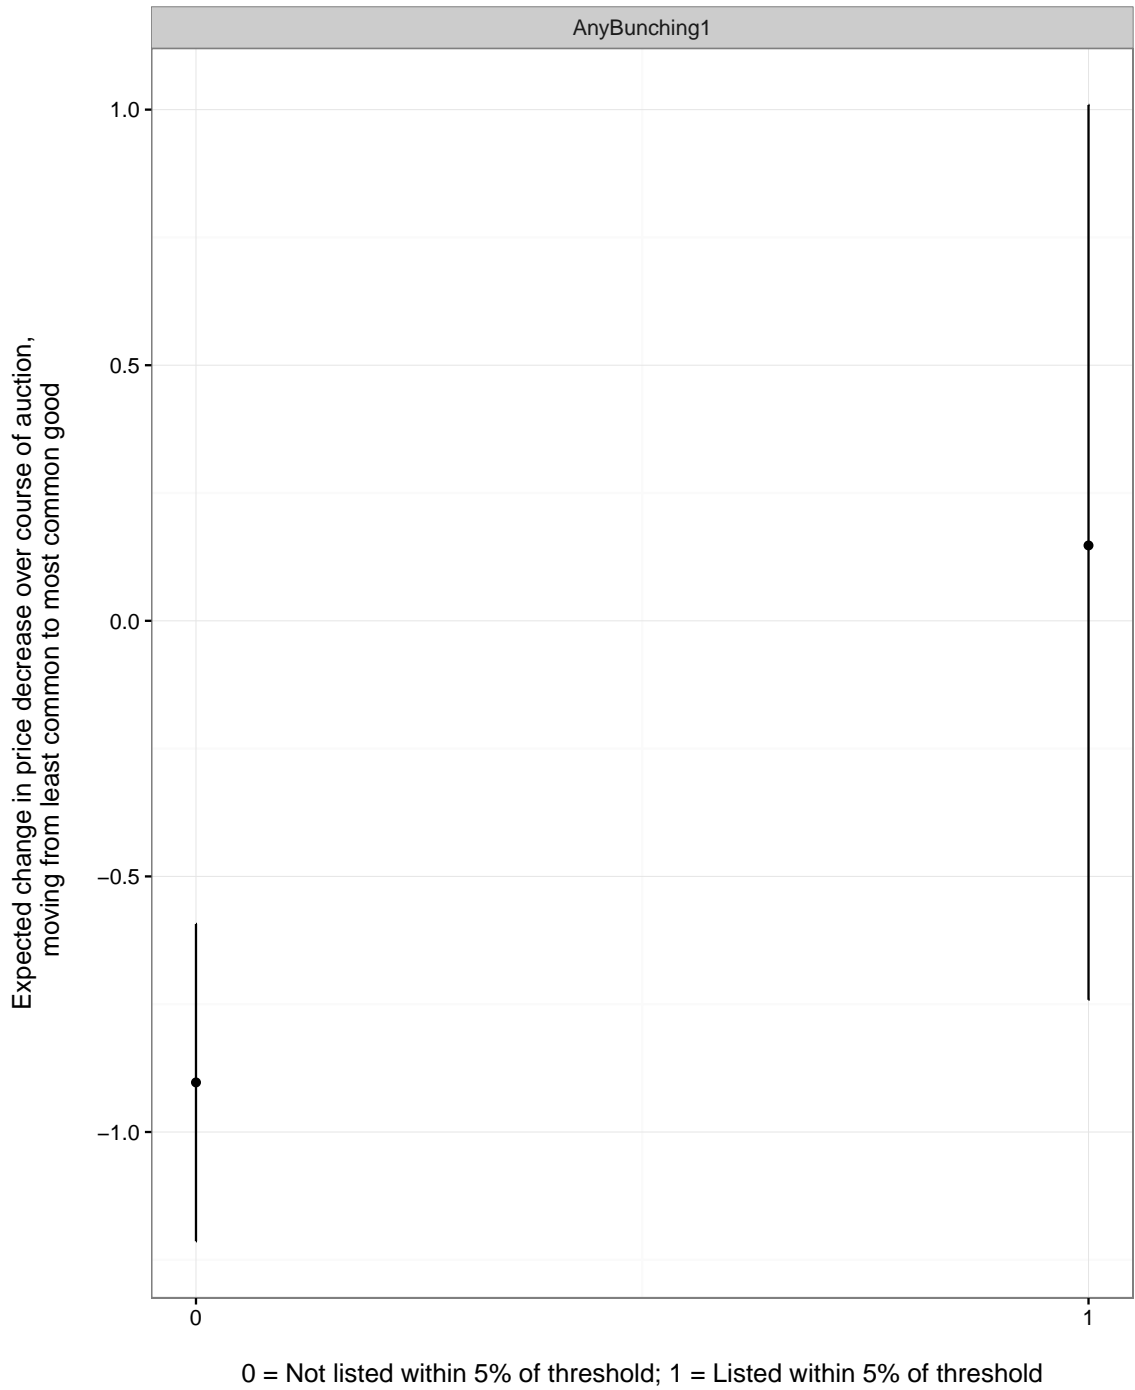
Figure 4.2: Buying more common goods reduces corruption opportunities, but only in auctions without 'bunching'.

pliers in greater competition, increases auction efficiency and reduces the opportunity for corruption.[10] Conversely, for purchases at high suspicion of bunching,[11] there is no such relationship. Though the coefficient estimate is above zero, there are fewer purchases in this subset (12%) and it is poorly identified. Nevertheless, the association between product commonness and auction efficiency for these purchases is different: when corruption is likely absent (left) outcomes resemble a well-functioning procurement market, when corruption is likely present (right) this mechanism breaks down. Thanks to the left-hand bar, we can at least conclude from this graph that at the level of individual purchases three things go together: absence of corruption, efficient auctions, and higher competition between suppliers of more common goods.

4.1.3 Revisions to purchases during the bidding process

A central idea of public procurement is that potential suppliers should all learn about a given purchase at the same time and that it should be well-enough described to fully inform their bids. Most guidance for practitioners recommends that in the event of serious irregularities at any stage of the process, a purchase be canceled and re-run.[12] Of course, occasionally changes are required mid-procedure, but these should clarify rather than fundamentally change the nature of what is being purchased.

A concrete case illustrates how this power of revision can be easily abused. In 2011, the Ministry of Emergency Situations announced an *open electronic auction* for construction of a new Moscow branch of the "A.M. Nikiforov All-Russian Center of Emergency and Radiation Medicine", a specialist hospital for victims of natural disasters.[13] From the initial announcement to the date of the electronic auction, four

---

[10]Recall that low auction efficiency is correlated with 'padding', which creates the rent that is a necessary condition for the bureaucrat to benefit from corruption.

[11]I code a purchase as having 'any bunching' if the initial (maximum) price is within 5% of the two most stark thresholds where rules change. The same pattern holds for a 1% threshold.

[12]This general advice appears in OECD (2016) and was given to Russia specifically by the World Bank before the system under study was implemented (The World Bank, 2006).

[13]The full detail of this case is available online at: `http://zakupki.gov.ru/pgz/public/action/orders/info/changes_and_explanations/show?notificationId=1145109`

revisions were made to the notification. Each time the auction date was extended but more importantly the initial (maximum) price increased five-fold, from US$34 million on June 6th, to US$185 million in the final revision on July 12th. The most sympathetic alternative explanation to 'corruption via extreme padding' is that new information came to light that increased the anticipated price. It seems unlikely that such an increase in absolute terms (US$150 million in five weeks) could have been completely unanticipated based on the building plans, only to be revealed during the procurement process in which suppliers may ask clarifying questions. Furthermore, after the final price hike but before the auction proceeded, two suppliers who were later disqualified from bidding filed official complaints. In these documents,[14] they allege that the price increase was due to the inclusion of additional requirements in the purchase, unrelated to construction, such as provision of medical equipment, soft furniture, and mirrors. Under Russian public procurement law, purchases are broken in to 'lots' which should only contain one class of good or service, ideally with a completely separate procedure for each. As such, the late revisions to the hospital's requirements, which could be characterized as 'ramming in' all manner of ancillary items, put the agency clearly out of compliance. I can observe that the complaints were received, but they did not change the outcome. Two bids were received and the lowest, at the extraordinarily precise amount of 5,521,309,287.20 rubles, was declared winner automatically by the auction system.[15]

To capture the likelihood that such manipulation has occurred, I count how many times a purchase is revised from the time of initial notification to the market, to after the conclusion of the contract. I cannot observe in fine detail exactly what changed with each revision, but the case above and Figure 4.3 tell a similar (and stronger)

---

[14]Available as PDFs in the online system, but only on a case-by-case basis.

[15]Maintaining the appearance of competition is crucial, as such a high-value purchase with a five-fold increase might be easily spotted by external auditors (e.g. the Federal Antimonopoly Service) and automatically canceled if the system showed only one bidder. This further motivates the senior bureaucrat's concerns in my model when presented with a share of the proceeds: the junior could be pulling the same scheme on the senior that they collectively must pull on the manager.

Effect of 'mid–procedure revisions' red flag on association between
product commonness and auction efficiency, Moscow agencies
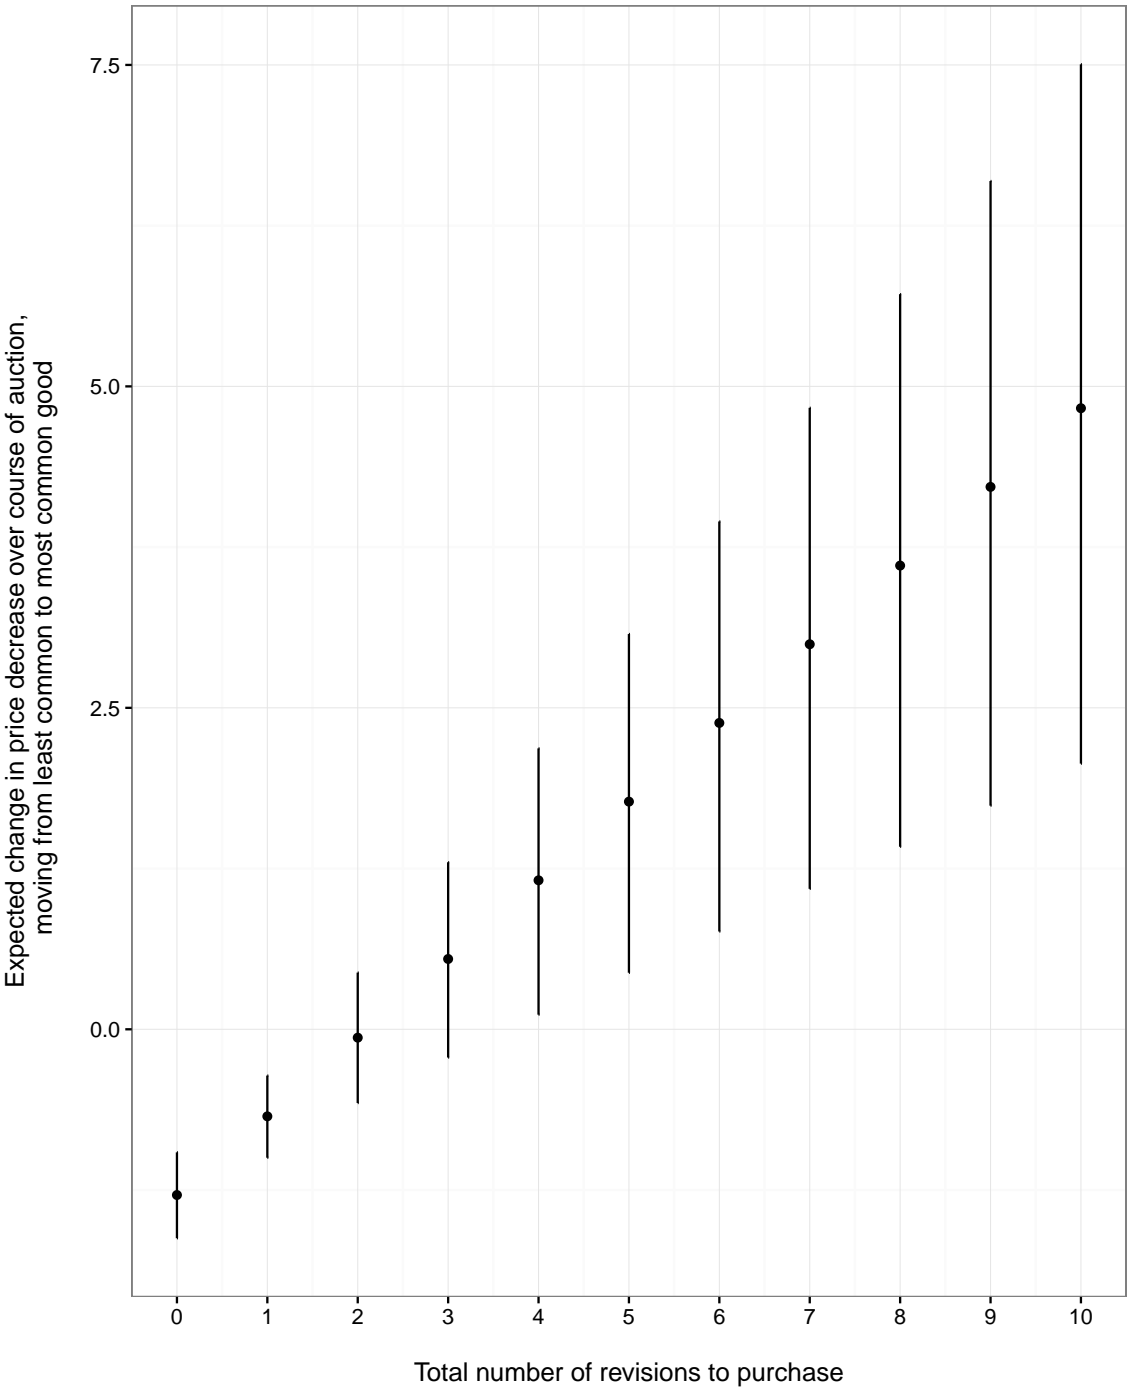
Figure 4.3: Buying more common goods reduces corruption opportunities, but only
when terms of purchase not substantially revised mid-procedure. As the
likelihood of corruption increases (further right) the association reverses.

story to what I showed with the 'bunching' red flag measure above.[16] When purchase details are not modified, or only a few times as might be reasonably expected, the market-driven explanation holds: more commonly-purchased goods, with more suppliers, are associated with more efficient auctions (coefficient significantly below zero) which present lesser opportunities for corruption. When purchase details are modified repeatedly, the opposite is true: more common products are associated with *reductions* in auction efficiency (coefficient above zero), enhancing the opportunities for corruption by increasing the likelihood of 'padding'. For clarity, in the baseline auction in this model the price decreases 13 percentage points; for purchases modified ten times, on average the price decreases only 8 percentage points (intercept at -13 and coefficient estimate significantly above zero at $+5$). This cautiously supports the core implication of my theory, that the following should appear together: more common goods (alleviating problems caused by information asymmetry), lower-efficiency auctions (creating 'padding' to enable extraction), and corruption (providing an instrument to enable a transfer of rents from winner to bureaucrat).[17]

4.1.4 Disqualification of bidders

In a well-functioning procurement market, several bidders should stand a chance to win each auction, otherwise it is trivial for bureaucrats and bidders to collude. Sometimes this will not be possible, due to small markets for less common goods, which creates a corruption risk as outlined earlier.[18] I am able to observe two features of the bidding process that help strengthen the empirical case for my measures and theory. First, for each purchase, I calculate the number of suppliers who *applied to bid* in the first stage of the procedure. As Russian public procurement during this period

---

[16]The underlying model, as with 'bunching', controls for the size of the purchase.

[17]The utility of this measure is further supported by a key difference between Federal Law 94 and its successor: the latter banned the kind of price adjustments described in the case above.

[18]These points are widely accepted in the literature, and were summarized earlier in hypothesis *L2d: Agencies facing procurement markets with higher levels of collusion between bidders are at greater risk of corruption.* Collusion is the converse of competition in this context.

offered no discretion for pre-qualification of bidders,[19] this is an accurate measure of the number of suppliers who believe they could supply the product. Within Moscow, some auctions attract hundreds of bidders, while many attract only a handful, and a significant proportion (7.5%) attract none and are automatically canceled.

The second feature I observe is the number of applicants who are *admitted* to the final stage of the auction vs *disqualified*. Reasons for disqualification are myriad, but the case outlined above (in which two bidders lodged complaints and were disqualified) suggests that disqualifying bidders is a discretionary tool that helps bureaucrats craft the market to the advantage of their preferred bidder(s). If the bureaucrat running procurement colludes with one or more bidders, and they anticipate a lower bid from a competing supplier, disqualification on a technicality offers an elegant method to ensure the 'right' bidder wins.[20] As long as the physical paperwork backs up the decision to disqualify, there is little recourse available to the disqualified suppliers beyond making a complaint.[21] Figure 4.4 shows the same story for this red flag as for the *mid-procedure revisions* measure above. In auctions where few bidders are disqualified, standard competition-driven market logic applies: purchases of more common products have lower final prices, on average 2% points lower. Moving right along the graph, to the auctions where a greater proportion of bidders are disqualified, the effect reverses: more common goods are associated with lower auction efficiency

---

[19]In a well-functioning procurement market, pre-qualification ensures bidders meet basic requirements, like legal status and non-bankruptcy, before bureaucrats review bids. Its exclusion during the period under study was motivated by prior over-use, allegedly to exclude non-aligned bidders, but is obviously not a panacea given the many other levers wielded by bureaucrats to alter the composition of the pool of competitors. The system in place since 2015 (under Federal Law 44-FZ) reintroduced pre-qualification under some conditions.

[20]In Russia, laws are often written and enforced precisely so as to produce these technicalities. This is natural consequence of the 'dictatorship of law' (Gel'man, 2000).

[21]Besides introducing unrelated products, as described in the hospital example above, this crafting can also be achieved by requiring bidders satisfy a 'market presence' test. For example, in 2014 the Veliky Novgorod City Legislature listed a 2 million ruble contract for information services (US$60,000 at 2014 rates), with the requirement that bidders' broadcasting catchment reach 90% of the city population and 70% of the surrounding region. The second requirement effectively ensured that Novgorod Regional Television, the only organization with the required regional coverage, would be the sole bidder, despite the City Legislature bearing no responsibility for region-wide affairs.

Effect of 'disqualifications' red flag on association between
product commonness and auction efficiency, Moscow agencies
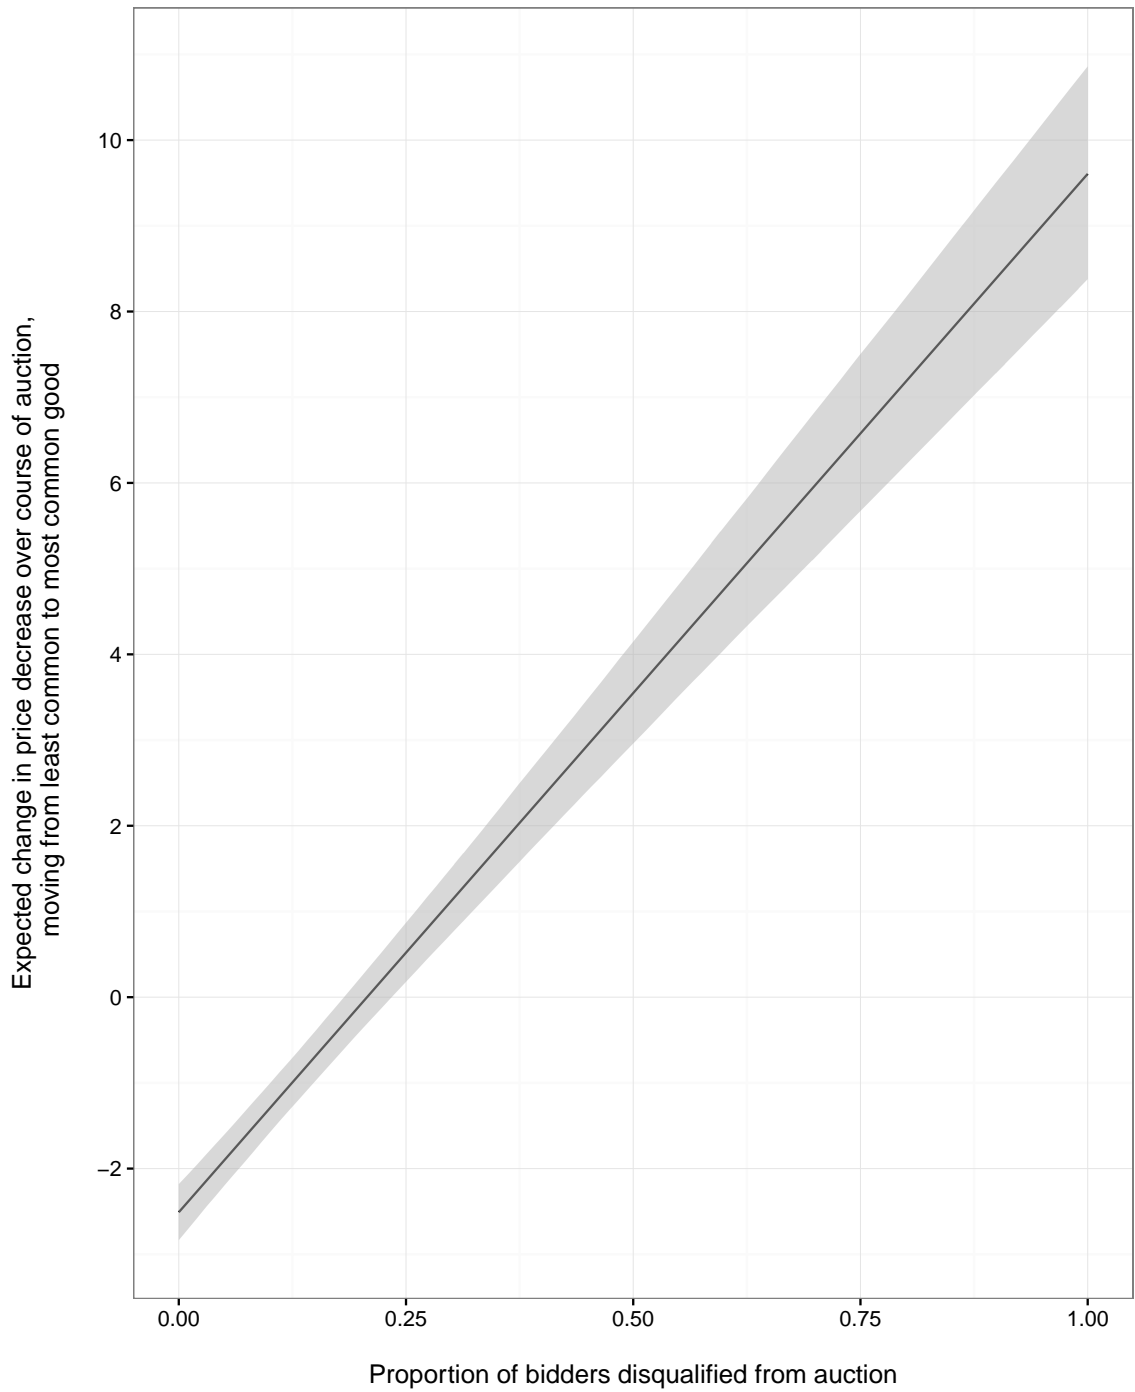
Figure 4.4: Buying more common goods reduces corruption opportunities, but only
when few bidders are disqualified. As likelihood of corruption increases
(further right) the association reverses.

(higher prices) creating ample opportunity for corruption. Once again common goods, less efficient auctions, and a corruption 'red flag' appear together.

4.1.5   Use of non-competitive single-supplier procedure

As discussed in Chapter 3, the rules governing Russian procurement reserve the single-supplier procedure for purchases where there is likely only one suitable supplier to begin with.[22] Nevertheless, as this judgment is made by the agency, it is open to abuse, and some agencies use it to such an extent that the only plausible explanation is avoidance of market pressures arising from competitive procedures.[23] As under this procedure, the supplier and agency proceed directly to contract negotiation, these purchases are missing the information I used above to validate 'red flags' at the purchase level. There is no initial (maximum) price, no field of bidders that can be whittled down, and revisions are not published. Instead, when estimating models on data aggregated to the agency level, I simply use the proportion of each agency's purchases under this most-discretionary procedure as a proxy for corruption. Across the 145,503 agencies in the all-Russia dataset, most use it sparingly (mean = 2.1% of purchases; median = 8.6%) while 227 use it for 80% or more of purchases.

4.1.6   Other 'red flag' measures based on price changes

In Chapter 3, when first testing the usefulness of my approach at the agency level, I developed two corruption measures based on price changes: the proportion of agency purchases with no price decrease, and the proportion with a dramatic price decrease (more than 50%). These measures were meaningful at the agency level because agencies should experience both rarely. At the purchase level, however, they are correlated with my measure of corruption opportunities, the percentage point

---

[22]For example, works of art and other cultural artifacts.

[23]As noted in the previous chapter, prior literature in Russia and elsewhere has established this procedure as the most at risk of enabling corruption.

difference between initial (maximum) and final (winning) price. Furthermore, there are many alternative explanations, including a poorly-set initial price (either too high or too low). This rules out the validation strategy I employed above for the red flag measures that capture specific interventions in the process by bureaucrats: bunching, revisions, and disqualifications. The advantage of these is their lack of plausible alternative explanations, whereas using price movements as purchase-level corruption measures would generate many alternatives. They were appropriate for illustrative purposes in Chapter 3 but I do not use them further.

4.1.7 From demonstrating effects of corruption, to explaining corruption

When validating three purchase-level measures of corruption above, I relied on a simple model of what should be observed in a well-functioning procurement market. I then showed how the presence of 'red flags' perturbs these expectations at the level of individual purchases within Moscow. As a prelude to the agency-level models below, where I show more conclusive evidence in support of my theory, I estimate a linear model explaining variation in each red flag measure, of the following form:

$$CorruptionProxy = \beta_0 + \beta_1 ProductCommonness + \beta_2 MaximumPriceLog$$
$$+ \beta_3 (ProductCommonness \times MaximumPriceLog)$$
$$+ \beta_4 ProcedureType + \epsilon$$

My product commonness measure is 0 when, in a given region of Russia, no agency buys that good or service. It is 1 for the most common (i.e. most-purchased) product in that region.[24] As there are thousands of product categories in the data, many individual purchases have product commonness close to zero, so I observe many

---

[24]These measures are calculated at the regional level to account for significant variation in market density and agency-level needs across Russia. I never specify, *a priori*, which goods and services any particular agency 'should' buy. Instead I define agencies empirically by what they choose to buy.

purchases at extreme corruption risk according to standard market logic (they have few suppliers, with whom collusion should be easy).

Table 4.1: Differences in corruption propensity by procedure type, Moscow

	Disqualifications	Revisions	Bunching
	(1)	(2)	(3)
Open electronic auction	0.090***	0.321***	0.307***
	(0.0004)	(0.002)	(0.0004)
Open tender	0.077***	0.544***	
	(0.001)	(0.004)	
Request for quotes	0.078***	0.150***	0.435***
	(0.001)	(0.002)	(0.001)
Olympic construction	0.123***	0.289**	
	(0.030)	(0.124)	
Preliminary selection	0.046***		
	(0.014)		
N	486,982	399,696	378,727
R^2	0.148	0.135	0.740
Residual Std. Error	0.206	0.765	0.214

*p < .1; **p < .05; ***p < .01
Note: Some procedures use no initial price, so 'bunching' is ruled out

I include the initial (maximum) price of the purchase[25] as I expect corruption is more likely when the 'prize' is larger, and I expect this to change how the mechanism of my theory works. A cynical interpretation of corruption in public procurement is that bureaucrats are following the maxim 'spend money to take money': the more they spend on a given purchase, the easier it is to include 'padding' as described earlier. The creation of this rent is a necessary condition for bureaucrats to benefit from colluding with each other and any favored bidder to manipulate the procedure. Conversely, purchases of lower value still require similar effort to run through the online system, while presenting only minimal opportunity for gains. I interact the maximum price with product commonness as I expect the information asymmetry

---

[25]Transformed with log base 10, due to right skew caused by relatively few very large purchases.

Table 4.2: Determinants of purchase-level corruption, Moscow (odds) and all regions (evens)

	Disqualifications		Revisions		Bunching	
	(1)	(2)	(3)	(4)	(5)	(6)
Open electronic auction	−0.028***	−0.005**		−0.217***		−1.813***
	(0.001)	(0.002)		(0.006)		(0.002)
Open tender		0.014***	0.090***	−0.050***		
		(0.002)	(0.005)	(0.006)		
Request for quotes	0.016***	0.016***	0.076***	−0.199***	0.286***	−1.478***
	(0.001)	(0.002)	(0.003)	(0.005)	(0.0005)	(0.002)
Olympic construction	0.031	0.035***	−0.054	−0.229***		
	(0.031)	(0.005)	(0.121)	(0.016)		
Product commonness	−0.458***	−0.175***	−2.303***	−0.634***	1.512***	−0.484***
	(0.012)	(0.003)	(0.046)	(0.008)	(0.012)	(0.003)
Maximum price (log)	0.025***	0.016***	0.213***	0.067***	0.339***	0.365***
	(0.001)	(0.0002)	(0.002)	(0.0004)	(0.0005)	(0.0002)
Commonness x Max price	0.077***	0.028***	0.423***	0.123***	−0.249***	0.090***
	(0.002)	(0.001)	(0.007)	(0.001)	(0.002)	(0.001)
Constant	−0.073***		−1.069***		−1.677***	
	(0.003)		(0.014)		(0.003)	
Region fixed effects	No	Yes	No	Yes	No	Yes
N	446,170	4,858,718	359,799	4,376,584	378,722	5,136,054
R^2	0.022	0.129	0.098	0.114	0.670	0.843

$^*p < .1$; $^{**}p < .05$; $^{***}p < .01$

effect at the core of my theory to be weaker in less valuable purchases.[26] I also control for the type of procedure as these have varying underlying propensities for corruption (reported for Moscow in Table 4.1)[27] but have no clear expectations about which is worse for corruption, and the 'leader' is different for each corruption proxy.

Table 4.2 shows the estimates for six variants of the model described above. Odd-numbered models are estimated on data from Moscow only, while even-numbered models use every purchase in the dataset and include fixed effects for the 84 Russian regions.[28] Naive fit (R^2) improves in the fixed effects models relative to Moscow-only models, and region-specific intercepts are both positive and negative. This suggests cross-regional variation is important in explaining variation across the 4-5 million purchases for which these corruption proxies can be computed. In all models but (5), the unconditional effect of more common products (i.e. moving product commonness from near 0 to 1, for very small purchases) is to decrease the respective corruption proxy. Increases in the maximum price are (unconditionally, for the least common product) associated with increases in the respective corruption proxies. To fully understand the implications of these models, however, we must plot marginal effects, as I predicted these effects depend on each other.

Figure 4.5 demonstrates how the effect of product commonness changes across the sample of purchases, from least to most expensive based on initial (maximum) price, based on the estimates from fixed effects model (2) above. This is powerful evidence in support of my theory at the level of individual purchases: for cheaper purchases, more commonly-purchased goods are associated with less corruption according to the 'disqualifications' measure. For more expensive purchases, when the potential 'spoils' of corruption are far greater relative to bureaucrats' other income, more common

---

[26]Less commonly purchased products are associated with greater information asymmetry between senior and junior bureaucrat, as the latter must do more 'leg work' to understand opportunities for corruption, which he may choose to keep private during bargaining over the 'spoils'.

[27]Model is $CorruptionProxy = \beta_p ProcedureType$ without an intercept, yielding conditional means. The latter two, *Olympic construction* and *Preliminary selection*, are used very infrequently.

[28]Several micro-regions have no procurement data for this period.

Effect of product commonness on 'disqualifications' corruption proxy
for different purchase values, all regions (model 2, region fixed effects)

Expected change in purchase–level corruption proxy
moving from least common to most common good

Starting price of reverse auction (a.k.a. initial/maximum price) chosen by agency, log base
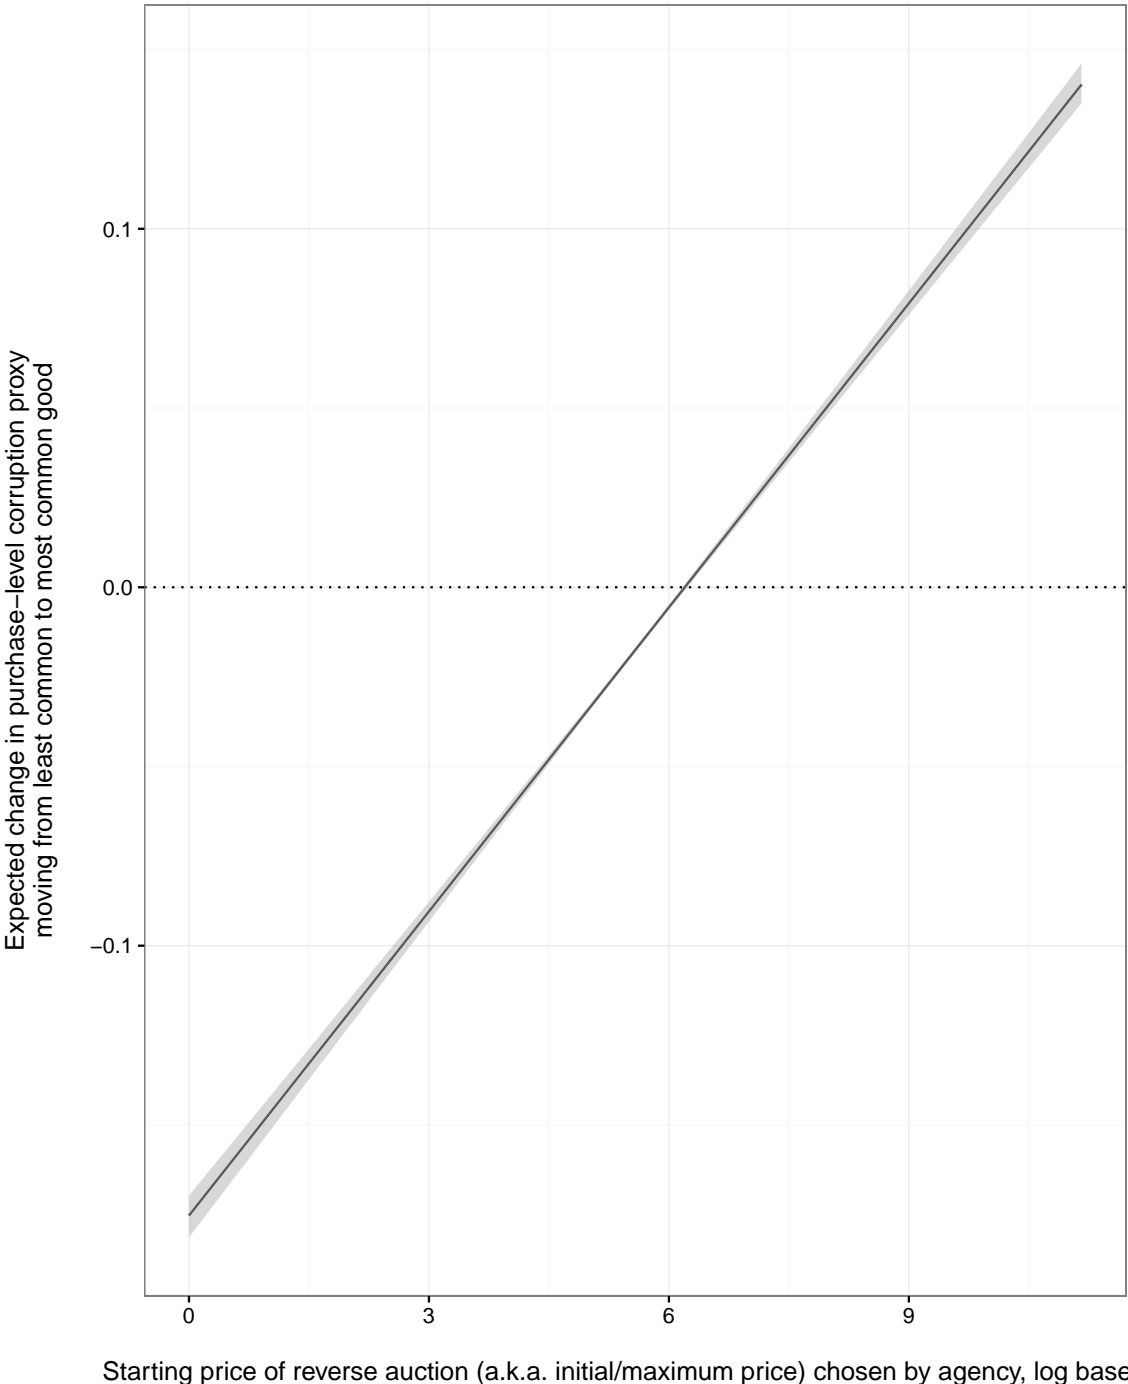
Figure 4.5: More common goods are associated with lower corruption, but only for
lower-value purchases. The association reverses for large purchases.

goods are associated with greater corruption. Note that the tight confidence interval is not a mistake, but rather a product of the large volume of data available across all 84 regions. For clarity, the median initial (maximum) price across the millions of purchases falls just below 6 on the x-axis (log base 10 scale), at 300,000 rubles or US$10,000 at the prevailing exchange rate across the period under study. The line representing the conditional coefficient estimate crosses zero at about 1,550,000 rubles or US$50,000. Overall this graph, and those for the other models in Table 4.2,[29] support the predictions of my theory at the level of individual purchases.

If this model shows 'congealed wisdom' about corruption accumulated by Russian bureaucrats over many purchases, it suggests that new entrants to the game should focus on less common goods when spending smaller amounts (left half), and more common goods when spending more (right half).[30] Of course, no bureaucrat (not even a new entrant) is completely free to choose what products to purchase: they operate in the context of an agency with a mission/purpose that defines and constrains its purchasing decisions. I cannot observe where these differences in mission come from, but I can conjecture as to how differences between them might be inferred. If agencies were purely shell companies for extracting rents, this purchase-level pattern would also be the cross-agency pattern. As I have now rigorously established that the predictions of my theory are borne out at the level of individual purchases, any systematic variation in the product 'mix' purchased by agencies should flow through to the incidence of 'red flags' across their purchases. I predict this will show further evidence of one of the broader implications of this study, that agencies must do at least *some* of what they are charged with, in order to get away with corruption.

---

[29]I do not present them here, but marginal effects plots for all models but (5) show the same pattern, including (6), the fixed effects companion to (5). This can be inferred by comparing the unconditional and conditional coefficient estimates across models in Table 4.2. I report the disqualifications measure as its link to the logic of procurement is clearest.

[30]I thank my father-in-law, Dr. Alexey Klyuev, Director of the Institute of Public Management and Entrepreneurship at Ural Federal University, for the insight that this dissertation is as much an instruction manual for upwardly-mobile bureaucrats keen to improve their extractive capabilities, as it is an account of why some agencies 'under-perform' in corruption relative to their peers.

In one sense then, the next section will repeat the analysis that closed out Chapter 3, which showed tentative evidence in support of my theory using a very conservative aggregation approach. In this section I have developed more detailed measures at the purchase level that are both theoretically plausible and empirically robust; this enables me to aggregate across entire agencies more aggressively than before,[31] and thus estimate more sophisticated models. Now I turn to agencies themselves.

4.2 How does corruption vary across agencies?

Having validated several plausible purchase-level corruption measures in detail, I now aggregate these measures to the agency level and report the results of tests of cross-agency variation using Moscow data. These are comprehensive tests of the core implications of my theory, building on the tentative tests at the end of Chapter 3. In the following section I extend these tests to the regional level, to establish in which regions the predictions of my theory hold up better or worse. This leads to the final tests of this chapter, where I integrate variation in agency corruption within regions, with variation in economic and political context across regions.

According to my theory, corruption is more likely the more common the product being purchased; these require less specialist expertise to buy, with corruption or without. This reduces information asymmetry between those authorizing corruption (the senior bureaucrat in my model) and those organizing it (junior bureaucrat). This is counter-intuitive as the logic of procurement markets suggests that less commonly purchased products, which have fewer suppliers by definition, should be more susceptible to collusion on both sides of the transaction.[32] I showed above, however, that

---

[31] For example, in Chapter 3 I limited attention to Moscow agencies that made at least 100 purchases across 2011-2015, as my agency-level measures (various means, medians, and shares across all purchases) were naturally sensitive to unusual purchasing patterns. In this chapter I have grounded measurement more firmly in the behavior of the bureaucrats in my model, rather than agency-level generalizations, allowing me to include all agencies across all regions in the remaining tests.

[32] At the very least, in the case of a heavily 'padded' contract for an uncommon product, if audited the organizing bureaucrat can claim a lack of deep knowledge of the market price as the reason.

many of these opportunities are left wanting: less common goods are only associated with greater corruption for smaller purchases.

4.2.1 Reproducing Chapter 3 results with finer-grained data

I first conduct high-level tests of my theory similar to those concluding Chapter 3, aggregating from detailed purchases to entire agencies. I estimate an agency-level model in Moscow (or any other given region) of the form:

$$CorruptionProxy = \beta_0 + \beta_1 MedianProductCommonness + \beta_2 MedianAuctionEfficiency$$
$$+ \beta_3(MedianProductCommonness \times MedianAuctionEfficiency)$$
$$+ \beta_\mathbf{A} AgencyControls + \beta_5 \epsilon$$

Compared to the tests in Chapter 3, these differ in three important respects, which combine to make them less conservative and more ambitious in using the variation in the data: all agencies are included (not just those with 100 or more purchases); product commonness is calculated at the individual level and aggregated (rather than across the agency's entire purchase profile at once); more agency-level controls are included based on the purchase-level calculations. Table 4.3 shows three variants of the model above, estimated against each agency's *mean proportion of bidder disqualifications per purchase* as the dependent variable.[33] Model (1) includes only the *median product commonness* across agencies' purchases. Its association with corrup-

---

This justification for a wildly inaccurate initial (maximum) price would be unwarranted in most cases. Under Russian procurement guidelines, bureaucrats are required to obtain price quotes from several potential suppliers for purchases over a certain value, and this requirement becomes more comprehensive and harder to avoid the more is being spent. Unfortunately I do not observe these quotes. This is another reason to control for the initial (maximum) price of each listed purchase, and often to interact it with another variable of interest as in 4.1.7 above.

[33]I generally present this measure over its peers, as its scale is intuitive and the conceptual link to more/less likely corruption is very clear. Models estimated against the other 'red flag' proxies, *bunching* and *revisions* give similar results, although the *bunching* measure is less robust to specification (3). This is at least in part due it its complex structure, measuring proximity to two thresholds depending on procedure type; while this presented no problems at the individual purchase level, at the agency level it has gone through several summarizing operations that reduce its sharpness.

tion is positive, as my theory suggests. Agencies whose 'median' purchase is of a more common good disqualify more bidders, though the effect is substantively small. Model (2) adds an interaction with *auction efficiency*, i.e. the percentage point price change over the course of the auction, which has a negative scale.[34] As this measure increases, a purchase is increasingly likely to become a source of rents; in Chapter 3 I defended it as a proxy for 'corruption opportunities', as inefficient auctions are a prerequisite for procurement corruption. This model accounts for agencies' opportunities for corruption, as well as their conversion of these in to likely corruption, as detected by my red flag measures. The effect of agency-level purchase commonness depends on average auction efficiency, which is consistent with my findings in Chapter 3, and in 4.1 above at the purchase level.

Table 4.3: Determinants of agency-level corruption, Moscow

	Mean Disqualifications			
	(1)	(2)	(3)	(4)
Agency spend (log base 10)	0.010***	0.009***	0.006***	−0.0001
	(0.001)	(0.001)	(0.002)	(0.001)
Mean listing duration			0.001***	0.001***
			(0.0002)	(0.0002)
Mean bidders applied			0.019***	0.018***
			(0.001)	(0.001)
Mean supplier favoritism			−0.004***	−0.002**
			(0.001)	(0.001)
Median product commonness	0.068***	0.102***	0.070***	0.062***
	(0.007)	(0.010)	(0.010)	(0.010)
Median auction efficiency		0.0001	0.001***	0.002***
		(0.0001)	(0.0001)	(0.0001)
Commonness x Efficiency		0.003***	0.003***	0.002***
		(0.001)	(0.001)	(0.001)
Constant	−0.012	−0.001	−0.042***	
	(0.009)	(0.010)	(0.011)	
N	4,014	4,014	4,013	4,013
R^2	0.034	0.043	0.142	0.447

*p < .1; **p < .05; ***p < .01

---

[34] A more negative number represents a more efficient auction, as the price has dropped further.

Model (3) adds controls for several agency-level features of procurement, summarized from their respective purchases, that are plausibly correlated with corruption and one of the other independent variables. As before, I include a spending measure, as I expect agencies that spend more will engage in more corruption, and buy a more diverse set of products. *Mean listing duration* is the average time elapsed between initial notification of a purchase to the market and the signing of a contract with the winner. Shorter time-frames might have two effects on corruption: they give bureaucrats increased discretion to push through a contract to their favored bidder; conversely given the complexity of corruption schemes in general, the effect could be opposite if urgency causes breakdowns in collusion. Regardless of which effect dominates, the time available to bid is plausibly correlated with the next independent variable, *mean bidders applied*.[35] The relationship of this variable to corruption is complex: if disqualification were not commonplace in Russia, the average number of bidders might cleanly represent the competitiveness of the particular local market for that product. Strategic bidders, however, facing bureaucrats who they know might be corrupt, could go either way: those keen to access the potential rents offered by padded contracts would bid even in the marginal case, whereas a clean firm, knowing the 'price' of admission, would avoid these opportunities altogether. In this model, the 'feeding frenzy' story seems to dominate. The more bidders that apply, the more (proportionally) are disqualified before the bidding even begins.[36] If the two types of firms react in Russia as I described, this association may indicate firms have good information about the rent-potential of agency purchases in their area of the market. This knowledge would increase the anticipated returns of non-clean firms and encourage them to bid in more cases.[37]

The final agency-level control I introduce in model (3) is a measure of *favoritism*

---

[35] In Moscow the correlation is 0.251.

[36] Correlation with dependent variable is 0.33 in Moscow.

[37] As public procurement represents approximately 30% of all government spending, it is a significant source of revenue for firms.

towards particular suppliers of a product. It is calculated at the purchase level within regions, based on the number of suppliers who ever win for each agency-product combination, and the total number of suppliers who have won for that product across all agencies in the region.[38] This measure improves on the 'shares of winning suppliers' measure I employed in Chapter 3, as it employs much more purchase-level detail before being aggregated by agency here. As well as its correlation with the dependent variable, if my theory is correct then the extent of favoritism should be correlated with auction efficiency, as the whole point of favoring suppliers is to dampen market competition and produce less efficient auctions with greater 'spoils' to extract.[39]

Including these controls improves the fit of the model markedly, with the biggest change in coefficient being the constant (intercept). As this has become significantly different from zero, and negative for a proportion measure (not technically possible), I re-estimate the same model without an intercept (4) to check what changes. This naturally improves the naive fit, but leaves other coefficients substantively unperturbed. As this is another interaction model, the evidence it provides for my theory is best examined with a marginal effects plot of model (4).

We have seen the pattern in Figure 4.6 before, in the tentative evidence at the end of Chapter 3. The left side of the graph represents the effect of moving an agency's median purchase in terms of product commonness from least common to most common, when auctions have been highly efficient.[40] For agencies that tend to run highly competitive auctions, and comparing one that buys relatively common

---

[38]The purchase-level version of this measure takes a value of zero (no favoritism) when an agency has engaged exactly as many suppliers for a given product as the number of purchases made of that product, as in these cases suppliers could have been selected from a uniform distribution. When the number of agency purchases of a given product exceeds its number of suppliers, there is the potential for favoritism, which I calculate as the ratio of available suppliers to winning suppliers (capturing the proportion of suppliers ignored), times the number of purchases over and above the number of winning suppliers (capturing the 'excess' purchases for which we can make inferences about favoritism). I then take the log as the measure is highly skewed.

[39]I resisted using favoritism as a 'red flag' measure as auction efficiency is an important part of my empirical modeling strategy, and the two are somewhat related (correlation in Moscow = 0.213).

[40]The relatively imprecise estimate of the marginal effect at this end is due to very few procedures concluding for almost 100% less than the initial (maximum) price.

Effect of agency–level product commonness on 'disqualifications' corruption proxy for different levels of corruption opportunity, Moscow (model 4)

Expected change in agency–level corruption proxy moving from least common to most common good

Average corruption opportunities
(change in price over course of auction; more negative = less opportunity)
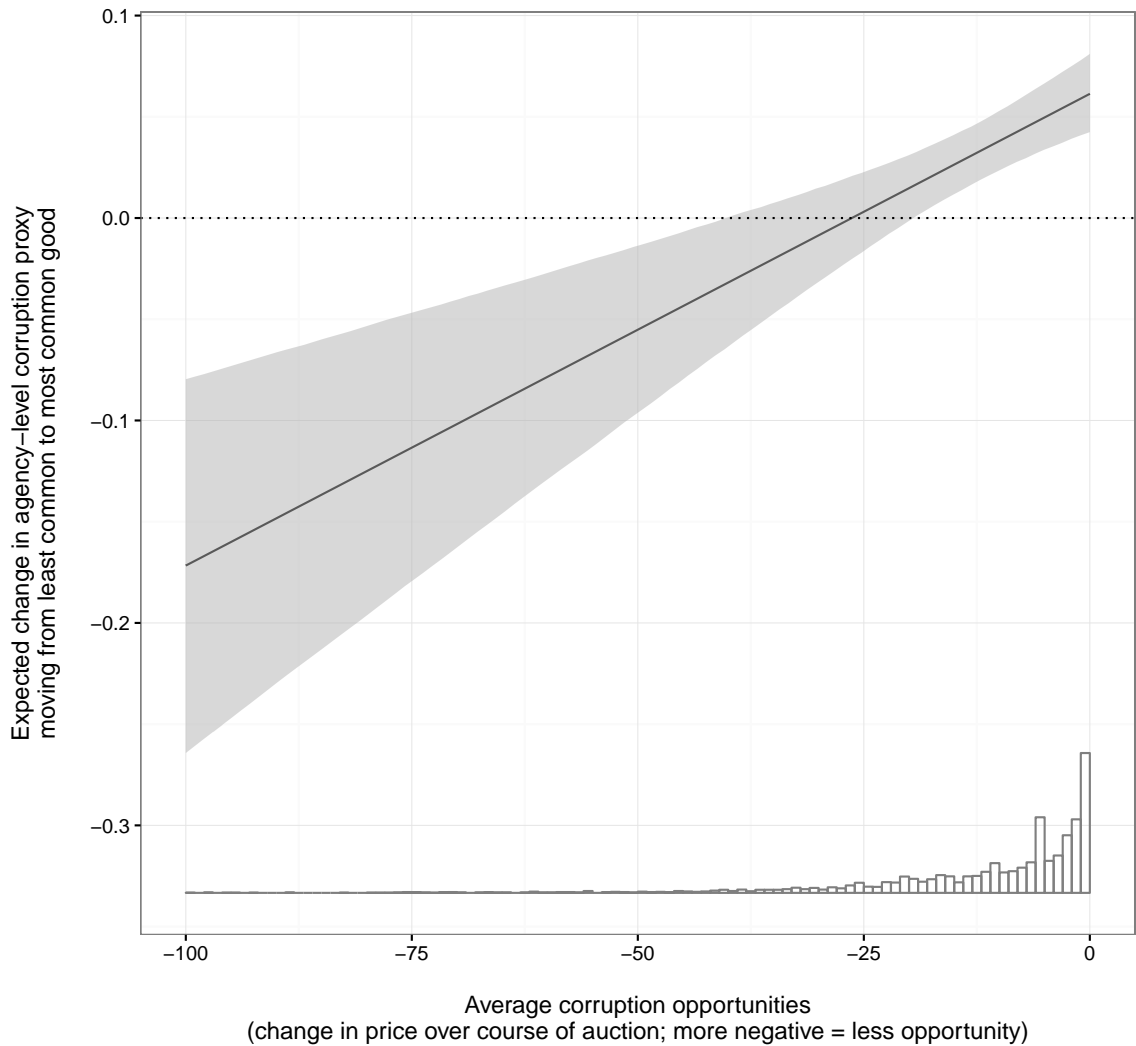
Figure 4.6: More common goods are associated with lower corruption in agencies running higher-efficiency auctions. The association reverses for agencies running less efficient auctions, where purchasing more common goods is associated with more corruption.

goods to one that buys less common goods, the latter engages in less corruption by this measure. For agencies where there are competitive auctions, the refuge of the corrupt bureaucrat is in the less common goods. On the right side of the graph the coefficient estimate reverses and becomes significantly different from zero.[41] Here reside the agencies that tend to run less efficient auctions, which generate greater opportunities for corruption under the logic developed in Chapter 3 and reiterated and tested at the purchase level in Section 4.1 above. Given that most advice to procurement practitioners recommends agencies aim for average auction efficiency in the range of -10% to -20%[42] this graph makes an additional point not clear from the results table above: in agencies that seem to follow (whether consciously or not) this recommendation, this model implies *corruption in these agencies is unrelated to the kinds of goods purchased.* This 'sweet spot' is where we would expect the majority of agencies to be in non-corrupt states. There, agency-level corruption is potentially more a function of individual bureaucrats' whims than the systematic agency-level story I have demonstrated here.

4.2.2 The answer: agencies in which corruption fails

The short answer provided by this section to the puzzle motivating this study is, at long last: clean agencies in corrupt states are those that purchase more common goods and services in efficient auctions, *and* those that purchase less common goods in less efficient auctions. The former set are predicted by standard market logic: corruption fails in these cases because competitive auctions force corruption-chasing bureaucrats to seek rents in the purchase of less common goods, of which there are fewer suppliers, protecting their rents from competition. The latter set are predicted by my own theory: corruption fails in these cases *despite* the fact the needs of the

---

[41]Weighting agencies in model (4) by their number of purchases strengthens the estimated effects both for more efficient (left) and less efficient (right) auctions.

[42]This advice is common in advice from international institutions, including OECD (2016) (general advice to all states) and The World Bank (2006) (specific guidance to Russia).

agency, revealed in its purchasing profile, are associated with less efficient auctions overall. According to the mechanism in my theory, this leaves corruption-seeking bureaucrats with little choice but to pass up the potentially lucrative rents that might be possible with the less-supplied products and focus on more common products.

Ministry name	Disqual.	Revisions	Bunching	Wage
Housing	5	1	1	1767
Education and Science	1	4	3	1896
Industry and Trade	8	2	8	1507
Internal Affairs	2	3	13	2520
Transport	3	6	11	1581
Communications	4	15	4	1653
Energy	11	10	5	1517
Finance	10	16	2	2209
Emergency Situations	15	8	7	3238
Foreign Affairs	12	13	6	1883
Regional Development	7	14	10	2219
Natural Resources and Ecology	19	5	9	1555
Economic Development	6	17	12	1679
Defense	9	12	18	1677
Culture	17	9	15	1572
Health	16	11	16	2707
Health and Social Development	20	7	17	2707
Sport, Tourism, and Youth	14	20	14	1294
Justice	13	19	19	1799
Labor and Social Protection	18	18	20	1820
Agriculture	21	21	21	1413

Table 4.4: Rankings of federal ministries by various corruption measures, and 2012 average monthly wage in USD

Table 4.4 shows how Russia's federal ministries, for purchases within Moscow, rank on the different red flag measures. In a few cases, such as Housing or Industry and Trade, we can conjecture the ranking is warranted by their deeper intervention in the economy,[43] but this is ad hoc. Based on the 'degree of intervention' hypothesis, the Ministries of Labor and Social Protection, Defense, and Economic Development should rank higher, while Communications and Foreign Affairs should rank lower. Yet across hundreds or thousands of purchases, they systematically frustrate these

---

[43]Russia maintains significant subsidies for social housing and the development of industry.

predictions. This demonstrates the shortcomings of modeling agency-level corruption solely as a function of the jobs that agencies do. While this might be true at the street level, where the Ministry of Internal Affairs (which includes the police) interacts with citizens and firms at traffic stops and document inspections, at the level of 'grand corruption' studied here, these notions fade away. The Ministry of Transport's high ranking fits nicely with my theory and with most literature's priors: though roads require specialist equipment to build, there are many suppliers of this product in Russia and this Ministry is not the only agency that purchases roads. While the propensity for corruption in large infrastructure projects is well-known empirically, my theory provides a detailed account for why this occurs that does not rely on anything particular to roads, dams, or building construction. What matters for the ability to extract of those up the chain of command in control of procurement, are the two well-described dimensions in my theory: what products agencies buy (largely determined by their missions) and how they buy those products (much more discretionary).

In the language of my theory, Housing does so well out of corruption because it purchases relatively common products (this is easy to imagine given all other agencies also purchase construction services), so the expertise required is low enough that the information asymmetry between senior and junior bureaucrat does not threaten collusion.[44] In Defense, on the other hand, the expertise requirement is plausibly far greater, as the products purchased are far more unusual. While under standard market logic this creates extensive opportunities for corruption (due to the small number of possible suppliers), the lower-than-expected ranking on my red flag measures is consistent with corruption failing for the reason identified by my theory: the information asymmetry between generals and procurement functionaries is too great. Understanding this important distinction between agencies is made possible by my novel empirical approach, which compares them across a common activity.

---

[44]Put simply, anybody can work out how to overpay for construction materials or services, and pocket the difference. This makes it difficult for the junior bureaucrat to short-change the senior.

4.2.3 Robustness: re-validating 'red flags' at the agency level

In Section 4.1 I validated my red flag measures using purchase-level data, and found that purchases with these present are significantly different to those without. To generate the agency-level data used in this section, I aggregated these purchase-level measures to represent the mean level of each red flag across each agency's purchases. While this was warranted, it does create another logical stepping stone between the purchase-level behavior where my theory and data reside, and safe inferences about corruption at the agency level here.

To guard against the risk that aggregation reduces the reliability of my red flag measures, I build upon the approach I used to validate purchase-level data. Recall that this began with the assumption that a handful of factors should explain price changes over the course of an auction in a well-functioning procurement market. If we assume the same is true at the agency level for aggregate measures of their purchasing behavior, we can estimate the following model and use its residuals:

$$MedianAuctionEfficiency = \beta_0 + \beta_1 NumberOfPurchases + \beta_2 AgencySpend$$
$$+ \beta_3 MeanListingDuration$$
$$+ \beta_4 MeanBiddersApplied + \epsilon$$

If an agency's procurement is free of corruption, then we would expect auction efficiency (the median change between initial and winning price) to depend on the number of purchases made by an agency. The more it makes, the more its bureaucrats should become proficient, increasing auction efficiency. The same goes for agencies spending more, as this is another proxy for agency-level 'state capacity'. If an agency lists its purchases for longer on average, it gives more time for potential suppliers to make a bid, reducing final (winning) prices to a greater extent, captured by the mean number of bidders. I estimated this model for Moscow and for all regions (with fixed

effects), and found that the signs on the coefficients were as expected.

To the extent an agency deviates on average from this market model, its residual is greater. Suppose an agency's median auction efficiency is under-explained by the model above, and its residual is positive, corresponding to lower auction efficiency. This makes the agency a candidate for potential corruption, and if my measures are still reliable then they should collectively explain variance in this residual. I estimated a model for these cases with the (positive) residual as the dependent variable, and my three (aggregated) agency level red flag measures as independent variables, with no intercept. Increases in all three measures are positively associated with increases in an agency's residual from the market model above, giving further confidence that aggregation has not deadened their impact or introduced unwanted noise.[45]

4.3   How does corruption vary across regions?

My general approach in this study has been to test implications using data from Moscow, as the largest and most competitive regional market, and then replicate the analysis elsewhere. I continue that pattern here, estimating the models reported in Table 4.3 on the entire sample of agencies, including fixed effects for regions. I am interested in whether the conclusions about agency corruption made above, on the basis of Moscow data, hold generally within regions.

Figure 4.7 shows a similar (though not identical) pattern of average agency-level variation within each region.[46] The estimated effect on corruption of moving from least to most common good still depends on agency-level auction efficiency, even when the model estimates each region separately. Though the results weaken substantively,

---

[45]For completeness, I also estimated the market model at the purchase level, and regressed its positive residuals on the purchase-level flags, finding the same results in Moscow and across regions.

[46]The model is identical to (4) above, except for the fixed effects. All control variables are of the same sign and more precisely estimated, including agency spend which in Moscow was indistinguishable from zero. In the regional fixed effects model it has the expected sign: more spending is associated with a higher mean rate of disqualifications by the agencies in a given region. As with model (4) above, weighting agencies by purchases does not change the substantive results.

Effect of agency–level product commonness on 'disqualifications' corruption proxy
for different levels of corruption opportunity, all regions (region fixed effects)
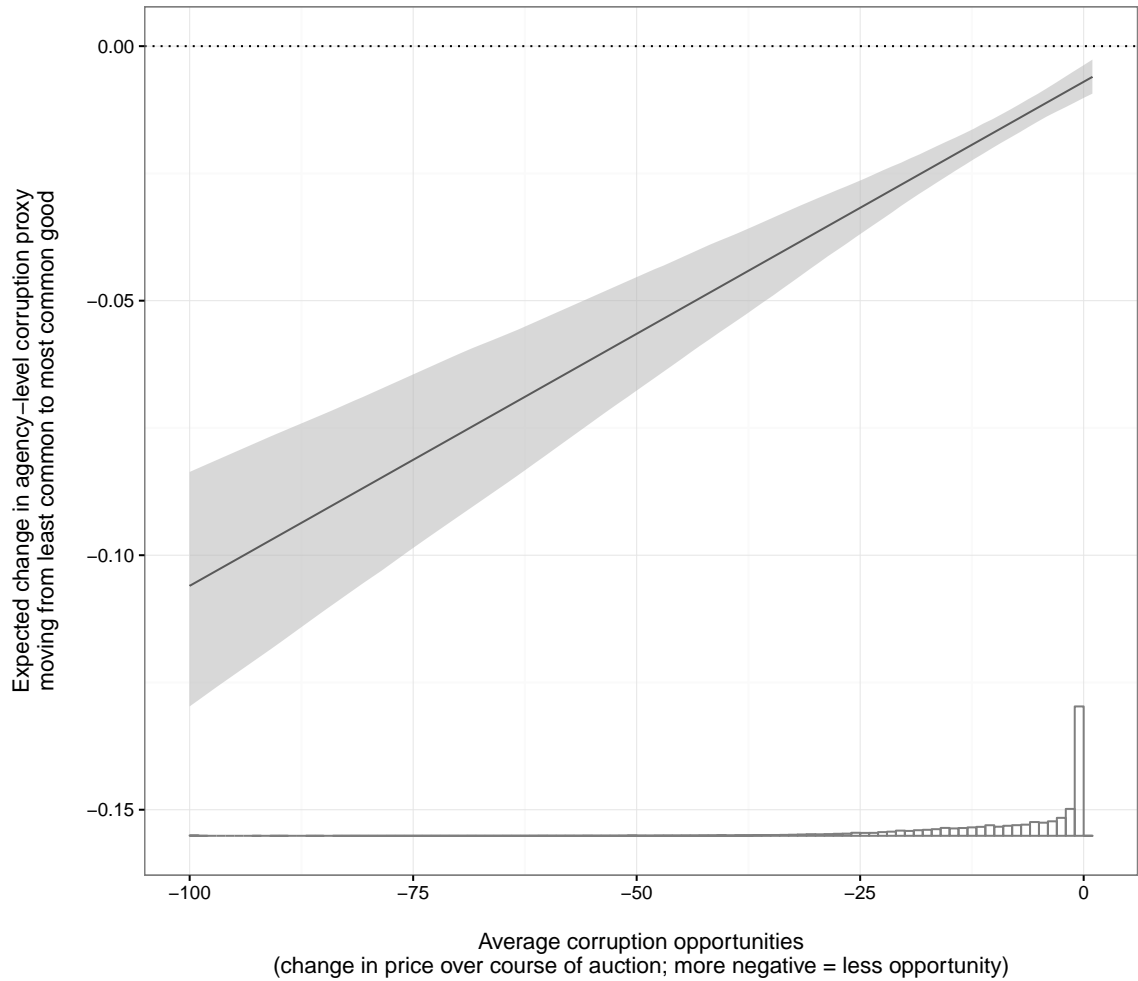
Figure 4.7: More common goods are associated with lower corruption, but only for
higher-efficiency auctions. The association weakens for lower-efficiency
auctions, consistent with my findings in Chapter 3.

it remains the case that for agencies with more efficient auctions (left), less common goods are more tightly associated with corruption than more common goods. In agencies with less efficient auctions (right) the association wears off. This supports one side of the conclusion about agency-level corruption drawn above. While the coefficient estimate does not cross above zero, the slope of the marginal effects line is the same: the less efficient are an agency's auctions, the less we observe an association between common goods and corruption. This builds further confidence that what I observed in Moscow was not peculiar to the capital, but a general pattern that might hold across Russia's regions, although to varying degrees.

4.3.1 Making full use of variation across Russia's regions

Having demonstrated a similar pattern within regions generally as seen in Moscow, I could plausibly leave the tests there. Fixed effects models are a blunt instrument, but the data survived, which I did not expect before estimation. As a blunt instrument, however, they ignore variation in agency behavior resulting from contextual factors: instead, all underlying heterogeneity is held hostage by the region-specific intercept. This presents a problem for drawing useful conclusions about regional variation, and an opportunity to take the study further. The problem is Russia's regions are clearly different when it comes to corruption: most region-specific intercepts are positive, some negative, and several are indistinguishable from zero. Despite this, we just saw that the predictions of my theory are broadly supported by the variation within these regions overall. The opportunity this presents is to model some of the regional variation directly, rather than leave it to a fixed effects model.

As noted earlier, the variation across Russia's regions not only helps explain agency-level corruption in that country, it can also inform wider debates. The chapter concludes by testing implications of my theory alongside those derived from prior literature in a unified empirical model, which nests agencies inside Russia's 79 regions

Map 4.1: Russian Regional GDP (PPP) in US dollars, 2013. The northern regions benefit from access to natural resources. Source: Wikimedia

with procurement data for this period.[47] As the richest region (Tyumen) has GDP per capita (PPP) equivalent to Norway, and the poorest region (Ingushetia) is equivalent to Iraq,[48] I am able to test implications of my agency-level theory, while accounting for variation across institutional contexts that has been the preserve of cross-national designs. Taking the regions as pseudo-states, I find that my within-state (agency-level) model is unperturbed by the introduction of region-level variables, even though regional variation does explain a significant portion of the agency-level variance. At the same time, estimates of several variables representing prominent cross-national theories have the expected sign based on prior literature. This offers an additional lens on old debates, and some assurance that the mechanisms of my agency-level theory add to the literature, rather than simply replacing what is already understood.

The dominant findings from the cross-national literature[49] are that corruption: *decreases* with increasing economic development; *decreases* with increasing political

---

[47]In future work, I would like to test three levels: purchases inside agencies inside regions.

[48]This variation is summarized in Map 4.1.

[49]Most clearly summarized by Treisman (2007) based on what remained robust in the years since his seminal earlier study (Treisman, 2000).

competition; *decreases* with increasing bureaucratic capacity; *decreases* with increasing media freedom; *increases* in reliance on unearned income (natural resources, transfers) for revenue; *decreases* with increased bureaucratic wages (efficiency wage).[50]

Thanks to the work of the *International Center for the Study of Institutions and Development* (ICSID),[51] good measures are available for all of these concepts across Russia for the years under study. This brings me to the question of variation across time, which I have avoided until now. The primary reason for collapsing all purchases for an agency in to one time period is that four years is simply insufficient to construct any meaningful panel. I can observe the detailed start and end dates of each purchase procedure but am not comfortable creating, for example, a 16-quarter or 48-month panel given the seasonality of agency procurement. Another reason is that the purchasing profiles of agencies do not change markedly over time.Finally, as the changes in these institutional variables would have a lagged affect on procurement, the main benefit of a panel (to make a causal case) would be neutralized.

Left with a choice to use one year, or average across them, I favor the former and choose 2011 as I can make a stronger case for these 'initial conditions' mattering for procurement over 2011-15.2011 is a good choice of 'baseline' for several reasons. First, it is the year the system I study came online. More consequentially, Russia's parliamentary elections in late 2011 were heavily contested, and *United Russia* (the party of then-President Medvedev, and then-once-and-future-President Putin) won by the narrowest of margins under allegations of widespread fraud.[52] This uncovers substantial variation in the dominance of the dominant party in Russia, which is often obscured in semi-authoritarian regimes: vote share for United Russia (UR) in Yaroslavl Oblast was 29%, while in Chechnya it was (allegedly) 99.48%.[53] The overall

---

[50]These hypotheses are common in the cross-national literature, and in this case are derived from Treisman (2000), Treisman (2007), Keefer (2007) and Svensson (2005).

[51]The ICSID is based at the Higher School of Economics in Moscow, whose Political Science department I was visiting in 2011-13. See: https://iims.hse.ru/en/csid/

[52]For detailed accounts see Bader and Ham (2015) and Kalinin (2016).

[53]That these shares were affected by fraud is problematic for explanations of election outcomes;

157
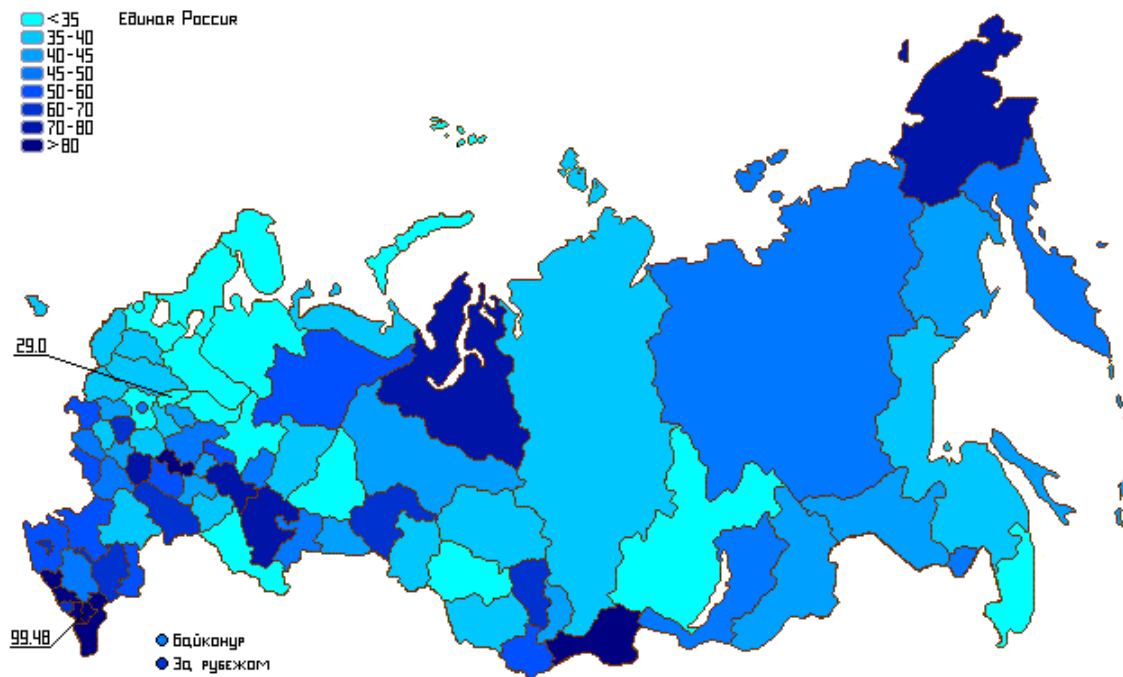
Map 4.2: United Russia's vote share by region in the 2011 parliamentary elections. Darker blue indicates greater share. Source: Wikimedia

pattern is shown in Map 4.2.

4.3.2   Testing within- and across-region variation in one model

Though contextual features may attenuate or accentuate aspects of my theory, I do not expect, given the robust evidence from several angles shown earlier, that this regional variation will fundamentally change the way agency-level corruption works. Nevertheless, the hypotheses from cross-national literature listed above provide clear empirical implications for how this could occur. If I have been wrong to this point, and macro-level institutional variables *can* explain better why corrupt states have clean agencies, this should become apparent in the following tests. Though I do not have a panel, 79 units is more than many of the leading cross-national studies,[54] and in most cases the ICSID measures of Russia's regions are at least as close to the underlying concept as their cross-national peers.

To test my theory's key implications along with those from the cross-national literature outlined above, I estimate this hierarchical mixed-effects effect model:

$$Corruption_{ar} = \beta_{0r} + \beta_{1ar} MedianProductCommonness + \beta_{2ar} MedianAuctionEfficiency$$

$$+ \beta_{3ar}(MedianProductCommonness \times MedianAuctionEfficiency)$$

$$+ B_{\mathbf{A}r} AgencyControls + B_{\mathbf{R}} RegionalVariables + \epsilon_{ar}$$

$$\beta_{0r} = \gamma_{00} + u_{0r}$$

$$\beta_{xar} = \gamma_{x0} + u_{xr}(Z_{ar})$$

The first equation specifies the fixed-effects relationships between agency-level independent variables and agency-level corruption for agency a in region r, and is specified the same as model (3) of Table 4.3 except for the addition of a vector of

as I use UR vote share as a proxy for regional political competition, the fact it is likely higher in every region due to fraud biases against the political competition hypothesis.

[54]Furthermore, in panel studies of corruption most variation is between units rather than across time, as both 'causing' and 'caused' variables change infrequently in these samples.

region-level variables $\mathbf{R}$ and the indexing of agency-level controls $\mathbf{A}$ by region. The second equation specifies a region-specific intercept for agency-level corruption, as a function of the mean level of corruption across all agencies in all regions γ_{00} and a region-specific deviation u_{0r}.[55] Finally, the third equation specifies the relationship between agency and regional levels for any given variable x, and is composed of an average slope across all agencies in all regions γ_{x0}, and a deviation u_{xr} that may (or may not)[56] depend on the level of some other variable Z in each agency a.

The objective of the succession of models estimated according to these equations, shown in Table 4.5, is to 'partial out' progressively more of the agency-level variation arising from the nesting of agencies inside regions. I estimate these models using my preferred corruption measure derived from a purchase-level 'red flag': the *mean proportion of bidders disqualified by an agency across its purchases*.[57] Model (1) reproduces the fixed effects model underlying Figure 4.7 as a baseline, and the corresponding marginal effects graph (not shown) evaluating the core implication of my theory is the same as Figure 4.7: less common goods, of which there are fewer suppliers (posing greater corruption risk), are only associated with greater corruption when an agency's average auction is efficient. The less efficient the average auction, the weaker the association. Model (2) allows the slopes of the key variables in my theory, *commonness* (median product commonness at agency level) and *efficiency* (median auction efficiency at agency level), as well as their interaction, to vary by region, whereas model (1) allowed only the intercepts to vary. As this does not change the unconditional effects of these two variables, it makes interpreting the change in the marginal effect easier: it strengthens. This means that accounting for how the core of my theory might vary by region sharpens the association between efficient auctions, more common goods, and less corruption. While this does not imply the converse (an

---

[55]This is in contrast to the region fixed-effects used in the model underlying Figure 4.7, which specified a completely separate intercept for each region.

[56]In this case, the random effect just allows for unconditionally varying slopes by region.

[57]The logic linking this observed behavior to corruption is clearest.

Table 4.5: Mixed-effects models of agency corruption across all regions

	Corruption (mean proportion disqualified by agency)			
	Baseline	Agency RE	Age RE/Reg FE	Age RE/Reg RE
	(1)	(2)	(3)	(4)
AgencySpendLog	.001***	.001***	.001***	.001***
	(.0003)	(.0003)	(.0003)	(.0003)
ListingDuration	.0001**	.0001	.0001*	.0001*
	(.00004)	(.00004)	(.00004)	(.00004)
BiddersApplied	.012***	.012***	.012***	.012***
	(.0002)	(.0002)	(.0002)	(.0002)
SupplierFavoritism	−.003***	−.003***	−.003***	−.003***
	(.0002)	(.0002)	(.0002)	(.0002)
Commonness	−.009***	−.008*	−.008*	−.009**
	(.002)	(.004)	(.004)	(.004)
Efficiency	.001***	.001***	.001***	.001***
	(.00003)	(.0001)	(.00005)	(.00005)
GRPConstantPrices			−.002***	−.002***
			(.001)	(.001)
PartyDominance			.00001	.0001
			(.0001)	(.0001)
TaxCapacity			−.010***	−.010***
			(.003)	(.003)
Newspapers			0.00000	0.00000
			(0.00000)	(0.00000)
GRPFromMining			.0003**	.0004***
			(.0002)	(.0001)
BureaucratWages			−.013	−.021*
			(.013)	(.013)
Commonness:Efficiency	.0005***	.001***	.001***	.001***
	(.0001)	(.0003)	(.0003)	(.0003)
Constant	.023***	.023***	.264***	.270***
	(.003)	(.003)	(.078)	(.074)
Random Effects				
# of Regions	79	79	79	79
- Regions std. dev.	0.014	0.015	0.014	0.069
Agency level predictors				
- Commonness		0.034	0.034	0.032
- Efficiency		0	0	0
- Commonness:Efficiency		0.002	0.002	0.002
Region level predictors				
- GRPConstantPrices				0.001
- PartyDominance				0
N	114,734	114,734	114,734	114,734

Standard errors in parentheses   *p < .1; **p < .05; ***p < .01

association between less efficient auctions, more common goods, and more corruption, as we saw at the purchase level in Moscow in 4.6), it does not contradict it.[58]

Model (3) adds several region-level independent variables to test implications of the wider corruption literature at the same time as my own agency-level theory. This model assumes the effect of these is constant across regions, i.e. none are introduced as random effects. In this model the estimates for variables testing my within-agency theory are fundamentally unchanged, but several of the regional-level variables are statistically significant explanators of the region-level variation that was constrained to the constant in models (1) and (2). The higher the level of economic development in a region, as represented by *GRPConstantPrices*,[59] the lower is average agency corruption. Political competition, measured by the 2011 parliamentary vote share of United Russia in *PartyDominance*, is not associated with regional variation in this corruption measure. There are few charitable interpretations of this finding: the model suggests bureaucratic corruption is invariant to the electoral fortunes of the dominant party, which is to be expected in a dominant-party regime like Russia.

A long-standing argument in comparative politics is that states with higher capacity[60] are more resistant to corruption. *TaxCapacity* measures the revenue-raising ability of Russia's regions, and is a proxy for regional-level bureaucratic capacity. In line with state capacity explanations of corruption, regions with greater ability to tax have lower average agency corruption. The argument that more extensive media penetration is associated with lower corruption is not supported by these data, as measured by *Newspapers*.[61] A prominent theory in political economy is that natural resource wealth has a deleterious effect on governance, particularly in less democratic

---

[58]It is possible that nesting purchases in agencies in regions, using all the variation in the data at once, could reproduce the very strong evidence seen in Moscow for both parts of the story: less corruption with more common goods in more efficient auctions (which these models reproduce), and more corruption with more common goods in less efficient auctions. I leave this to future work.

[59]This measure is in constant prices, so accounts for inflation.

[60]Defined variously as the power to take action, or the power to resist action taken by others.

[61]Defined as newspapers per 1000 residents. This is perhaps not surprising given the Kremlin's ever-tightening control over Russia's media landscape.

settings. *GRPFromMining* is the share of Gross Regional Product derived from the natural resources sector, and regions with higher shares have higher average agency corruption. The final variable in the model returns us to the start of Chapter 2, and the discussion of the efficiency wage hypothesis. *BureaucratWages* is the ratio of the average state employee's wage to the average wage of residents in a region, and captures the 'outside option' of the bureaucrats if they leave the bureaucracy. The measure ranges from just below 1.0 in Moscow, meaning bureaucrats are paid the same on average as everyone else, to almost 2.0 in Dagestan (a conflict-torn region in the Caucasus, neighboring Chechnya).[62] While no causal claims are possible with my research design, and the association here is only tentative, in model (4) on average regions with higher wages have lower average agency corruption, accounting for other factors. The final item of interest in model (3) is the marked change in the constant compared to model (2). This now represents average agency-level corruption across all regions, taking in to account all the other variables and allowing my core variables different slopes across regions.[63] In this model, the average agency disqualifies 26.4% of bidders in its average purchase.

The final model presented in this study (4), allows the slopes of two additional variables to vary by region: *GRPConstantPrices* and *PartyDominance*. These were selected because they account for contextual variation across regions not already scaled to other regional-level features.[64] A region's level of economic development changes a lot about the context of my theory, in particular the outside options for revenue- and rent-seeking firms engaging in procurement. And the level of political

---

[62]This makes intuitive sense, given the availability of labor in Moscow. It also makes intuitive sense that bureaucrats would be better compensated in remote, conflicted, and poverty-stricken regions. A concrete example of non-wage benefits for bureaucrats in remote regions was the 10% of the local agriculture census budget spent on a Toyota with wood paneling and premium audio system in Dagestan and Ingushetia. See: `http://transparency.org.ru/goszakupki/rosstat-zakupaet-dorogie-inomarki-dlia-selskokhoziaistvennoi-perepisi`

[63]This is the term γ_{00} in the second equation above.

[64]The other regional-level variables are already scaled: *TaxCapacity* is relative to Gross Regional Product, *Newspapers* is relative to population, *GRPFromMining* is obviously relative to GRP, and *BureaucraticWages* is relative to regional wage rates.

competition (or party dominance, in the case of United Russia) changes the incentives for regional whistleblowers and thus, potentially, the incentives of the actors in my theory. In this model, despite accounting for these region-level possibilities, the estimates of other parameters are unchanged from (3) except for *BureaucratWages*. Across all models, the association between auction efficiency, product commonness, and corruption, follows the pattern identified earlier and described for model (4) in Figure 4.8. The bounds of the confidence interval expand moving across models (1)-(4),[65] as the introduction of additional model features (controls, random effects) reduces the variance explained by this interaction. Nevertheless, the pattern remains.

The other important feature of the mixed effects framework is its partitioning of variance across the two levels, which can give some idea of the relative importance of my agency-level theory versus the existing cross-national wisdom, tested here by treating Russian regions as 'pseudo-states'. This is assessed by calculating the proportion of all variance across agencies that is explained by variation in the intercept (constant) of the model.[66] According to this model, just over 40% of the variance in agency-level corruption is explained by variation across regions, which demonstrates the importance of thinking about both levels at once. Importantly, as in Chapter 3 and earlier tests here, my core agency-level story is robust to regional differences; the estimated coefficients differ significantly from zero, in the predicted directions.

4.4   Discussion: the state of procurement in Russia

In this chapter I found empirical support for two explanations of clean agencies in corrupt states. Some clean agencies are buying more common goods in more efficient auctions, while others are buying less common goods in less efficient auctions. The former cases are intuitive and predicted by market logic: more common goods have

---

[65]Other plots not shown here, but observable in the increasing standard errors.

[66]In the second equation above, the intercept (constant) is β_{0r} and its region-specific variance is a function of u_{0r}.

Effect of agency–level product commonness on 'disqualifications' corruption proxy for different levels of corruption opportunity, all regions (mixed effects model (4))
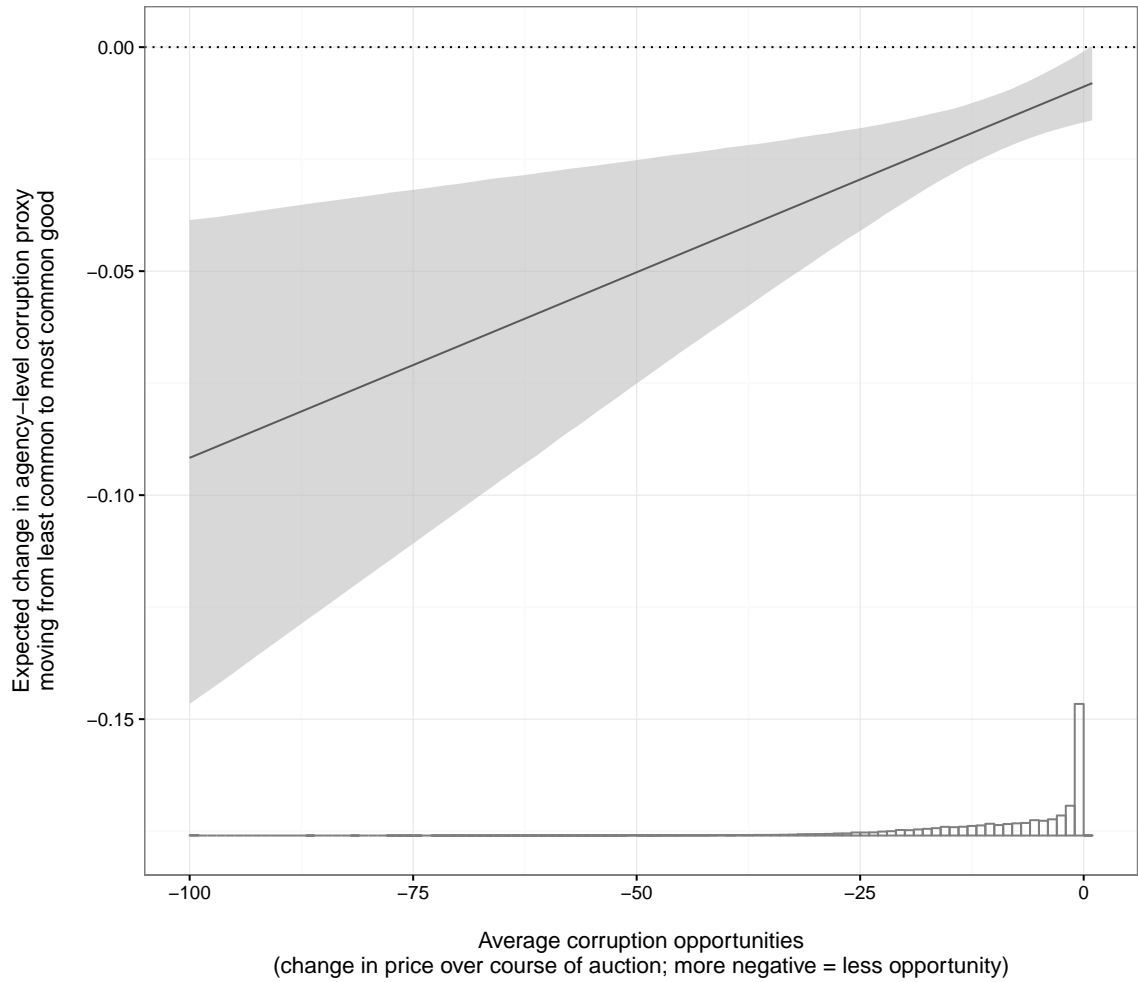
Figure 4.8: Taking in to account both agency- and regional-level sources of variation, more common goods are associated with lower corruption, but only for higher-efficiency auctions. The association weakens for lower-efficiency auctions. This is entirely consistent with findings in Chapter 3.

more suppliers, while efficient auctions unleash competition between them. The latter cases are counter-intuitive, but predicted by my theory: even though less common goods have fewer suppliers and weaker competitive pressures, corruption opportunities presented by less efficient auctions are not typically realized. If you accept that the more/less common distinction is a proxy for the information asymmetry between senior and junior bureaucrat in my theory, then in the latter cases it implies a clean agency is one where *corruption fails more than good government succeeds.*

Some interesting inferences can be drawn about the general state of procurement in Russia from 2011-2015, based on the evidence presented in this chapter. On average, procurement does work: the average purchase decreases in price over the auction and does not trigger my 'red flag' measures. Nevertheless, corruption affects a significant proportion of the purchases of many agencies, in patterns that are not explicable by other theories, and I offered compelling evidence that my three red flag measures detect this manipulation. The patterns of association uncovered between product commonness, auction efficiency, and corruption are broadly consistent across three very different levels of analysis and several techniques for linking them: 5-6 million individual purchases, 115,000 agencies, 79-82 regions. This increases confidence that the empirical support for my theory is general rather than tied to any one representation of these data. Finally, I found support for several implications of cross-national corruption research, tested at the regional level using variation across agencies as raw material. This demonstrated that while unobserved local differences matter for procurement corruption, there is a general regional pattern consistent with existing explanations, as well as an agency-specific pattern common to regions.

CHAPTER V

Conclusions and future work

"Everyone knew that the restaurant at Haji Ali, like every other illegal nightspot in town that faked a close, would reopen in less than half an hour. Everyone knew about the bribes that were paid and taken. Everyone knew about the warning phone calls. Everyone profited, and everyone was well pleased. The worst thing about corruption as a system of governance, Didier once said, is that it works so well."

Gregory David Roberts, Shantaram: A Novel[1]

"The Tsar is good; his advisors are the bad ones." Old Russian proverb

This dissertation has innovated in three ways, none of which is unique in itself, but which together create a project that makes several useful contributions to the wider literature.

5.1 Three innovations

First, my puzzle takes seriously an empirical regularity and theoretical concern, cross-agency variation in corruption, that many have noted but none have tackled to the extent I have attempted here. In one sense, it is trivially true that many

---

[1]Roberts (2004), p.186

phenomena varying across states are likely to vary within states. The question is whether that variation is systematically different at these two levels of analysis. If not, within-state variation in corruption is just another data source for testing the same very blunt theories of corruption from twenty years ago. The specific framing of my puzzle, 'why do corrupt states have clean agencies', makes clear that I believe within-state variation is not just another data source, but actually the foundation of what a somewhat blinkered measurement approach has observed cross-nationally. As Professor Clark once put it 'maybe low variation in corruption across agencies is the hallmark of good government'. I have noted this at a few points in my analysis.

The second innovation is in data. The fundamental problem studying corruption, like any illicit behavior, is the social science version of the Heisenberg uncertainty principle: if you can guarantee that what you observed was corruption, corruption already failed. I have joined a growing group of people who examine run-of-the-mill activity and build a theory about what should 'go wrong' (or right, from the corrupt bureaucrat's perspective) under corruption. Others have used administrative data from public procurement to study corruption, but none at the scale I do here. The variation in the data has given me flexibility to take several different approaches to testing parts of my theory. I can test the more ambitious implications of the model by aggregating to the agency level, which discards a lot of information but minimizes the impact of unobservable heterogeneity across purchases. I can use the full variation in the dataset, at the level of individual purchases, to test implications of my theory that depart less from existing literature. And because Russia's regions range from Norway to Iraq in terms of GDP PPP, I can employ the most robust of my measures to test within-'state' and across-'state' explanations at the same time. This was not easy, and the innocent paragraph describing how I wrote software to make sense of the raw purchase-level data represents two years of failure and learning.

Finally, my puzzle and these data have allowed me to push forward theory. Within

our field a lot of work has identified what is associated with corruption, often turning up the usual 'correlation of all good things'. In political economy the focus has been on identifying and testing specific interventions that aim to reduce corruption. I believe that we now know 'what works', which leaves us with a much tougher question for future work on within-state variation in behavior and outcomes: if we all know a certain brand of alarms deters bank robberies better than others, why don't all banks install this brand? I was unable to include as much of the politics of that choice in the present study as I hoped, but by taking as given that 'random audits work', I was forced to confront the question of where agency-level variation comes from. In my story *an agency is what it buys,* first and foremost. I have a theory for why agencies are different that can be operationalized given appropriate data, and an empirical approach built from my theory's expectations of behavior, rather than the endogenous features of agency structure others have used to explain my puzzle.

5.2   Wider implications and future work

I have demonstrated in this study the general importance of explaining variation in bureaucratic phenomena across agencies. Failure to do so has held back our understanding of what bureaucracies do, and not just in corruption; this blindness in turn hinders our understanding of politicians' design and monitoring choices. For example, an implication of my study, completely unrecognized in the literature, is that agencies present different levels of political risk to their elected overseers. This must change the calculus at the foundation of all politics: who gets what, when, and how. By examining a narrow period of time, over which design and monitoring choices are unlikely to have changed, I avoided this as a confound but left this potential additional explanation untapped. I challenge other scholars to take the existence proof of agency variation presented here, and my agency-centric explanation, and place it within a wider political context.

My findings challenge the implications of developmental and cultural explanations of corruption. By selecting one state, Russia, and modeling as much as possible its regional diversity, I avoided by design arguments that are frequently voiced in Russia and other corrupt states: 'corruption is in our culture.' I have shown that many agencies, staffed with people from the same culture, fail to extract all they can from the state budget. They fail not for a noble reason (a 'better organizational culture'), but as a result of their own greed, laying bare the functionalism of cultural arguments. These hold back scientific discovery and meaningful policy change in equal measure.

Explanations of corruption tied to economic development are likewise challenged by aspects of this study. In the lead-up to the period under study, Russia's economy grew remarkably on the back of windfall oil and gas revenues. While this obviously made the richest even richer, the benefits were relatively widely spread and officials' salaries in particular improved markedly. Bureaucratic pay increased further during the period under study as Putin's electoral authoritarianism sought reliable support. Yet despite tremendous growth in the productiveness of the golden goose that is the Russian economy, and a theoretical reduction in the temptation to extract due to efficiency wages, corruption never slowed. Instead it metastasized from widespread bribery of officials to 'get things done', in to the foundation of an entire system of rule, with procurement as a core institution for organizing the war over rents. This should be a further nail in the coffin of the naive hope of some international organizations that the path out of high-corruption equilibria is unbridled economic development. Without meaningful political competition, increased economic growth in Russia seems to have simply fueled the growth of ever more deep, dark and dangerous corruption. This entrenchment of corruption in procurement is a serious problem for Russia's leaders and citizens, and one they recognize well: that I have found plausible and systematic evidence of corruption, as well as variation across agencies, would surprise none of these actors.

Further support for this study's findings is provided by the regulatory changes at the end of the period under study, during which a new public procurement law was introduced. Pre-qualification of bidders was reintroduced, allegedly in response to rampant unjustified disqualification of bidders; revisions to details of purchases over their lifetime, especially prices, were restricted; the somewhat arbitrary thresholds around which rules changed were softened. In other words, the behavior generating each of my plausible and well-validated 'red flag' measures was seen as some kind of threat, and acted against. If we are optimistic, these changes were enacted to make corruption more difficult. If we are pessimistic, these changes simply upended the chess board to keep rent-seeking bureaucrats and firms on their toes.

Despite framing my puzzle in terms of corrupt states, there are wider lessons here. Corruption is still too often treated as a pathology of governance, that afflicts only developing states stuck in the 'wrong' equilibrium. Despite the focus of my study being this wrong equilibrium, there is nothing in my theory that cannot be applied to analyze bureaucratic decision-making in any context. Procurement is a key source of rents and corruption risk everywhere in the world, and I anticipate my theory would have useful explanatory power if procurement data in 'clean' countries were as open and accessible as it is in Russia.

Finally, those interested in minimizing procurement corruption, in any context, face important choices about which agencies to focus on and where in the process to cast their gaze. My measures are focused on what we can observe during the heat of the bidding, but the more decisive work to guarantee a desired outcome is often done beforehand, creating a *fait accompli* despite the outward appearance of competition. This is where effort should be focused in developed democracies with relatively robust bureaucratic accountability: as corruption during the process is much harder to get away with, firms need to exert influence much earlier. Detecting this kind of influence will be much more difficult than what I have achieved here.

APPENDIX

APPENDIX A

Notes on Russian procurement data

The data used in this study are downloaded legally from the Russian government procurement portal at `www.zakupki.gov.ru`. Complete instructions on downloading the data, and the code I use to parse it, clean it, and make it useful, are available online at `https://github.com/shaunmcgirr/shaun-mcgirr-dissertation/`

By following my instructions and using my code, any interested party can repeat my entire analysis and change it at any step. Other data comes from variety of sources, as cited in the bibliography and referred to in my online instructions.

Below I briefly describe the data quality steps taken to make the data ready for analysis, using Moscow as an example. Statistics for all regions are available online at the address above. I show the approximate percentage of cases lost with each successive decision.

1. Total number of purchase notifications downloaded from online system: 598,442

2. Removing duplicate leaves 550,375 notifications (8% of records discarded)

3. As I rely on comparing features from both notification and contract stages (e.g. price change), we are interested only in those from the 550k that match:

- Notifications with matching contract (core dataset): 399,696 (72.6%)

- Notifications without matching contract (aborted procedure): 148,740 (27.0%)

- Contracts without matching notification (non-competitive): 1,939 (0.4%)

4. Some of the 399,696 matched records describe purchases of several different kinds of product, or purchases for multiple agencies, or purchases from multiple suppliers, in the same procedure. I generally discard these because separating out the behavior for each is impossible. This removes 15.6% of the matched records, leaving 345,489 records available for the empirical models that use the most complicated measures of my variables of interest.

5. Some models, for less complicated measures using less of the detailed purchase-level data, can be estimated on 446,170 records (e.g. model (1) in Table 4.2; there disqualifications of bidders are relevant even in aborted procedures).

Further details on these decisions are available as comments in the code online.

BIBLIOGRAPHY

BIBLIOGRAPHY

Aidt, Toke S. 2003. "Economic analysis of corruption: a survey." *The Economic Journal* 113(491):F632–F652.
URL: *http://doi.wiley.com/10.1046/j.0013-0133.2003.00171.x*

Aidt, Toke S. and Jayasri Dutta. 2008. "Policy Compromises: Corruption and Regulation in a Democracy." *Economics & Politics* 20(3):335–360.
URL: *http://doi.wiley.com/10.1111/j.1468-0343.2008.00332.x*

Andersson, Staffan and Paul M. Heywood. 2009. "The Politics of Perception: Use and Abuse of Transparency International's Approach to Measuring Corruption." *Political Studies* 57(4):746–767.
URL: *http://psx.sagepub.com/lookup/doi/10.1111/j.1467-9248.2008.00758.x*

Bader, Max and Carolien van Ham. 2015. "What explains regional variation in election fraud? Evidence from Russia: a research note." *Post-Soviet Affairs* 31(6):514–528.

Balsevich, Anna and Elena Podkolzina. 2014. "INDICATORS OF CORRUPTION IN PUBLIC PROCUREMENT : THE EXAMPLE OF RUSSIAN REGIONS." *BASIC RESEARCH PROGRAM, WORKING PAPERS SERIES: ECONOMICS* (WP BRP 76/EC/2014).
URL: *http://www.hse.ru/en/org/hse/wp/prepfr_EC*

Bashina, Alexandra. 2013. "Factors related to informal payments in public procurement in Russia." *International Journal of Economic Policy in Emerging Economies* 6(4):403–430.
URL: *http://www.inderscience.com/link.php?id=57907*

Bauhr, Monika. 2016. "Need or Greed? Conditions for Collective Action against Corruption." *Governance* 00(00):Early View.
URL: *http://doi.wiley.com/10.1111/gove.12232*

Becker, Gary S. 1968. "Crime and Punishment: An Economic Approach." *The Journal of Political Economy* 76(2):169 – 217.
URL: *http://www.jstor.org/stable/1830482*

Becker, G.S. and G.J. Stigler. 1974. "Law Enforcement, Malfeasance, and Compensation of Enforcers." *J. Legal Stud.* 3(1):1–18.
URL: *http://www.jstor.org/stable/724119*

Block, F. 1977. "The ruling class does not rule: notes on the Marxist theory of the state." *Socialist Revolution* 77((May-June)).

Brambor, T., William Roberts Clark and Matt Golder. 2005. "Understanding Interaction Models: Improving Empirical Analyses." *Political Analysis* 14(1):63–82.
URL: *http://pan.oxfordjournals.org/cgi/doi/10.1093/pan/mpi014*

Bueno De Mesquita, Bruce, A. Smith, R.M. Siverson and J.D. Morrow. 2005. "The logic of political survival." *MIT Press Books* .
URL: *http://econpapers.repec.org/bookchap/mtptitles/0262524406.htm*

Cadot, O. 1987. "Corruption as a Gamble." *Journal of Public Economics* 33(2):223–244.
URL: *http://linkinghub.elsevier.com/retrieve/pii/0047272787900752*

Campos, Nauro F. and Francesco Giovannoni. 2006. "Lobbying, corruption and political influence." *Public Choice* 131(1-2):1–21.
URL: *http://www.springerlink.com/index/10.1007/s11127-006-9102-4*

Carpenter, Daniel P. 2000. "State Building through Reputation Building: Coalitions of Esteem and Program Innovation in the National Postal System, 1883–1913." *Studies in American Political Development* 14(2):121–155.
URL: *http://www.journals.cambridge.org/abstract_S0898588X00003382*

Charron, Nicholas, Carl Dahlström and Victor Lapuente. 2016. "Measuring Meritocracy in the Public Sector in Europe: a New National and Sub-National Indicator." *European Journal on Criminal Policy and Research* pp. 1–25.
URL: *http://link.springer.com/10.1007/s10610-016-9307-0*

Clark, WR and S McGirr. 2010. Asset Specificity, Inequality, and Limited Government. In *APSA 2010 Annual Meeting Paper*. Number September.
URL: *http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1643535*

Compte, O, Ariane Lambert-Mogiliansky and T. Verdier. 2005. "Corruption and competition in procurement auctions." *Rand Journal of Economics* 36(1):1–15.
URL: *http://www.jstor.org/stable/10.2307/1593751*

Coppedge, Michael, John Gerring, Staffan I. Lindberg, Svend-Erik Skaaning, Jan Teorell, David Altman, Michael Bernhard, M. Steven Fish, Adam N Glynn, Allen Hicken, Carl Henrik Knutsen, Kyle L. Marquardt, Kelly McMann, Farhad Miri, Pamela Paxton, Daniel Pemstein, Jeffrey Staton, Eitan Tzelgov, Yi-ting Wang and Brigitte Zimmerman. 2016. "V-Dem [Country-Year/Country-Date] Dataset v6.1.".
URL: *https://v-dem.net/en/data/data-version-6-1/*

Corruption Research Center Budapest. 2015. "No Public Procurement Corruption in Hungary? Analysis of the Decisions of the Hungarian Court between 2009-2014.".
URL: *http://www.crcb.eu/?p=832*

Dawisha, Karen. 2014. *Putin's kleptocracy : who owns Russia?* 1 ed. New York: Simon & Schuster.

Di Tella, Rafael and Ernesto Schargrodsky. 2003. "The Role of Wages and Auditing during a Crackdown on Corruption in the City of Buenos Aires*." *The Journal of Law and Economics* 46(1):269–292.
URL: *http://www.journals.uchicago.edu/doi/abs/10.1086/345578*

Escresa, Laarni and Lucio Picci. 2015. "A New Cross-National Measure of Corruption." *World Bank Policy Research Working Paper No. 7371* .
URL: *http://wber.oxfordjournals.org/content/early/2015/07/23/wber.lhv031.abstract*

Fazekas, Mihály, István János Tóth and Lawrence P. King. 2013. "Corruption Manual for Beginners: 'Corruption Techniques' in Public Procurement with Examples from Hungary." *SSRN Electronic Journal* .
URL: *http://www.ssrn.com/abstract=2333354*

Fazekas, Mihály, István János Tóth and Lawrence Peter King. 2016. "An Objective Corruption Risk Index Using Public Procurement Data." *European Journal on Criminal Policy and Research* pp. 1–29.
URL: *http://link.springer.com/10.1007/s10610-016-9308-z*

Ferejohn, John. 1986. "Incumbent performance and electoral control." *Public Choice* 50(1-3):5–25.
URL: *http://link.springer.com/10.1007/BF00124924*

Ferraz, Claudio and Frederico Finan. 2008. "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes *." *Quarterly Journal of Economics* 123(2):703–745.
URL: *http://www.mitpressjournals.org/doi/abs/10.1162/qjec.2008.123.2.703*

Gailmard, Sean and John W. Patty. 2007. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51(4):873 – 889.
URL: *http://www.jstor.org/stable/4620105*

Gel'man, Vladimir. 2000. "The Dictatorship of Law in Russia: Neither Dictatorship, Nor Rule of Law.".
URL: *https://www.gwu.edu/~ieresgwu/assets/docs/ponars/pm_0146.pdf*

Gerring, John. 2001. *Social science methodology: a criterial framework.* Cambridge: Cambridge University Press.

Gingerich, Daniel W. 2013. "Governance Indicators and the Level of Analysis Problem: Empirical Findings from South America." *British Journal of Political Science* 43(03):505–540.
URL: *http://www.journals.cambridge.org/abstract_S0007123412000403*

Grzymala-Busse, A.M. 2007. *Rebuilding Leviathan: Party competition and state exploitation in post-communist democracies.* Cambridge Univ Pr.
URL: *https://books.google.com/books?isbn=1139464922*

Grzymala-Busse, Anna. 2010. "The Best Laid Plans: The Impact of Informal Rules on Formal Institutions in Transitional Regimes." *Studies in Comparative International Development* 45(3):311–333.
URL: *http://www.springerlink.com/index/10.1007/s12116-010-9071-y*

Huber, John D and Charles R Shipan. 2002. *Deliberate discretion : the institutional foundations of bureaucratic autonomy.* Cambridge, UK: Cambridge University Press.

Hunt, Jennifer. 2005. "Why are some public officials more corrupt than others?" *NBER Working Paper* (11595).
URL: *http://www.nber.org/papers/w11595*

Ichino, N, JK Faller and A Glynn. 2013. "Electoral Systems and Corruption." *APSA 2013 Annual Meeting Paper* .
URL: *http://people.fas.harvard.edu/~nichino/*

Ishii, Rieko. 2009. "Favor exchange in collusion: Empirical study of repeated procurement auctions in Japan." *International Journal of Industrial Organization* 27(2):137–144.

Kalinin, Kirill. 2016. "Validating Precinct-Level Measures of Fraud: Evidence from the Russian Electoral Cycle 2011-2012." *SSRN Electronic Journal* .
URL: *http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2767713*

Keefer, Philip. 2007. "Clientelism, Credibility, and the Policy Choices of Young Democracies." *American Journal of Political Science* 51(4):804–821.
URL: *http://doi.wiley.com/10.1111/j.1540-5907.2007.00282.x*

Kelly, William A. 1981. "A Generalized Interpretation of the Herfindahl Index." *Southern Economic Journal* 48(1):50–57.
URL: *http://www.jstor.org/stable/1058595*

Klašnja, Marko and Joshua a. Tucker. 2013. "The economy, corruption, and the vote: Evidence from experiments in Sweden and Moldova." *Electoral Studies* 32(3):536–543.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0261379413000553*

Kramer, W. Michael. 2012. "The most common procurement fraud schemes and their primary red flags.".
URL: *http://iacrc.org/procurement-fraud/the-most-common-procurement-fraud-schemes-and-their-primary-red-flags/*

Krueger, A.O. 1974. "The political economy of the rent-seeking society." *The American Economic Review* 64(3):291–303.
URL: *http://www.jstor.org/stable/1808883*

Laakso, Markku and Rein Taagepera. 1979. "The "Effective" Number of Parties: A Measure with Application to West Europe." *Comparative Political Studies* 12(1):3–27.
URL: *http://cps.sagepub.com/content/12/1/3.full.pdf+html*

Lambert-Mogiliansky, Ariane. 2011. Corruption and Collusion: Strategic Complements in Procurement. In *International Handbook on the Economics of Corruption, Volume 2*, ed. Susan Rose-Ackerman and Tina Søreide. Northampton, MA: Edward Elgar chapter 4, p. 587.

Lambert-Mogiliansky, Ariane and Konstantin Sonin. 2006. "Collusive market sharing and corruption in procurement." *Journal of Economics & Management Strategy* 15(4):883–908.
URL: *http://onlinelibrary.wiley.com/doi/10.1111/j.1530-9134.2006.00121.x/full*

Lambsdorff, J.G. 2002. "Corruption and rent-seeking." *Public Choice* 113(1):97–125.
URL: *http://www.springerlink.com/index/J11U6548X515V739.pdf*

Lui, Francis T. 1986. "A dynamic model of corruption deterrence." *Journal of Public Economics* 31(2):215–236.
URL: *http://linkinghub.elsevier.com/retrieve/pii/0047272786900198*

Mauro, Paolo. 1995. "Corruption and Growth." *The Quarterly Journal of Economics* 110(3):681 – 712.
URL: *http://www.jstor.org/stable/2946696*

McMann, Kelly M., Daniel Pemstein, Brigitte Seim, Jan Teorell and Staffan I. Lindberg. 2016. "Strategies of Validation: Assessing the Varieties of Democracy Corruption Data." *V-Dem Working Paper* 2016(23).
URL: *http://www.ssrn.com/abstract=2727595*

Menezes, Flavio M. and Paulo Klinger Monteiro. 2006. "Corruption and auctions." *Journal of Mathematical Economics* 42(1):97–108.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0304406805000698*

Méon, Pierre-Guillaume and Laurent Weill. 2010. "Is Corruption an Efficient Grease?" *World Development* 38(3):244–259.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0305750X09001193*

Mohtadi, H. 2003. "Democracy, rent seeking, public spending and growth." *Journal of Public Economics* 87(3-4):445–466.
URL: *http://dx.doi.org/10.1016/S0047-2727(01)00135-9*

Morozov, Ilya and Elena A. Podkolzina. 2013. "Collusion Detection in Procurement Auctions." *Higher School of Economics Research Paper No. WP BPR 25/EC/2013*.
URL: *http://www.ssrn.com/abstract=2221809*

Myerson, RB. 1993. "Effectiveness of Electoral Systems for Reducing Government Corruption: A Game-Theoretic Analysis." *Games and Economic Behavior* 5(1):118–132.
URL: *http://linkinghub.elsevier.com/retrieve/doi/10.1006/game.1993.1006*

Niehaus, Paul and Sandip Sukhtankar. 2013. "Corruption Dynamics: The Golden Goose Effect." *American Economic Journal: Economic Policy* 5(4):230–269.
URL: *http://pubs.aeaweb.org/doi/abs/10.1257/pol.5.4.230*

Nye, J.S. 1967. "Corruption and political development: a cost-benefit analysis." *The American Political Science Review* 61(2):417–427.
URL: *http://www.jstor.org/stable/1953254*

OECD. 2016. Corruption in Public Procurement. Technical report Secretary-General.
URL: *http://www.oecd.org/gov/ethics/Corruption-in-Public-Procurement-Brochure.pdf*

Olken, BA. 2007. "Monitoring corruption: evidence from a field experiment in Indonesia." *Journal of Political Economy* 115(2):200–249.
URL: *http://www.journals.uchicago.edu/doi/abs/10.1086/517935*

Olken, Benjamin A. 2009. "Corruption perceptions vs. corruption reality." *Journal of Public Economics* 93(7-8):950–964.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0047272709000310*

Olken, Benjamin A and Rohini Pande. 2012. "Corruption in Developing Countries." *Annual Review of Economics* 4(1):479–509.
URL: *http://www.annualreviews.org/doi/abs/10.1146/annurev-economics-080511-110917*

Olson, M. 1993. "Dictatorship, democracy, and development." *American Political Science Review* 87(3):567–576.
URL: *http://www.jstor.org/stable/10.2307/2938736*

Ostrovnaya, M. and E. Podkolzina. 2015. "ANTITRUST ENFORCEMENT IN PUBLIC PROCUREMENT: THE CASE OF RUSSIA." *Journal of Competition Law and Economics* 11(2):331–352.
URL: *http://jcle.oxfordjournals.org/cgi/doi/10.1093/joclec/nhv013*

Persson, T., G.E. Tabellini and I. Brocas. 2000. *Political economics.* MIT press.
URL: *http://www.econ.cam.ac.uk/faculty/aidt/policy/2010/outline2010.pdf*

Plautus, Titus Maccius. 1992. *Amphitryon - The Comedy of asses - The Pot of gold - The Two bacchises - The Captives.* 10 ed. Cambridge, MA: Harvard University Press.

Porter, RH and JD Zona. 1999. "Ohio school milk markets: an analysis of bidding." *The RAND Journal of Economics* 30(2):263–288.
URL: *http://www.jstor.org/stable/10.2307/2556080*

Recanatini, Francesca, Alessandro Prati and Guido Tabellini. 2005. "Why Are Some Public Agencies less Corrupt than Others? Lessons for Institutional Reform from Survey Data." *Working Paper* October.
URL: *http://www.imf.org/external/np/res/seminars/2005/arc/pdf/reca.pdf*

Roberts, Gregory. 2004. *Shantaram : a novel.* 1 ed. New York: St. Martin's Press.

Rose-Ackerman, S. 2006. *International handbook on the economics of corruption.* Cheltenham, UK; Northampton, MA: Edward Elgar.

Rose-Ackerman, Susan. 1999. *Corruption and government.* Cambridge: Cambridge University Press.
URL: *http://ideas.repec.org/b/cup/cbooks/9780521632935.html*

Rothstein, Bo and Jan Teorell. 2008. "What Is Quality of Government? A Theory of Impartial Government Institutions." *Governance* 21(2):165–190.
URL: *http://doi.wiley.com/10.1111/j.1468-0491.2008.00391.x*

Russian Federation. 2011. "Federal Law No. 94-FZ "On State and Municipal Procurement of Goods, Works and Services" (Official English translation).".
URL: *http://en.fas.gov.ru/documents/documentdetails.html?id=13920*

Scott, James C. 1969. "The Analysis of Corruption in Developing Nations." *Comparative Studies in Society and History* 11(03):315–341.
URL: *http://www.journals.cambridge.org/abstract_S0010417500005363*

Shamrin, Alexander and Andrei A Yakovlev. 2009. "Analysis of the Public Procurement System in Russia and Development of Recommendations on its Improvement.".
URL: *https://www.hse.ru/en/org/projects/13315037*

Shepsle, K. A. 1989. "Studying Institutions: Some Lessons from the Rational Choice Approach." *Journal of Theoretical Politics* 1(2):131–147.
URL: *http://jtp.sagepub.com/cgi/doi/10.1177/0951692889001002002*

Sundell, Anders. 2014. "ARE FORMAL CIVIL SERVICE EXAMINATIONS THE MOST MERITOCRATIC WAY TO RECRUIT CIVIL SERVANTS? NOT IN ALL COUNTRIES." *Public Administration* 92(2):440–457.
URL: *http://doi.wiley.com/10.1111/padm.12077*

Svensson, J. 2005. "Eight questions about corruption." *Journal of Economic Perspectives* 19(3):19–42.
URL: *http://www.atypon-link.com/AEAP/doi/abs/10.1257/089533005774357860*

Svolik, Milan. 2008. "Authoritarian Reversals and Democratic Consolidation." *American Political Science Review* 102(02):153–168.
URL: *http://www.journals.cambridge.org/abstract_S0003055408080143*

The World Bank. 2006. Russian Federation - Country Procurement Assessment Report (A Limited Update). Technical Report December.
URL: *siteresources.worldbank.org/INTRUSSIANFEDERATION/Resources/cpar2007eng.pdf*

The World Bank. 2010. Most Common Red Flags of Fraud and Corruption in Procurement in Bank-Financed Projects. Technical report Integrity Vice Presidency.
URL: *http://pubdocs.worldbank.org/en/324351449169329391/Red-Flags.pdf*

Tirole, Jean. 1986. "Hierarchies and bureaucracies: On the role of collusion in organizations." *Journal of Law, Economics, & Organization* 2(2):181–214.
URL: *http://www.jstor.org/stable/765048*

Tirole, Jean. 1994. "The internal organization of government." *Oxford Economic Papers* 46(1):1–29.
URL: *http://www.jstor.org/stable/10.2307/2663521*

Tollison, R.D. 1982. "Rent seeking: A survey." *Kyklos* 35(4):575–602.
URL: *http://onlinelibrary.wiley.com/doi/10.1111/j.1467-6435.1982.tb00174.x/pdf*

Transparency International. 2013. "Global Corruption Barometer.".
URL: *http://www.transparency.org/research/gcb/*

Treisman, Daniel. 2000. "The causes of corruption: a cross-national study." *Journal of Public Economics* 76(3):399–457.
URL: *http://dx.doi.org/10.1016/S0047-2727(99)00092-4*

Treisman, Daniel. 2007. "What Have We Learned About the Causes of Corruption from Ten Years of Cross-National Empirical Research?" *Annual Review of Political Science* 10(1):211–244.
URL: *http://arjournals.annualreviews.org/doi/abs/10.1146/annurev.polisci.10.081205.095418*

Tullock, G. 1967. "The welfare costs of tariffs, monopolies, and theft." *Economic Inquiry* 5(3):224–232.
URL: *http://onlinelibrary.wiley.com/doi/10.1111/j.1465-7295.1967.tb01923.x/abstract*

Tullock, Gordon. 1990. The costs of special privilege. In *Perspectives on Positive Political Economy*, ed. James E. Alt and Kenneth A. Shepsle. Cambridge: Cambridge University Press chapter 8, p. 195–211.
URL: *https://books.google.com/books?isbn=0521398517*

Williamson, OE. 1981. "The Economics of Organization: The Transaction Cost Approach." *American Journal of Sociology* 87(3):548–577.
  URL: *http://www.jstor.org/stable/10.2307/2778934*

Wilson, James. 1989. *Bureaucracy: what government agencies do and why they do it.* New York: Basic Books.
  URL: *http://www.worldcat.org/title/bureaucracy-what-government-agencies-do-and-why-they-do-it/oclc/19920848&referer=brief_results*

Winters, Matthew S. and Rebecca Weitz-Shapiro. 2013. "Lacking Information or Condoning Corruption: When Do Voters Support Corrupt Politicians?" *Comparative Politics* 45(4):418–436.
  URL: *http://openurl.ingenta.com/content/xref?genre=article&issn=0010-4159&volume=45&issue=4&spage=418*

Yakovlev, Andrei. 2010. "The system of public procurements in Russia: on the road of reform.".
  URL: *https://www.hse.ru/data/2010/10/12/1223048786/Public procurement in Russia.pdf*

Yakovlev, Andrei. 2012. "From Competition to Dominance: Political Determinants of Federal Transfers in the Russian Federation." *Working Paper* March.
  URL: *http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2005710*