The Structure of Costs and
the Duration of Supplier Relationships[*]

Alexander MacKay[†]

January 12, 2016

JOB MARKET PAPER
[Download the latest version](#)

Abstract

The duration of a supply relationship depends on two types of costs: (i) the transaction cost of switching suppliers and (ii) the cost of being matched to an inefficient supplier when the relationship lasts too long. I develop a model of optimal contract duration that captures this tradeoff, and I provide conditions that identify the underlying cost structure. Latent transaction costs are identified even when the exact supplier selection mechanism is unknown. For a typical procurement good, I estimate the model and find that transaction costs are a significant portion of total costs. I conduct two counterfactual exercises. In the first, I estimate the effects of changing the maximum allowable contract duration, which is a common contracting friction. In the second, I evaluate the impact of reducing transaction costs. This second counterfactual illustrates why quantifying transaction costs is important for accurate welfare analysis.

1 Introduction

When buyers select sellers, they select not only who but also how long. Supply relationships govern a great portion of transactions for commodities, inputs, and services. The duration of these relationships varies within and across sectors. For example, contracts for industrial goods such as aluminum, paper, and glass range from three to twelve months,[1] whereas utilities purchase residential electricity in contracts of one to ten years. Contracts of varying duration are typical for building leases, hospital and health insurer contracts, and service contracts for nurses, security guards, and janitors, among other goods.

The prevalence of long-term supply relationships in these markets suggests that transaction costs are meaningful, as the products share similar features to those sold efficiently in spot markets.[2] Transaction costs[3] can constitute a large fraction of total costs and, through supply relationships, mediate how prices respond to changes in the economic environment. Thus, transaction costs may be of first-order concern when conducting welfare analysis. Yet quantification of these costs has often proved difficult, and measures of transaction costs are few.

In this paper, I develop a model of optimal contract duration that identifies latent transaction costs. The optimal duration depends on two types of costs: the transaction cost of re-selecting a supplier, and the opportunity cost of being matched to an inefficient supplier for part of the duration of the contract. Even while maintaining an agnostic approach to how suppliers are selected, transaction costs are identified. I then examine the case of procurement auctions, which are the primary form of competitive bidding for supply contracts. The added structure from auction theory allows for the identification of the joint distribution of supply costs.

Using a specified version of the model, I perform a within-industry analysis of a typical procurement good: janitorial contracts. I find that transaction costs are meaningful in this setting, comprising 10 percent of overall costs. I conduct a policy counterfactual in which I change the cap on maximum contract duration, which is a common contract friction. I find that the existing five-year maximum is close to the optimal (capless) policy, but that shorter caps can be costly. In a second counterfactual, I consider the consequences of omitting transaction costs from welfare analysis. For a policy that reduces transaction costs by 20 percent, an analysis of prices would capture only 56 percent of the effect on total costs.

For some intuition behind the tradeoff between transaction costs and an inefficient match, consider the case of university rates at local hotels. Universities negotiate annually with hotels to make rooms available at a special rate. Negotiations are costly; universities would often prefer to pay a slightly higher rate for a two-year term instead. However, hotels face large variation in year-to-year demand on account of conferences and travel shocks. For a two-year deal, a hotel would charge a significant premium to cover potential high future demand, even with low demand in the current

---

[1]In a seminal survey, Stigler and Kindahl (1970) found that about half of the industrial goods in their sample are purchased with explicit contracts.

[2]For the examples given, the products are homogenous, performance is observable, and there are many suppliers. Indeed, some products are sold both in spot markets and via long-term contracts.

[3]Transaction costs include switching costs and search costs.

year. Annual contracts strike a balance between costly negotiations and lower prices, as universities can select efficiently among hotels with available capacity in that year.

To illustrate how supply costs and transaction costs determine optimal duration, I develop a stylized two-period model in Section 2. The model shows how key features of the economic environment - variance in supply costs and the number of suppliers - affect the optimal duration. For example, greater autocorrelation in supplier costs leads to longer contracts, as the identity of the low-cost supplier is more persistent. In the case of perfect autocorrelation, the low-cost supplier does not change. Intuitively, the more similar the future is to today, the greater the benefit of long-term contracts.

Though autocorrelation and duration have a straightforward relation, other features of the economic environment have non-monotonic or ambiguous effects. In the simple model, the optimal duration is U-shaped in competition, where competition is indexed by the number of suppliers bidding on the contract. With few suppliers, the procurer chooses a long-term contract, as the benefit of re-selecting a supplier is small. With greater competition, the low-cost supplier changes more frequently,[4] increasing the benefit of shorter contracts. When competition is very intense, supply costs are sufficiently low that the benefit of switching suppliers is small relative to transaction costs.

As the simple model demonstrates, the optimal duration policy is inherently empirical. This motivates the development of the structural empirical model in Section 3. The model consists of three stages. First, the buyer determines contract duration after observing contract-specific characteristics and supplier participation cost shifters. Second, suppliers decide to participate in a supplier selection mechanism (e.g. an auction) after observing duration, contract characteristics, and participation costs. Third, the selection mechanism determines the price and the selected supplier.

Starting with a general version of the duration-setting problem, I provide a set of assumptions under which key structural components, including transaction costs and duration-dependent prices, are identified. Transaction costs are identified even when maintaining an agnostic approach to how suppliers are selected.[5] This result is useful for analyzing contracts where the underlying selection process may be obscure, such as contract negotiations between television providers (Comcast) and content producers (ABC Disney). The identified components are needed to evaluate policies that impact contract duration, as well as those that affect entry and participation by suppliers.

When the means of supplier selection is known, additional components of the model are identified. I consider the case of the procurement auction, which is perhaps the most common form of competitive bidding and maps to the empirical setting of my application. In Section 3.3.3, I demonstrate that an auction model with unobserved heterogeneity and endogenous entry is identified when only the winning bid is observed. This is a practical contribution to the auction identification literature, as the result depends on readily available data and allows for a particular form of selective entry.

---

[4]For a given process of supply costs.

[5]That is, the result does not depend on the auction framework, which I implement later.

For the auction model, identification of private costs and unobserved heterogeneity generalizes to auction models beyond the one in my paper, as the duration-setting problem is not used in the proof. Moreover, when there is no selection on unobservables, the model with unobserved heterogeneity is identified even without explicitly modeling entry. Identification is obtained by using exogenous variation in the number of bidders - either through participation cost shifters or by assumption - to disentangle the private cost component from the unobserved heterogeneity. Intuitively, variation in the number of bids shifts the distribution of the private component in a known way, while the distribution of auction-specific heterogeneity is unaffected. Previous approaches to this challenging problem relied on observing either multiple bids per auction or a reservation price.

For the empirical work of the paper, I have collected a detailed dataset for a typical procurement good: janitorial services. Janitorial services are a significant industry in the United States; federal government contracts, from which my sample is collected, comprise $500 million annually. In Section 4, I describe the data and provide some motivating evidence for the assumptions of the structural analysis.

Janitorial contracts are well-suited for understanding the optimal duration problem. First, they are determined by sealed-bid offers, which maps into the auction model of Section 3. Second, they are a homogenous good, which allows for a direct comparison across contracts. Third, transaction costs are meaningful, as the government faces significant labor costs when re-competing each contract.[6] Fourth, the contracts are issued in simple fixed-duration, fixed-price terms.[7] Finally, there is substantial variation in duration, as the contracts vary from three months to five years. The results from this particular study may provide some insight to procurement costs more generally, as the contracting officers responsible for janitor contracts also procure other goods.[8]

In Section 5, I present the main empirical findings of my paper. The empirical strategy proceeds in two stages. In the first stage, participation and bidding are jointly estimated via maximum likelihood. In the second stage, the optimality condition of the duration-setting problem is combined with the estimated parameters from the first stage to construct distribution-free bounds on transaction costs.

Once these bounds are obtained, I apply a prior over the bounds in order to estimate transaction costs and construct counterfactuals. I project transaction costs onto variables outside of the model, and I find that transaction costs are higher for locations with a greater overall volume of procurement, which indicates an upward-sloping supply curve rather than economies of scale.

The contracts in my dataset exhibit a typical contracting friction: there is a cap on contract duration of five years. Forty-three percent of the contracts are at the cap. In a counterfactual analysis, I find that the cap is close to the optimal policy. However, shorter caps would be costly. A three-year cap would increase total costs by 2.2 percent.

---

[6]One senior contracting officer estimated that issuing a simple contract would take about three weeks of full-time work for a government employee.

[7]Additionally, ex post incentive problems, which are a large focus of the contract literature, are not a first-order concern. The suppliers in this market are generally well-established firms.

[8]In fiscal year 2014, the federal government issued over one million fixed-term, fixed-price competitive contracts.

Finally, I examine the consequences of omitting transaction costs from welfare analysis in Section 5.4. In addition to their direct contribution to total costs, transaction costs mediate how prices respond to policy changes. Procurers may improve welfare by adjusting duration in response to a policy change. For a reduction in transaction costs of 20 percent, procurers issue 10 percent shorter contracts and realize 1.3 percent lower prices. Overall, total costs fall by 2.3 percent. An analysis of prices alone would capture only 56 percent of the full impact.

Related Literature To the best of the author's knowledge, this is the first paper to analyze the cost of inefficient supplier match over the duration of a contract. Previous literature on contract duration has focused on ex post coordination problems, primarily through costly renegotiation (Gray (1978); Masten and Crocker (1985)) and relationship-specific investments (Joskow (1987)). For many commodities, a first-order concern is not the proper alignment of buyer and seller incentives, but rather that buyers and sellers are efficiently matched.

For clarity, I abstract away from features of ex post incentive problems that have been studied previously, including risk-sharing, principal-agent relationships, the holdup problem, and incomplete contracting.[9] My work is complementary to these existing models. The tradeoff in this paper between transaction costs and price is closely related to the models of contract duration of Dye (1985) and Gray (1978), who take the stochastic price process as given. The innovation of this paper is to use tools of industrial organization to model primitives of the price process and explore its implications.

As a permanent contract is analogous to vertical integration, this paper is complementary to a large empirical literature that examines the determinants of the decision to vertically integrate. For a summary of this literature, see Lafontaine and Slade (2007). The cost factors that lead to longer contracts should also increase the propensity to vertically integrate, as noted by Coase (1960). Economists have recognized that a key factor in the decision to integrate is transaction costs, but most papers are concerned with testable implications of these costs, rather than their direct estimation.[10]

Previous papers that estimate transaction costs have focused on consumer markets, which is a different economic environment from the one analyzed here. A key feature of consumer markets is an inability to contract on future prices, leading to models that weigh an "investing" effect versus a "harvesting" effect (Farrell and Klemperer (2007)).[11] When buyers and sellers agree on future prices, as in this paper, these effects are competed away. Further, switching costs in previous empirical studies can be inferred from posted prices (e.g. Dubé et al. (2009) for orange juice and margarine, Elzinga and Mills (1998) for wholesale cigarettes[12]), whereas procurement prices are idiosyncratic to the buyer-seller match.

[9]For some examples, see Holmstrom (1983), Guriev and Kvasov (2005), and Rey and Salanie (1990).

[10]For examples of the testable implications approach, see Monteverde and Teece (1982), who consider proxies for asset specificity, and Walker and Weber (1984), who include proxies for demand uncertainty and technological uncertainty.

[11]For recent papers on this subject, see Cabral (2015) and Rhodes (2014).

[12]The wholesale market in the analysis of Elzinga and Mills (1998) mirrors a consumer market in that pricing, though nonlinear, is uniformly applied.

The effect of transaction costs on welfare has primarily been addressed in terms of their impact on equilibrium prices.[13] In procurement markets, transaction costs may be a sizable portion of total costs and should be accounted for in addition to any price effects.[14] Moreover, models in the consumer switching literature take supply costs as fixed (see, for example, Beggs and Klemperer (1992)), whereas variation in supply costs is a key factor in the decision to switch suppliers in procurement. In a recent paper, Carlton and Keating (2015) emphasize the role of transaction costs in welfare analysis when the affected variable is not simply the price level. In their setting, they examine a firm's ability to implement nonlinear pricing. Similarly, I present a case in which omitting the ability to adjust along one dimension of the contract (duration) would lead to a misleading analysis of welfare.

My contribution to the auction identification literature is most closely related to Krasnokutskaya (2011), who solves the problem of disentangling private costs from auction-specific heterogeneity by relying on two bids per auction. Since then, other authors have developed results for somewhat more general settings, by relying on three bids per auction (Hu et al. (2013)) or an observable reserve price in addition to the winning bid (Roberts (2013)). Aradillas-López et al. (2013) exploit variation in the number bids for second-price auctions, though the identification results of their paper are limited to constructing bounds on surplus. In this paper, I demonstrate point identification of surplus for both first-price and second-price auctions and partial identification of the full joint distribution of costs.

This model of this paper is equivalent to a simultaneous bundling problem, and is therefore complementary to the bundling literature. Most of the theoretical work deals thus far is concerned with a price-setting monopolist, rather than bundling items for auction.[15] The closest model to the one developed in this paper is that of Palfrey (1983).[16] Palfrey looks at a seller choosing between independent auctions and a single, bundled auction for all goods. I advance his analysis by allowing for intermediate degrees of bundling and by introducing transaction costs. Salinger (1995) and Bakos and Brynjolfsson (1999) note that bundling affects prices by reducing the variance of average valuations. In the setting of Bakos and Brynjolfsson (1999), this is an information advantage to the seller. However, as I demonstrate in this paper, the smaller variance induced by bundling reduces total surplus when there are no transaction costs. Cantillon and Pesendorfer (2006) share this insight in their analysis of combination bidding for multi-unit auctions.

---

[13]See Klemperer (1995).

[14]In consumer switching models, consumers so rarely switch that measuring prices alone suffices to quantify welfare effects.

[15]For examples, see Stigler (1963), Adams and Yellen (1976), Schmalensee (1982), and McAfee et al. (1989).

[16]Armstrong (2000) and Avery and Hendershott (2000) consider optimal auctions with multiple products when probabilistic bundling is allowed.

2 A Simple Model

In this section, I develop a simple model to illustrate how the structure of supply costs and transaction costs interact to determine duration.

Suppose a buyer is seeking the supply of a good for two periods. There are N suppliers in the market, and the set of suppliers stays the same across both periods. The buyer can issue either a single two-period contract, or sequential single-period contracts. The cost of each contract is δ, which captures search and information costs, as well as a complementarity of supplying for two periods. Assume that the buyer is risk-neutral.

The game proceeds in three steps. First, the buyer selects the contract duration: either two sequential contracts, or a bundled two-period contract. Second, suppliers realize cost draws for both periods. Third, suppliers participate in an auction for each contract, i.e. an efficient mechanism. [17]

When the buyer issues two single-period contracts, the cost for each contract is the random variable $c_{1:N}^S$, where the subscript indicates that it is the minimum from N draws. For a single two-period contract, the total cost is $2c_{1:N}^B$, where c is defined in terms of the per-period cost and the superscript B indicates that the distribution of per-period changes when the contracts are bundled.

Remark 1. As long as the separate costs c^S are not perfectly correlated across periods, $c^B \neq c^S$.

For one supplier, let c and c' denote the first and second period cost realizations, respectively. Assume they have the same marginal distributions. With no supply-side complementarity, $c^B = \frac{1}{2}(c + c')$. Then $Var(c^B) = \frac{1}{4}(Var(c) + 2 \cdot Cov(c, c') + Var(c'))$, and $Var(c^B) = Var(c) \iff Cov(c, c') = Var(c) \iff c' = c$.

In general, bundling will affect the per-period cost distribution.

2.1 The Efficient Case

At first, consider the case where the buyer's objective is to minimize total costs, i.e. the buyer internalizes any surplus the suppliers receive.[18] In the efficient case, the buyer's problem is

$$\min\{\underbrace{2E[c_{1:N}^S] + 2\delta}_{separate}, \underbrace{2E[c_{1:N}^B] + \delta}_{bundled}\}$$

It is optimal to bundle if $\frac{\delta}{2} > E[c_{1:N}^B] - E[c_{1:N}^S]$. Intuitively, if the amortized transaction cost of separate contracts is greater than the lost surplus from bundling, bundling is efficient.

From this simple model, we obtain a set of four predictions.

Prediction 1 Lower autocorrelation induces shorter contracts.

---

[17]In this simple model, suppliers have perfect foresight about future costs. A more general setup with imperfect information shares the same qualitative features of this model, though there is an additional ex post inefficiency arising from imperfect information.

[18]In Section 2.3, I consider the case where the buyer maximizes his own surplus.

Suppose that $\tilde{c}$ is a cost process with lower autocorrelation than c, but the same per-period marginal distribution, i.e. $E[\tilde{c}_{1:N}^S] = E[c_{1:N}^S]$. Then it follows that

$$
\begin{aligned}
E[\tilde{c}_{1:N}^B] &> E[c_{1:N}^B] \\
\implies E[\tilde{c}_{1:N}^B] - E[\tilde{c}_{1:N}^S] &> E[c_{1:N}^B] - E[c_{1:N}^S]
\end{aligned}
$$

The marginal cost of bundling is decreasing with the autocorrelation of the cost process. With less autocorrelation, bundling is more costly, and the rate of separate contracts increases.

Prediction 2 Greater variance in costs across suppliers induces shorter contracts.

For a simple case, consider location-scale transformations of $\tilde{c}$, such that $\tilde{c} = a + bc$. The optimality condition for bundling becomes

$$
\frac{\delta}{2} > b \left(E[c_{1:N}^B] - E[c_{1:N}^S] \right)
$$

As b increases, separate, shorter contracts become more desirable.

Prediction 3 When costs are bounded from below, greater variance may induce longer contracts.

If we impose the reasonable restriction that costs are bounded from below, greater variance may induce longer contracts by pushing the expected bundled price and expected separate price closer to the lower bound. Suppose c has a lower bound at 0, and let σ represent its standard deviation. Then,

$$
\lim_{\sigma \to \infty} E[c_{1:N}^B] = \lim_{\sigma \to \infty} E[c_{1:N}^S] = 0.
$$

As $E[c_{1:N}^B] - E[c_{1:N}^S] \to 0$ bundling becomes optimal in the limit. This effect tends to dominate as N gets large, as more draws brings the minimum closer to the lower bound.

Prediction 4 When costs are bounded from below, the number of suppliers has an inverse U-shape effect on the marginal cost of bundling. Therefore, duration may be decreasing, increasing, or U-shaped with N.

Intuitively, the cost of bundling is low when N is small, as averaging across periods has less effect on the first-order statistics when they are close to the mean. As N increases, the first-order statistics move away from the mean, increasing the cost of bundling. As N gets large, the first-order statistics approach the lower bound, which shrinks the cost of bundling.

2.2 A Numerical Example

To illustrate the above predictions, I present a simple case in which per-period costs are independent and uniform on $[0, 100]$. Figure 1a plots the expected costs for separate and bundled contracts, and Figure 1b displays the difference between the two lines in the first panel, i.e. the marginal cost of bundling. The dashed line in the second panel indicates a transaction cost of 18. When the blue line falls above this dashed line, the cost of bundling exceeds the transaction costs saved, and separate

contracts are optimal. Figure 1c plots optimal duration as a function of N. In this example, it is U shaped.

Key mechanics from the model can be illustrated with this simple graph. As transaction costs increase, the region of the blue line lying above the dash line shrinks, indicating a higher rate of bundled contracts. With high enough transaction costs, bundling is always optimal.

Figure 2 displays the change in marginal costs from increasing the variance of the cost distribution by 15 percent. For $N < 10$, greater variance increases the cost of bundling. When $N \geq 10$, the winning supplier's cost is close enough to the lower bound to reduce the cost of bundling.

2.3   The Buyer-Optimal Case

Earlier in this section, I considered the case where the buyer minimizes social costs. Typically, the buyer will instead choose to maximize his own surplus and will consider the transaction price instead of the supply cost. To analyze the buyer-optimal case, it is necessary to take a stance on the allocation mechanism. For exposition, assume the buyer uses an auction, and that costs are symmetric and independent across suppliers and periods. Under these conditions, the expected winning bid is equal to the expected second-order statistic of costs. The buyer will bundle if $\frac{\delta}{2} > E[c_{2:N}^B] - E[c_{2:N}^S]$.

Figure 3 captures the lock-in costs to the buyer, illustrated with the black line. This line is the sum of the total lock-in costs (blue line) and the seller surplus (red line). The buyer-optimal decision has the same qualitative features as the socially efficient contract, as it is an inverse U shape in the number of bidders.

The information rents from private costs drive a wedge between the buyer-optimal contract and the efficient contract. The buyer would prefer to issue separate contracts for $N \in \{8, ..., 55\}$, whereas it is efficient to issue separate contracts for $N \in \{5, ..., 19\}$.

The U shape arises in this model due to bundling's effect on the two channels of information rents and supply costs. When competition is low, bundling reduces information rents to the sellers by reducing the variance of cost draws. When competition is high, bundling increases information rents; when the second-order and first-order statistics are near the minimum, the second-order statistic (i.e., the expected price) responds more strongly to a reduction in variance.

2.4   Supply-Side Lock-In

In this section, I have explicitly modeled one mechanism for demand-side lock-in, as the focus in this paper is modeling buyer behavior. On the supply side, lock-in can arise from a seller's beliefs about future contract opportunities. When a seller expects better opportunities to arrive at some stochastic rate, then his opportunity cost will be increasing in the duration of the contract, and the seller will charge a premium for a longer contract.[19]

---

[19]If there is match-specific learning-by-doing, the seller's opportunity cost may be decreasing with the duration of the contract.

Figure 1: A Numerical Example

(a) Expected Costs of Separate and Bundled Contracts
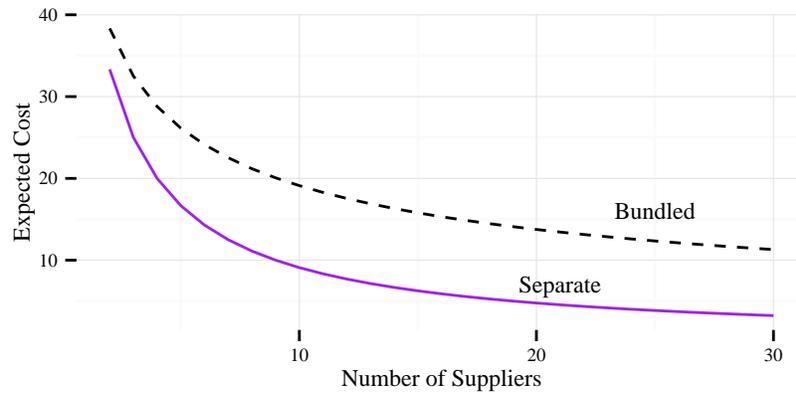
(b) Lock-In Costs of Bundling
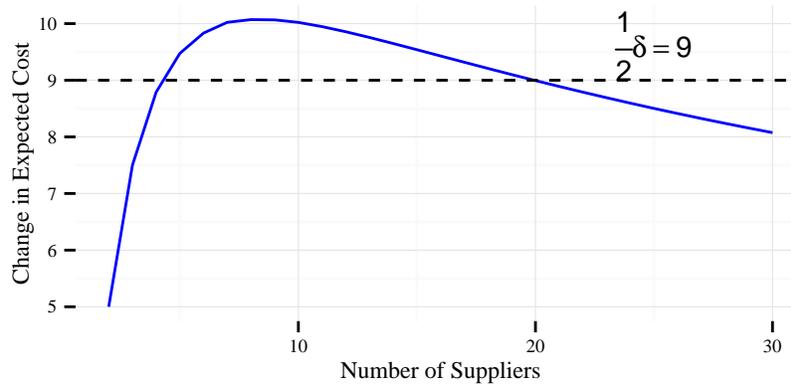
(c) Optimal Duration
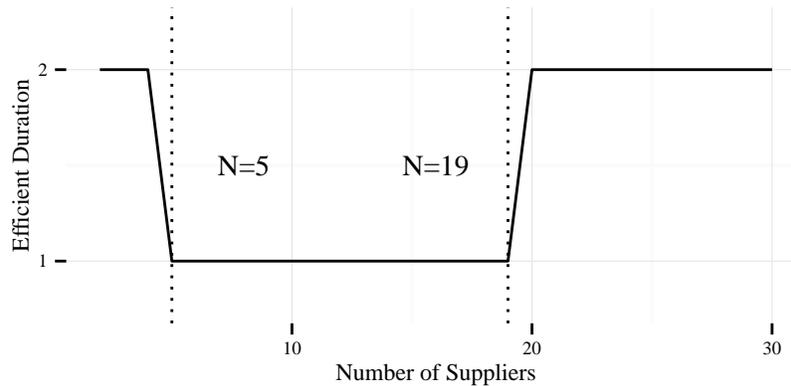
Notes: Panel (a) plots the expected per-period costs for bundled and separate contracts, as a function of the number of bids. The blue line in panel (b) is the difference between the two. The dashed line in panel (b) indicates a transaction cost of 18, which is the amount saved by bundling. For values of N where the blue line is above the dashed line ($N \in \{5, ..., 19\}$), separate contracts are efficient, as the increase in cost from bundling is greater than the savings in transaction costs.

Figure 2: Increased Variance in Cost
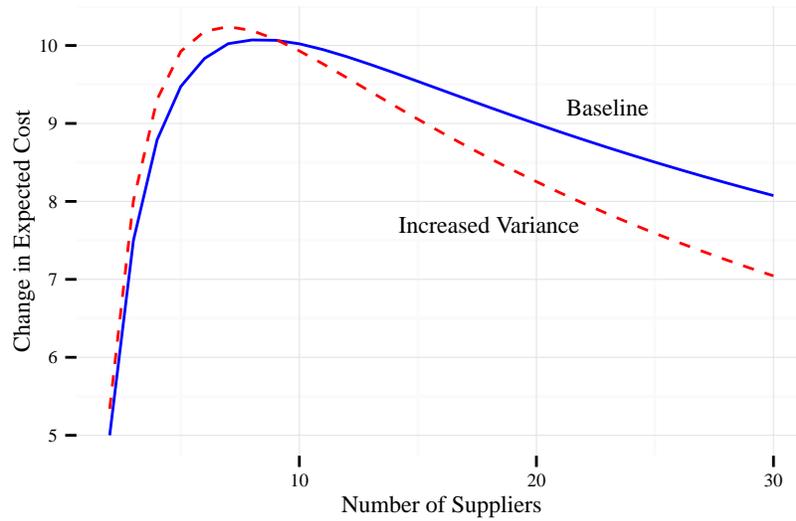
Notes: The blue line shows the increase in expected per-period costs from bundling, when per-period costs are independent and uniform on [0,100]. The red dashed line shows the cost of bundling when costs are drawn independently from symmetric beta distributions with the same mean and support, with a 15 percent greater variance.

Figure 3: Bundling's Effect on Price, Cost, and Surplus
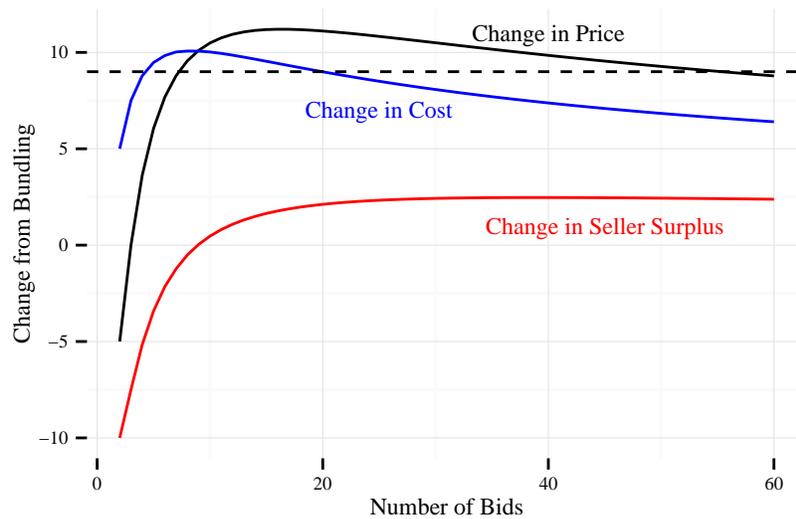
Notes: This figure shows the effect of bundling on different sides of the transaction. The blue line is the marginal social cost and is identical to the blue line in 1b. The black line is the change in price to the buyer, and the red line is the change in seller surplus. The dashed line indicates a transaction cost of 18, which is the amount saved by bundling. For values of N where the black line is above the dashed line, $N \in \{8, ..., 55\}$, the buyer would prefer to issue separate contracts, as the increase in price by bundling is greater than the savings in transaction costs. This range does not coincide with that of the efficient contract: $N \in \{5, .., 19\}$.

11

3 A Model of Supplier Selection

In this section, I first introduce a more general duration-setting procurement problem, and I show how the problem simplifies when the distribution of prices is stationary over time. I provide a set of conditions under which key components of the model, including transaction costs, are identified. I then specialize the model to an auction setting, which allows for identification of the joint distribution of costs when only the winning bid is observed. In the auction setting, as well as in the general model, I allow for unobserved heterogeneity and a form of selection on unobservables.

This theory is the basis for the empirical approach of Section 5.1.

3.1 The Duration-Setting Problem

Suppose that a buyer needs the supply of a good for a length of time S. The buyer decides on a policy T, which is the duration that the buyer will purchase from the seller. Each time the buyer decides to re-select a seller, he faces a transaction cost of δ. This transaction cost may represent the fixed costs of a relationship, search costs, or the cost of implementing a mechanism.

The game proceeds in three stages. First, the buyer determines duration T after observing contract characteristics X, entry cost shifters M, and the transaction cost δ. Second, N suppliers decide to participate in the supplier selection mechanism after observing (T, X, M). Third, a supplier is selected via a mechanism with a per-period stochastic price $P(N, T, X, M)$, where the price distribution may depend on the duration of the contract and the number of sellers.[20] One such mechanism is an auction, which I employ in the empirical analysis of section 5.

Let $\overline{P}$ denote the ex ante expected price conditional on (T, X, M), so that $\overline{P}(T, X, M) = \sum_{n=1}^{\overline{N}} \left(E[P(n, T, X, M)] \cdot \Pr(N = n | T, X, M) \right)$.

The buyer's problem is

$$\min_{J, \{T_j\}} \sum_{j=1}^{J} \left(T_j \cdot \overline{P}(T_j, x, m) + \delta \right) \; s.t. \sum_{j=1}^{J} T_j = S.$$

When $\overline{P}(T, X, M)$ is stationary, an optimal policy will have $T_j = T \, \forall j$. (Let S be sufficiently large to ignore the leftovers). Then the problem reduces to minimizing the average per-period price inclusive of transaction costs.

$$\min_{T} \overline{P}(T, x, m) + \frac{\delta}{T} \tag{1}$$

For a given mechanism, the procurer selects the optimal t satisfying the first order condition

$$\frac{d\overline{P}(T, x, m)}{dT} \Big|_{T=t} = \frac{\delta}{t^2} \tag{2}$$

For any interior solution t, $\frac{d\overline{P}(T, x, m)}{dT} \Big|_{T=t} > 0$. For the rest of this section, I assume that such interior solutions exist. In general, $\frac{d\overline{P}(T, x, m)}{dT}$ will be positive where lock-in costs exist. For example,

---

[20] The assumption that N is sufficient to describe P conditional on (T, X, M) rules out certain kinds of asymmetry.

if the per-period cost distribution is stable and the market is sufficiently competitive,[21] an increase in T causes suppliers to average cost draws across multiple periods. This shrinks the variance of the bundled cost distribution, which increases the expected minimum cost.[22] Thus, the buyer suffers from demand-side lock-in by failing to match with the lowest-cost supplier in each period.

Proposition 1. *When an interior solution exists, the optimal duration is increasing with transaction costs.*

Proof. See Appendix B. □

This model provides some predictions on the relationship between duration and observable characteristics X and M. For a particular application, it may be of interest to know if supply relationships will increase or decrease in response to lowered entry costs. This model relates the prediction to the second-partial of the pricing function, which can be estimated without modeling the buyer's decision or observing transaction costs.

Proposition 2. *Whether or not T is increasing with a variable in the pricing function is determined by the cross-partial of the expected price function.*

Proof. See Appendix B. □

3.2 A Procurement Model

In this section, I develop a three-stage procurement model, where the first stage is the duration-setting problem. I place restrictions on the general model that allow for nonparametric identification when only the transaction price, the number of participants, and cost shifters X and M are observed. I allow for an unobservable cost shifter, U, that may affect the participation decision, and I show that, even in the presence of selection on unobservables, the model is identified. Independence and multiplicative separability will be important restrictions that allow for identification.

1st Stage: Duration Setting The procurer observes (X, M, δ) and sets T to minimize the expected per-period price plus the amortized transaction cost. The price consists of a proportional offer B and common multiplicative cost shifters $h(X)$ and U, where U is unobserved by the procurer. The proportional offer is an equilibrium strategy when suppliers are risk neutral and private costs and common costs are multiplicatively separable.[23] The procurer's objective function is:

---

[21]If $N = 2$, this will generally not be the case. It is ambiguous for $N = 3$.

[22]In the limit, all suppliers' costs are equal to the long-run average.

[23]One example is the auction framework I discuss later.

$$\min_{T} \overline{P}(T, x, m) + \frac{\delta}{T}$$

$$= \min_{T} E[B \cdot U \cdot h(X) | T, x, m] + \frac{\delta}{T}$$

$$= \min_{T} \left(\sum_{n=1}^{\mathbb{N}} E[B \cdot U | n, T, x, m] \cdot \Pr(N = n | T, x, m) \right) h(X) + \frac{\delta}{T}$$

2nd Stage: Participation Potential entrants observe (U, T, X, M) and an common entry cost shock ε. Bidders enter if expected profits exceed entry costs. Let π_n denote the proportional expected profits for the n^{th} marginal entrant. The entry condition is given by

$$E[\pi_n | n, t] \cdot h(x) \cdot U - k(m) \cdot \varepsilon > 0 \iff N \geq n$$

3rd Stage: Supplier Selection After the participation decision, a mechanism is used to select a single supplier from the N participants. The mechanism has an stochastic price $B \cdot U \cdot h(x) | (N, T, X, M)$.[24] One example mechanism is a first-price auction, where B would be the lowest submitted bid. Another example is a challenger-incumbent game, in which suppliers submit take-it-or-leave it offers to the buyer that the incumbent can decide to match.

3.3 Identification

Identification in this model proceeds in two parts. In the first part, the participation and supplier selection components of the model are separated from the duration decision and nonparametrically identified. Thus, identification participation and price components holds even if T is not set optimally, and the results generalize to cases of supplier selection with no duration decision.

In the second part, I use the duration decision and previously identified components of the model to identify contract-specific transaction costs.

3.3.1 Identification of Entry and Offers

The econometrician observes the transaction price $P = B \cdot U \cdot h(X)$ as well as (N, T, M, X). The cost shocks U, ε, and C are unobserved by the procurer and the econometrician, but their distributions are common knowledge. Assume

1. *Conditional Independence:* $B \perp\!\!\!\perp U | (N, T, X, M)$ and $B \perp\!\!\!\perp \varepsilon$.

2. *Independence of Unobservables:* $(\varepsilon, U) \perp\!\!\!\perp (T, X, M)$.

3. $h(\cdot)$ and $k(\cdot)$ are continuous, and the range of $h(\cdot)$ or $k(\cdot)$ has broad support.

---

[24]As I mention in the description of the first stage, separability in B and $U \cdot h(X)$ arises from risk neutrality and separability in private costs and common costs.

Proposition 3. *When (P, N, T, X, M) is observed, the following components of the model are identified:*

1. *$E[B|N, T, X, M]$*
2. *$E[U|N, T, M, X]$.*
3. *$h(X)$ and $k(M)$, up to a normalization.*
4. *The distribution of $\frac{\varepsilon}{U}$.*
5. *Relative profits for n and n' participants: $\frac{E[\pi_n|n,T]}{E[\pi_{n'}|n',T]}$.*
6. *Relative profits for t and t' with n participants: $\frac{E[\pi_n|n,t]}{E[\pi_n|n,t']}$.*

Proof. See Appendix C. $\qquad\square$

Identification of these components of the model allow for the identification of contract-specific transaction costs, as I demonstrate below. Further, these components are useful for estimating the impact of counterfactuals, such as a reduction in participation costs. Importantly, identification is obtained even when the underlying selection mechanism is obscure. Thus, the model can be used for policy analysis while maintaining an agnostic approach to the supplier selection mechanism.

To conduct an efficiency analysis, we need to supplement with additional data on expected profits or put additional structure on the model. With the above assumptions, only relative profits are obtained. Data on profits for one (n, t) pair identifies the expected profit function and, therefore, the expected supply cost $E[C|N, T]$. When no data is present, specifying the selection mechanism can pin down seller surplus. For example, when a supplier is selected with an auction among symmetric bidders, surplus is identified. I explore this case in Section 3.3.3. Now, I turn to the identification of transaction costs.

3.3.2 Identification of Transaction Costs

Recall the buyer's objective function:

$$
\min_T \left(\sum_{n=1}^{\mathbb{N}} E[B \cdot U|N = n, T, x, m] \cdot \Pr(N = n|T, x, m) \right) h(X) + \frac{\delta}{T}
$$

$$
= \min_T \left(\sum_{n=1}^{\mathbb{N}} E[B|N = n, T] \cdot E[U|N = n, T, x, m] \cdot \Pr(N = n|T, x, m) \right) h(X) + \frac{\delta}{T} \quad (3)
$$

Where the second line is obtained under conditional independence. When T is continuous, point identification of δ is obtained directly from the first order condition. In many applications, such as the empirical one in this paper, duration is discrete, issued in monthly or yearly increments. In these cases, bounds for transaction costs can be obtained.

Proposition 4. *When T is continuous, δ is identified for each contract. When T is discrete, bounds for realizations of δ are identified.*

Proof. In the continuous case, δ is identified from the first-order condition of equation (3). In the discrete case, denote the duration choice set $\mathbb{T}$. Revealed preference for the chosen duration t provides a set of inequalities on transaction costs of the form:

$$
\begin{aligned}
(t'-t)\delta \quad < \quad & t \cdot t' \bigg(\sum_{n=1}^{N} E[B|n,t'] \cdot E[U|n,t',x,m] \cdot \Pr(N=n|t',x,m) - \\
& \sum_{n=1}^{N} E[B|n,t] \cdot E[U|n,t,x,m] \cdot \Pr(N=n|t,x,m) \bigg) h(x)
\end{aligned}
\tag{4}
$$

for all $t' \in \mathbb{T}\backslash t$. These inequalities provide upper bounds on δ when $t' > t$ and lower bounds when $t' < t$. The minimum upper bound and the maximum lower bound provide bounds on δ. $\square$

Even in the discrete case, the distribution of δ can be identified from additional assumptions on the relationship between δ and X or M. This distribution can be used as a prior over the bounds.

Proposition 5. *Assume δ is independent of X. When (i) $h(X)$ varies continuously with X, (ii) the range of $h(X)$ is $(0,\infty)$, and (iii) X has full support on the domain of $h(\cdot)$, then the distribution of δ is identified.*

Proof. As the bounds in equation (4) vary continuously with X, the cumulative distribution function of δ is identified. $\square$

3.3.3 Identification of the Auction Model

Placing additional restrictions on the structure of the supplier selection mechanism allows for the identification of seller surplus, and, in the case of auctions, partial identification of the joint distribution of outcomes. The auction model is the basis for the empirical analysis in Section 5.1.

In addition to the previous assumptions, further assume:

1. The selection mechanism is an auction (first-price or second-price).

2. *Conditional Independence:* The winning proportional bid B is determined by private costs $C_i|T \sim F_{i,T}$, where $C_i \perp\!\!\!\perp U|(N,T,X,M)$.

3. *Symmetry*: $F_i = F$ for all i.

4. F is continuous with positive support. $U \sim G$, where G has positive support.

5. Auctions with sequential values of $N \in \{\underline{N}, ..., \overline{N}\}$ are observed, with $\underline{N} < \overline{N}$.

Proposition 6. *When the supplier selection mechanism is an auction with symmetric bidders, seller surplus is identified.*

Proof. See Appendix C. $\square$

Symmetry and independence are standard assumption in auction models of unobserved heterogeneity.[25] We can obtain further identification of the distribution of private costs in the auction framework using variation in N.

Proposition 7. *The distribution of private costs is identified up to the first $(\overline{N} - \underline{N} + 2)$ expected order statistics of $\overline{N}$ draws from F.*

Proof. See Appendix C. □

Observe that if $\underline{N} = 2$ and $\overline{N} \to \infty$, the restrictions on expected order statistics approximate the quantile function, and F is exactly identified. The restrictions are powerful in that they may reject many classes of flexible distributions with $(\overline{N} - \underline{N} + 2)$ parameters.

Corollary 1. *The distribution of unobserved heterogeneity is determined by identification of F.*

Proof. By independence, we can use the characteristic function transform to write $\varphi_{\ln W_n}(z) = \varphi_{\ln B_n}(z) \cdot \varphi_{\ln U}$, where $W_n = Y_n / h(X)$ is the observed winning bid scaled by the observables. We can perform this exercise conditional on every realization of (N, T, X, M). Once the characteristic function of F is obtained, either by exact identification ($\overline{N} \to \infty$) or by flexible estimation methods, G is pinned down. □

Finally, the auction model provides an alternative identification result if no instrument M is available. With no selection on unobservables, the model is still identified.

Proposition 8. *First-price, symmetric auctions with independent unobserved heterogeneity and conditionally independent private values are identified with only the winning bid. In particular, seller surplus and the first $(\overline{N} - \underline{N} + 2)$ expected order statistics of $\overline{N}$ draws from F are identified. Identification is obtained without modeling entry as long as there is no selection on unobservables.*

Proof. See Appendix C. □

---

[25]See, for example, Aradillas-López et al. (2013) and Krasnokutskaya (2011). Work to relax these assumptions is in progress.

4 Empirical Application: Janitor Contracts

4.1 Data

The contracts analyzed in this paper are competitive contracts for janitorial services for the United States federal government. Janitorial services were chosen from all federal contracting goods and services because they are numerous, the product is homogenous, and there is a lot of variation in contract duration. Finally, demand is inelastic. That is, this market is a relatively clean setting to analyze the duration-setting decision of the procurer. Further, the market is sizable, totaling $5.1 billion from 2004 to 2013.

For each contract, I observe key variables including start and end dates, the total value of the project, the square footage of the building to be cleaned, the 9-digit ZIP code where the project is performed, the number of offers received for the contract, and variables indicating whether or not the contract is competitive. In addition, there are detailed variables on the winning bidder, including the DUNS number, which allows me to track whether an incumbent wins a follow-on contract.

I collected contract documents directly from an online portal used by the federal government to solicit bids.[26] For a clean comparison, I selected competitive, small-business contracts, which are the most numerous. I obtained the relevant contract documents (request for proposal, cleaning frequency charts, maps, etc.), and constructed detailed contract information directly from the documents. Variables include the square footage of the site to be cleaned, the number of buildings at the site, and the frequency of service.

I merged the contract data with detailed location, price, and vendor information maintained in the Federal Procurement Data System (FPDS).[27] By law, the FPDS keeps public records of all contracts for the U.S. federal government. This dataset contains of approximately 30,000 janitor contracts from 2000 to 2014.[28] Of the larger dataset, 7,558 observations were for competitive, small-business contracts, and 1,994 of these had solicitation identifiers that could be matched to the online portal with contract documents. The resulting sample contains observations that have data on square footage and cleaning frequency, as these variables are important for explaining variation in the price of the contracts.[29] After removing outlier contracts,[30] 420 observations remain.

I matched the contract-specific dataset with auxiliary datasets of 1) government contracting expenditures at the same location in related products and 2) local labor market conditions. The local labor-market conditions include county-level unemployment.

---

[26]The website, FedBizOpps.gov, is the most common posting location for competitive contracts, which, by law, must be posted publicly.

[27]These data were obtained from USASpending.gov.

[28]Data on other products in the FPDS were used to construct a measure of transaction cost and control variables for the size of the office.

[29]The majority of the 1,994 contracts either 1) did not list project details or 2) linked from the portal to a secure military server, so I was unable to obtain details for that subset.

[30]Outlier contracts were those with 1) a monthly price greater than $50,000, 2) duration shorter than 3 months, 3) frequency less than once per week, or 4) cleaning tasks outside of general janitorial services.

4.1.1 Institutional Details

Competitive contracts are contracts that are posted publicly and allow open competition from registered vendors.[31] Many of these contracts are posted on the centralized web portal FedBizOpps.gov, from which I collected the data in this analysis. On the website, a prospective supplier can view the contract details, including contract duration and the square footage of the building, requirements for the job, and a list of interested suppliers. From the portal, a supplier submits a bid to the contracting office that includes the total price over the duration of the contract. The contracting office determines the winning supplier primarily based on the lowest price. By law, the contracting office must justify selecting other than the lowest-price offer.[32]

Importantly, contract duration is determined locally by the local contracting officer. As several industry personnel described to the author, contract duration is a balance between minimizing the administrative costs of re-contracting and realizing the benefits from re-competing more frequently. Costs may be increasing with duration because suppliers charge a premium or because the procurer ends up locked in to a high-cost supplier. This motivates using this market as a case study for the model developed in this paper. Transaction costs and competition are key motivating factors for the procuring agencies.

4.1.2 Description of Services

Contracts include specifications for the tasks to be done and their frequencies. Tasks for janitors include cleaning and sanitizing restrooms, vacuuming carpets, dusting, and emptying trash cans. For an example list of specifications, see Appendix H.

The majority of the contracts are for services at a single site, with one or two buildings on the site. The majority of the contracts (267) are for office cleaning, though frequently an office includes an auxiliary building, such as an exercise room, a bunkhouse, or a small warehouse. For the empirical analysis of this paper, offices with auxiliary buildings were classified as Field Offices. Table 1 lists the frequency of each type of site, which are grouped into eight major categories.[33]

4.1.3 Summary Statistics

Summary statistics for the contracts are displayed in Table 2. Contracts vary in price, duration, and competition. As shown later, much of the variation in price can be captured by the square footage of the building and the cleaning frequency.[34] The maximum weekly frequency can exceed 7 as

---

[31]These contracts fall under three categories: Full and Open Competition, Full and Open Competition after the Exclusion of Sources, and Competed Under Simplified Acquisition. 86 percent of the contracts deemed Full and Open Competition after the Exclusion of Sources are listed as a small business set-aside. As 96 percent of the contracts are won by small businesses (as determined by the contracting officer), I ignore this distinction for the purposes of analysis. See Federal Acquisition Regulation (FAR) Part 5.

[32]Based on the guidelines established by FAR and conversations with local contracting offices, the contracting office will prefer suppliers that have an established history.

[33]For a breakdown of contracts by the issuing department or agency, see Appendix G.

[34]Cleaning frequency is encoded as the maximum required weekly frequency in the contract specifications. In cases where the maximum frequency changed between season, the approximate mean was used

Table 1: Count of Sites by Category

Category	Sub-Category	Count
Office (146)		
	Office	61
	Recruiting Office	85
Field Office (121)		
	Field Office	13
	Ranger District Office	82
	Ranger Station	11
	Reserve Fleet	2
	Work Center	13
Research (79)		
	Laboratory	19
	Plant Materials Center	5
	Research Center	11
	Weather Station	44
Visitors (23)		
	Cemetery	3
	Museum	2
	Recreation Area	11
	Restroom	2
	Visitor Center	5
Airport (14)		
	Airport	14
Services (13)		
	Service Center	7
	Vet Center	6
Technical (13)		
	Converter Station	1
	Data Center	2
	Equipment Concentration Shop	1
	Missile Tracking Annex	1
	Power Plant	8
Medical (11)		
	Clinic	8
	Medical Center	3
Total		420

Table 2: Summary Statistics: Janitor Contracts

	mean	min	p25	p50	p75	max
Annual Price (1000s)	23	1.1	7.1	12	20	543
Duration (Years)	3.9	.25	3	4.8	5	5
Number of Bids	7.2	2	4	6	9	31
Square Feet (1000s)	13	.3	4	6.8	13	267
Num. Buildings	2.3	1	1	1	2	32
Num. Sites	1.3	1	1	1	1	6
Weekly Frequency	3.5	.5	2.5	3	5	14
Firm Size	45	1	2	13	56	380
Observations	420					

multiple shifts are occasionally explicitly required.

One important source of variation in the analysis is in the number of bids received. The median contract receives 6 bids, though some contracts receive as many as 31 bids. This variation helps disentangle the effect of private costs from unobserved heterogeneity in the structural analysis.

In the last row, the table provides the number of employees for the winning firms. The winning firms in this dataset are quite small, with a median of 13 employees. Over 25 percent of the winning suppliers have 2 or fewer employees. In 1.4 percent of the contracts, firms have 250 or more employees, yet all winning firms are classified as small businesses by the contracting office.

Figure 4a displays a scatterplot of the logged values of the winning bids on the *y*-axis against the number of bidders on the *x*-axis. The pattern observed in the scatterplot - large variation in prices with clustering at the median price, rather than the minimum - motivates the assumption of unobserved auction-specific costs used in this paper.

In the second panel, Figure 4b, I display residualized values for the (log) winning bid. The residuals were constructed from a regression of price on duration, square footage, cleaning frequency, and agency and site type fixed effects. Even after controlling for observable characteristics, there is large variation in prices for auctions with many bidders. This variation supports the assumption of unobserved heterogeneity. Though much of the variation in prices can be explained by observables, there is still residual variation that is inconsistent with an independent private values model; the model with multiplicative common shocks fits far better.

These contracts have a great deal of variation in duration. Figure 5 plots the contract duration against the number of bids. The number of bids and contract duration are positively correlated. There is a good deal of variation in duration, ranging from 3 months to 5 years, though contracts tend to cluster at yearly increments. Additionally, 43 percent of contracts are for 5 years, which is the maximum contract duration imposed by federal budgeting regulations. These features of the data motivate the counterfactual analysis I perform in Section 5.4, where I consider relaxing the cap and the yearly increment constraint.

Figure 4: Price versus Number of Bids

(a) Annual Price
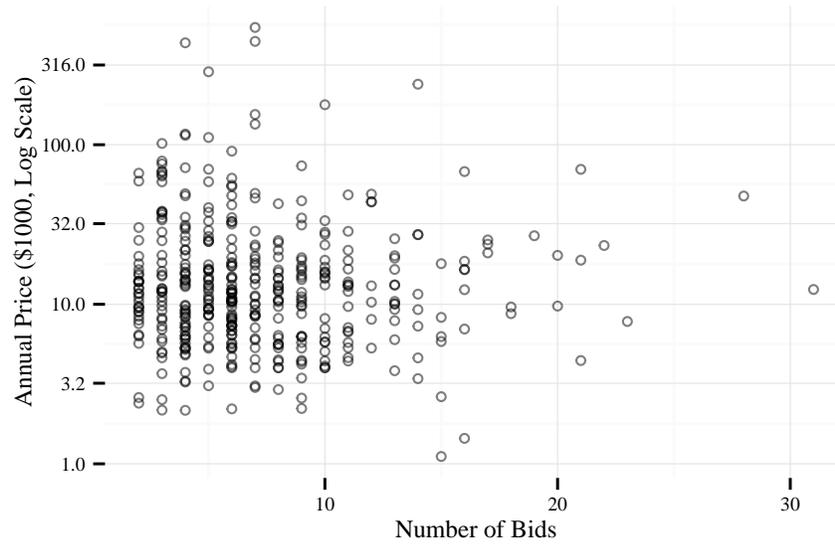
(b) Residualized Annual Price
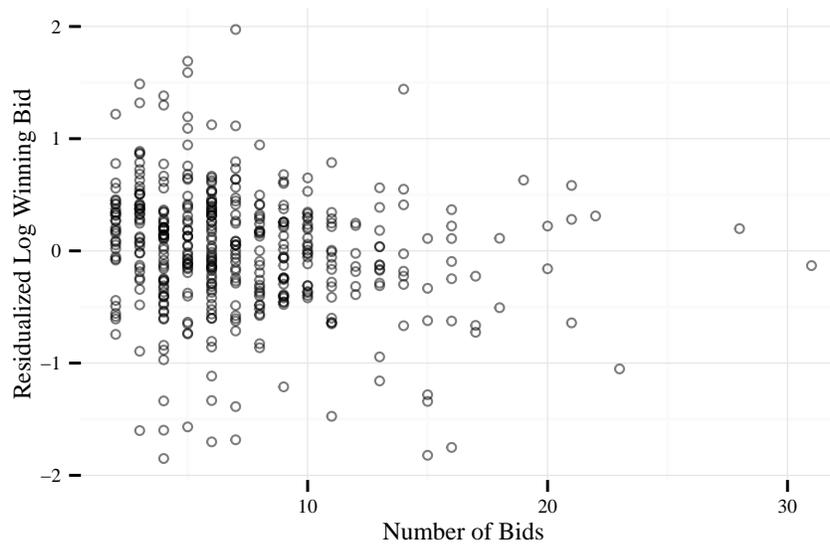
Notes: This figure plots the log annual price against the number of bids received for each contract. There is a great deal of variation in the annual price, much of which cannot be explained by observable variables. It is notable that some of the highest and lowest prices are realized with few bidders.

Figure 5: Contract Duration versus Number of Bids

Notes: This figure plots the contract duration against the number of bids received for each contract. 43 percent of the contracts are at the maximum duration of 5 years. Contracts are clustered in yearly intervals, though the support in between full years is relatively well-covered. A linear fit shows an increasing relationship between duration and the number of bids. Additionally, shorter contracts occur less frequently as the number of bids increases.
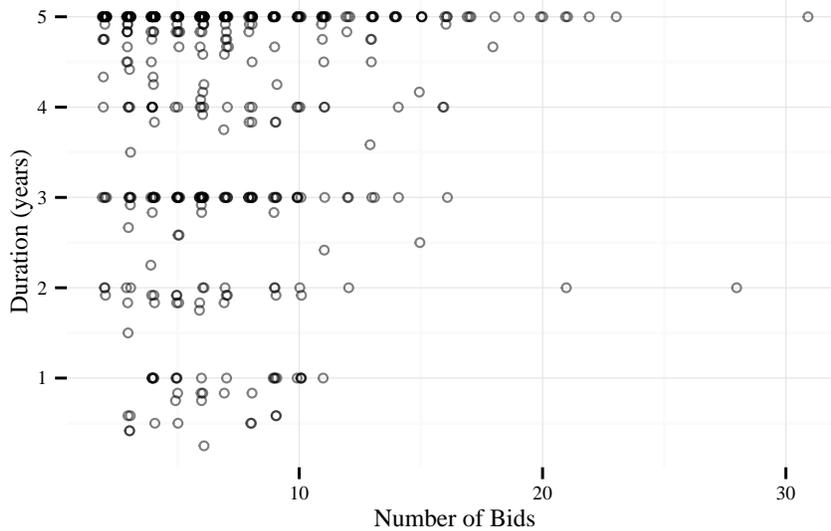
4.2 Descriptive Regressions

In this section, I present descriptive regressions to motivate the choice of variables and assumptions made in the structural estimation. Table 3 provides regressions of the log monthly price on the number of bids, duration, and controls. The main specification in this table is OLS-2. Square footage and cleaning frequency together capture 56 percent of the variation in price. Once these cost variables are controlled for, the price declines by 3 percent with each additional bid. This decline is consistent with a conditionally independent private values framework, and is intuitive: competition should lower the price.

To account for possible selection on unobservables (endogenous entry), I employ an instrumental variables approach using local labor market conditions. Higher unemployment, relative to baseline levels, increases entry into janitorial services. Using the exogenous variation in unemployment relative to 2004 levels[35] as an instrument for the number of bids, I estimate a more negative point estimate of -5 percent per bid, though it is not significantly different from the -3 percent coefficient of OLS-2. In Section 3.3.3, I developed a model that can correct for the presence of selection on unobservables. In Section 5, I estimate the structural model with endogenous entry, though results incorporating selection on unobservables are an extension in progress.[36]

The inclusion of fixed effects for fiscal year, agency, and site type in specifications OLS-3 and

---

[35]The first contract in my dataset was issued in 2005.

[36]Accounting for selection on unobservables adds computational difficulty to the already challenging problem of unobserved heterogeneity.

Table 3: Descriptive Regressions: ln(Price per Month)

	OLS-1	OLS-2	OLS-3	IV-1	IV-2	IV-3
Number of Bids	−0.008	−0.030***	−0.025***	−0.050**	−0.055	−0.041**
	(0.010)	(0.007)	(0.007)	(0.023)	(0.036)	(0.019)
Duration (Years)	0.059*	0.025	−0.011	0.031	−0.007	0.028
	(0.032)	(0.021)	(0.021)	(0.023)	(0.022)	(0.022)
ln(Square Footage)		0.530***	0.590***	0.545***	0.619***	0.539***
		(0.031)	(0.033)	(0.035)	(0.047)	(0.034)
Weekly Frequency		0.179***	0.115***	0.178***	0.117***	0.178***
		(0.018)	(0.022)	(0.018)	(0.022)	(0.018)
Site Type FEs			X		X	
Agency FEs			X		X	
Fiscal Year FEs			X		X	
Observations	420	420	420	420	420	420
R^2	0.01	0.57	0.66	0.57	0.64	0.57

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays regression results for regressions of log monthly price on auction characteristics. The variables from specification OLS-2 are included in the structural model. These regressions show that square footage and cleaning frequency are important in explaining variation in prices. Once square footage and cleaning frequency are accounted for, fixed effects for fiscal year, agency, and site type add little explanatory power. Specifications IV-1, IV-2, and IV-3 are two-stage least squares regressions, where the instruments for the number of bids are monthly county-level unemployment (IV-1 and IV-2) and monthly county-level unemployment, one-year moving average county-level unemployment, and the number of janitor establishments in the same 3-digit zip code (IV-3). Unemployment measures are the logged difference between current and 2004 values.

IV-2 have a low per-variable impact on R^2 and do not have a substantial effect on the estimated coefficients. I omit them from structural estimation. In IV-3, I consider an alternative in which I instrument with current unemployment, one-year average unemployment, and the number of janitorial establishments in the same 3-digit ZIP code.

The regressions capture a near-zero relationship between price and duration. In practice, the linear model does not properly account for the dual effects of duration on price and (via profits) on entry. In the structural model, I find a positive and significant relationship.

In Table 4, I display regressions of the number of bids on auction characteristics and local measures of unemployment. Specification (2) is equivalent to the first-stage regression in IV-1. In (3), I split the instrument into current unemployment and the baseline level. Though current unemployment is associated with more bids, higher baseline levels are associated with fewer bids.

Table 4: Descriptive Regressions: Number of Bids

	(1)	(2)	(3)	(4)
Duration (Years)	0.313**	0.347**	0.234*	0.117
	(0.154)	(0.147)	(0.139)	(0.144)
ln(Square Footage)	0.727***	0.695***	0.794***	0.846***
	(0.221)	(0.211)	(0.199)	(0.218)
Weekly Frequency	−0.064	−0.119	−0.291**	−0.079
	(0.130)	(0.124)	(0.119)	(0.147)
ln(Unemp./(2004 Unemp.))		4.198***		
		(0.657)		
ln(Unemployment)			4.019***	2.822***
			(0.617)	(0.761)
ln(2004 Unemp.)			−3.119***	−2.070***
			(0.633)	(0.781)
Site Type FEs				X
Agency FEs				X
Fiscal Year FEs				X
Observations	420	420	420	420
R^2	0.04	0.12	0.23	0.34
F-statistic	5.1	14.4	24.4	7.4

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays regression results for regressions of the number of bids on auction characteristics and local labor market variables. Specification (2) is equivalent to the first-stage regression of IV-1 in Table 3. In (3) and (4), I split the instrument into the current component and the baseline (2004) component.

5 Results from the Structural Model

Estimation of the structural model proceeds in three steps. First, I use a parametric maximum likelihood to perform joint estimation of entry and bidding. Second, using the duration decision of the procurer and estimated parameters from the first step, I construct distribution-free bounds for transaction costs. Third, I construct estimates of transaction costs by applying a prior over the bounds.

Using the estimated transaction costs and parameters of the model, I conduct policy counterfactuals in Section 5.4.

5.1 Estimation of Entry and Bidding

For the janitor contracts, I estimate the auction model of Section 3.3.3. Though I obtain nonparametric identification results, I employ a parametric approach for first-stage estimation in this paper. In this application, I have the added complication of estimating a duration-dependent distribution of private costs. This increases the number of parameters needed for any nonparametric approach. Parametric assumptions make estimation computationally tractable while capturing the first-order features of the model.

I employ parameterizations of the objects of interest, where

- Private costs: $c \sim$ Weibull, with mean $\mu(T) = \mu_0 + \mu_1 T$ and shape $\alpha(T) = \alpha_0 + \alpha_1 T$

- Unobserved heterogeneity: $U \sim \ln \mathcal{N}(-\frac{\sigma_U^2}{2}, \sigma_U^2)$. (Mean = 1)

- Observed heterogeneity: $h(X) = square\_footage^{\gamma_1} \cdot weekly\_frequency^{\gamma_2}$

- Entry costs:
  $k(M) = square\_footage^{\kappa_1} \cdot weekly\_frequency^{\kappa_2} \cdot unemployment^{\kappa_3} \cdot 2004\_unemployment^{\kappa_4}$

- Entry shock: $\varepsilon \sim \ln \mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2)$

The Weibull distribution is chosen for tractability and flexibility, as it allows the estimated probability density functions to be either convex or concave. For the distribution of unobserved heterogeneity, the log-normal distribution was chosen because it best fit the model out of several choices.

I allow the parameters of the private cost distribution to vary linearly with duration, which captures the first-order effects of interest in this model. As I am not taking a stand on the underlying cost process, I am in this sense estimating a "reduced-form" primitive for the cost distribution.[37]

Using the entry and bidding problems below, I estimate these parameters using maximum likelihood. For details of the likelihood, see Appendix E. Note that square footage and weekly frequency affect the entry decision by both increasing supply costs and affecting the fixed cost of entry.

---

[37]For a microfounded model, see Appendix D.

Table 5: Parameter Estimates

Group	Parameter	Variable	Estimate	95 Percent C.I.
Private Costs	μ_0	-	23.844	[20.767, 27.287]
	μ_1	Duration	0.708	[0.127, 1.339]
	α_0	-	3.659	[2.559, 5.716]
	α_1	Duration	0.161	[0.000, 0.389]
Heterogeneity	σ_U	-	0.597	[0.542, 0.642]
	γ_1	Square Footage	0.538	[0.484, 0.605]
	γ_2	Weekly Frequency	0.530	[0.405, 0.650]
Entry	μ_ε	-	-0.360	[-0.719, -0.131]
	σ_ε	-	0.641	[0.580, 0.693]
	κ_1	Square Footage	0.395	[0.299, 0.499]
	κ_2	Weekly Frequency	0.666	[0.469, 0.865]
	κ_3	Unemployment	-0.763	[-0.945, -0.591]
	κ_4	2004 Unemployment	0.608	[0.430, 0.788]

Notes: Parameter estimates are from joint maximum likelihood estimation. Confidence intervals are constructed via bootstrap.

$$Entry \quad N > n \iff E[\pi_n|t] \cdot h(x) \cdot E[U] - k(m) \cdot \varepsilon > 0$$
$$Bidder \quad \max_b (b - c) \, (h(x) \cdot u \cdot t) \Pr(b \, wins|n)$$

[Empirical results that allow for selection on unobservables, in which bidders observe U before entering, are in an extension in progress].

5.2 First-Stage Results

Table 5 displays the parameter estimates from the first-stage estimation. Prices increase with duration through an increase in the mean private cost and a reduction in the variance of cost draws. Square footage and weekly frequency are scaled by the mean, so that the estimate of μ_0 is interpreted as the mean annual private cost draw for a zero-duration contract at a typical location. The mean annual private cost is $23,800 and increases by 3 percent per contract year. For a visual representation of how costs depend on duration, I plot the density of private cost draws for a one-year and a five-year contract in Figure 6.

As expected, higher values for square footage and weekly frequency increase costs, and higher current unemployment lowers entry costs. Consistent with the findings in the descriptive regressions, higher baseline levels of unemployment correspond to higher entry costs. Square footage has a net positive effect on entry, as $\gamma_1 > \kappa_1$. Supply costs, which are positively correlated with profits, increase by more than entry costs for square footage. Weekly frequency, on the other hand, has a net negative effect on entry, as $\gamma_2 < \kappa_2$.

Figure 6: The Distribution of Bidder Costs

(a) Duration-Dependent Private Costs
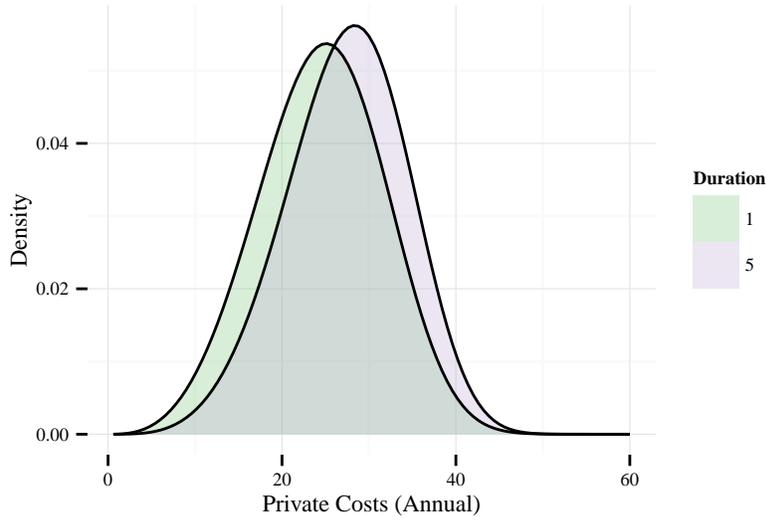
(b) Unobservable Auction-Specific Heterogeneity
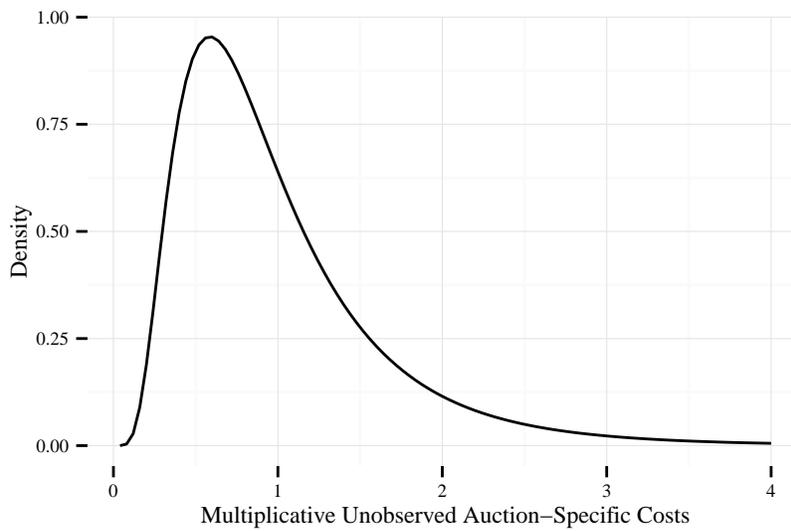
Notes: This figure plots the distributions of the unobservable components of bidder costs. Private costs are displayed in panel (a), and the density of unobserved auction-specific heterogeneity is displayed in panel (b). In panel (a), the density is plotted for a one-year contract and a five-year contract. The density shifts smoothly between these functions for intermediate values of duration.

Table 6: Bounds on Transaction Costs

T	$E[\underline{\delta}]$	$E[\overline{\delta}]$	$E\left[\overline{\delta}/\underline{\delta}\right]$	$\frac{T+1}{T-1}$
1	0.0	0.9	∞	∞
2	1.3	3.9	3.00	3.00
3	3.8	7.6	2.00	2.00
4	6.4	10.6	1.66	1.67
5	14.2	-	-	1.50

Notes: Transaction costs are bounded by the revealed preference of the procurer. For $T = 5$, only the lower bound is identified, as there is a cap at 5 years.

The model fits the data well. In Figure 7, I display actual values for annual prices compared to the predicted values. Though the correlation is very strong, unobserved heterogeneity is important to match the distribution of prices. Figure 6, Panel (b) plots the estimated distribution of unobserved heterogeneity. Unobservable costs are meaningful, in that they capture over 40 percent of the variance of log prices.

5.3 Estimating Transaction Costs

In this section, I develop distribution-free bounds for transaction costs based on the estimated parameters for entry and bidding. The procurer's decision problem is

$$\min_T \sum_{n=1}^{\overline{N}} E[B_n|n, T, x, m] \cdot E[U|n, T, x, m] \cdot h(x) + \frac{\delta}{T}$$

As demonstrated in Section 3.3.2, the optimality condition can be used to construct bounds for δ in the case where T is discrete and point estimates when T is continuous. In my data, contracts are either set to the nearest monthly or nearest yearly increment, providing a set of tight and loose bounds, respectively. For the results of this section, I assume that U is not observed prior to entry, so that $E[U|n, t, x, m] = E[U] = 1$. Table 6 provides the mean bounds for the yearly-increment contracts.[38]

As contracts are capped at 5 years, only a lower bound for δ for contracts at the cap can be obtained without additional assumptions. To proceed with counterfactual analysis, I construct conservative upper bounds on transaction costs for contracts at the cap. They are conservative in that they treat all contracts at the cap as optimal. As 43 percent of contracts are at the cap, many of the optimal contracts under a higher cap would be likely be longer, implying higher upper bounds on transaction costs than constructed here.

To construct the upper bounds, I use a simple heuristic based on the observed ratio of upper to

---

[38]Some of these yearly-increment contracts may arise from a procurer who can issue a monthly-increment contract. A kernel smoother estimates that these would be less than 7 percent of the contracts at yearly intervals. By taking a stance on which procurers are constrained, the bounds could be tightened for those contracts.

Figure 7: Model Fit: Annual Price

(a) Actual versus Predicted
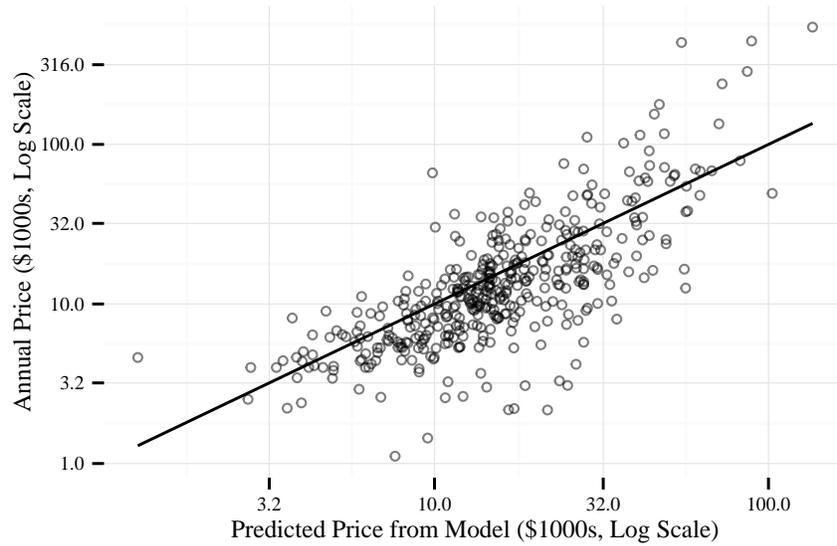
(b) Density with Unobserved Heterogeneity
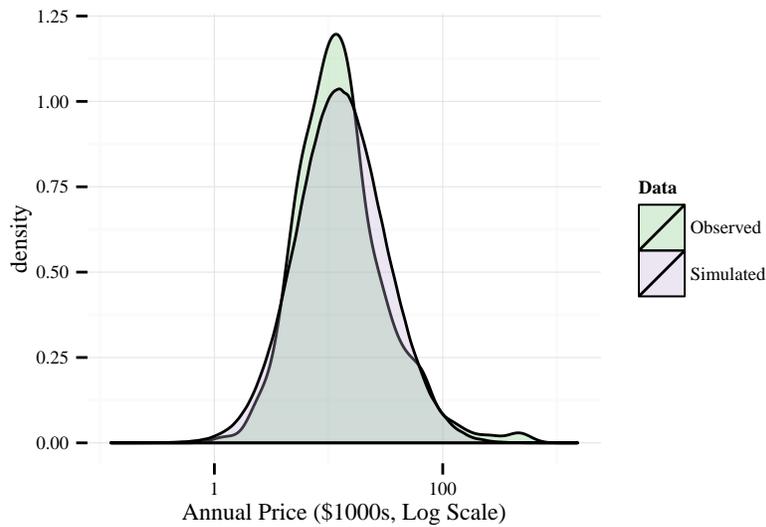
Notes: This plots display the model fit. In panel (a), prices are plotted against predicted prices from the model. In panel (b), prices are simulated using one million draws from the model to incorporate estimated unobserved heterogeneity. The green area represents the density for the observed data, and the purple area is the density from the simulation. The densities are plotted using a Gaussian kernel with Silverman's rule-of-thumb bandwidth.

Table 7: Estimated Transaction Costs

(a) In $1000s

	Mean	95 Percent C.I.	Min	25th Pct.	Median	75th Pct.	Max
$T < 5$	6.51	[1.64, 11.83]	0.03	1.56	4.67	8.26	63.76
$T = 5$	17.75	[4.76, 31.38]	0.86	9.39	12.65	23.47	109.40
All	11.27	[2.92, 20.38]	0.03	3.69	8.13	13.81	109.40

(b) As a Fraction of Total Cost (Percent)

	Mean	95 Percent C.I.	Min	25th Pct.	Median	75th Pct.	Max
$T < 5$	11.8	[3.7, 19.0]	0.5	6.5	10.5	15.0	54.4
$T = 5$	18.3	[6.3, 26.7]	1.7	11.4	17.1	22.0	57.6
All	14.5	[4.8, 22.2]	0.5	7.9	13.1	18.6	57.6

Notes: Estimated transaction costs are the expectation taken with a uniform prior over the distribution-free bounds identified from the duration decision of the procurer. For $T = 5$, conservative upper bounds are projected using the ratio of upper bounds to lower bounds for $T \in \{1, 2, 3, 4\}$.

lower bounds for $T \in \{1, 2, 3, 4\}$. These ratios move in parallel to the function $\frac{T+1}{T-1}$, whose values are included in Table 6. For $T = 5$, I estimate $\bar{\delta} = 1.5 \times \underline{\delta}$, following the pattern of the table.

Finally, to construct expected transaction costs and conduct counterfactuals, I apply a uniform prior for the density between the distribution-free bounds. Using the prior, I construct point estimates by taking the expectation. The uniform prior is appealing for its transparency.[39]

Transaction costs are significant in this setting, comprising 10.1 percent of annual costs.[40] Table 7 contains summary statistics for the estimated transaction costs. The first panel displays the values, and the second displays transaction costs as a percentage of total costs. Those at the maximum duration ($T = 5$) have much larger estimates of mean transaction costs: $17,750 versus $6,500 for those below.

For context, these estimates are not unreasonable given cost estimates provided to the author by a senior contracting officer. The officer estimated that a simple janitorial contract would take about three weeks of full-time work for an employee whose salary would be approximately $75,000 to $90,000. Based on 50 full-time workweeks, this gives a cost range of $4,500 to $5,400. Larger projects may take months of work and multiple officers.

One of the tradeoffs to taking the sequential, distribution-free approach to bounding transaction costs is that information on the duration decision is not used in the estimation. More precise estimates could be obtained at the expense of additional, restrictive assumptions. Thus, the magnitudes of the point estimates are sensitive to the estimation of the parameters μ_1 and α_1 in the

---

[39]The uniform prior is also appealing for the reason that the observed duration is optimal at the mean transaction cost when procurers can issue contracts in monthly increments. If a left triangular prior were used instead, the optimal monthly-increment contract would be shorter than the observed value for $T \in \{1, 2, 3, 4, 5\}$.

[40]The bootstrapped 95 percent confidence interval is [3.9, 16.7]. This differs from the estimate of 14.5 percent in Table 7 as the 14.5 percent is calculated as the mean fraction across contracts. The mean transaction cost is about half of the mean annual price.

Figure 8: Estimated Transaction Costs versus Percent of Total Cost
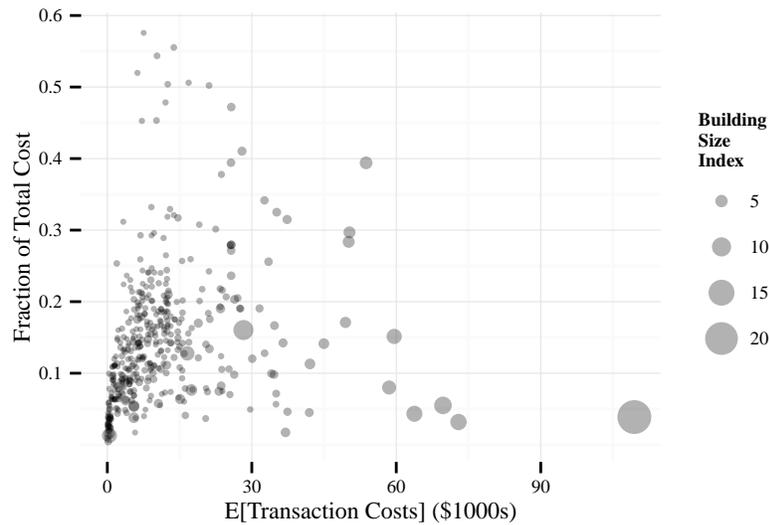
Notes: This figure plots the portion of total costs arising from transaction costs against the estimated transaction costs. The largest relative transaction costs occur at moderate levels of transaction costs. The largest transaction costs comprise a small portion of total costs.

first stage. The uncertainty in these parameters is captured by bootstrapping, which results in the confidence intervals in the table.

In some cases, the estimated transaction costs are quite large as a percent of total costs. As shown in Figure 8, this is not due to very high absolute costs, but rather due to moderate transaction costs realized for low-price projects. The highest transaction costs are estimated for very large projects, such as the Railroad Retirement Board's 267,000 square-foot headquarters in Chicago (the largest point on the graph). On the other hand, the highest fraction of total costs is obtained for small projects. All of the contracts where transaction costs are greater than 45 percent of the total costs have an annual price less than $6,000; these include the cleaning of ranger offices in national parks.

5.3.1 Verifying Estimated Transaction Costs

As an exercise to verify the estimated transaction costs, I project the estimates on other variables not used in the structural estimation. First, I categorize the estimates by building type and by department, in Table 8. As one would expect, the highest transaction costs are obtained for airports, technical facilities (e.g. power plants), and medical centers. Standard office cleaning has the lowest mean transaction costs. In the second panel, I group by government department. Homeland Security has the highest median transaction cost, at $23,600 per contract. Conversely, Agriculture and the Interior have low median transaction costs, at $6,300 and $4,000 respectively.

In Table 9, I regress the estimated transaction costs on variables excluded from the structural

Table 8: Estimated Transaction Cost ($1000s) by Category

(a) Building Type

Type	Median	CV	Mean	Std. Dev.	Count
Airport	22.1	0.5	22.2	11.9	14
Technical	20.8	0.7	25.2	18.2	13
Medical	17.3	0.4	16.6	7.3	11
Research	11.7	0.8	12.7	10.7	79
Visitors	11.0	1.0	15.9	15.9	23
Field Office	6.9	1.1	10.1	11.3	121
Office	6.2	1.4	8.3	11.4	146
Services	6.0	1.2	8.5	9.9	13

(b) Department

Type	Median	CV	Mean	Std. Dev.	Count
Homeland	23.6	0.6	24.7	14.5	19
Other	12.6	1.3	21.6	29.0	16
Commerce	11.8	0.7	12.7	8.6	56
VA	11.3	0.6	12.4	7.1	19
Defense	6.7	1.1	9.3	10.2	151
Agriculture	6.3	1.0	9.5	9.8	148
Interior	4.0	1.5	15.4	22.7	11

Notes: This table displays summary statistics for expected transaction costs estimated by the model. In panel (a), transaction costs are grouped by building type. In panel (b), they are grouped by contracting department.

Table 9: Projecting Transaction Costs on Variables Outside of the Model

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Num. Pages)	0.153***				0.100**	0.067
	(0.052)				(0.046)	(0.054)
ln(Related Expenditures)		0.073***			0.032**	0.028*
		(0.018)			(0.016)	(0.016)
Simplified Acq. Ind.			−0.363***		−0.179	−0.136
			(0.137)		(0.118)	(0.151)
ln(Square Footage)				0.518***	0.452***	0.486***
				(0.057)	(0.061)	(0.066)
Weekly Frequency				0.228***	0.232***	0.186***
				(0.033)	(0.033)	(0.042)
Site Type FEs						X
Agency FEs						X
Fiscal Year FEs						X
Observations	420	420	420	420	420	420
R^2	0.02	0.04	0.02	0.31	0.33	0.41

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays estimated coefficients from regressing estimated log transaction costs on variables outside of the model. These variables are (i) the (log) number of pages in the contract, (ii) log government procurement expenditures at the same 9-digit ZIP for maintenance, office furniture, and other housekeeping services, and (iii) an indicator for whether the contract falls under the federal government's simplified acquisition protocol. In all specifications, I include site type and agency fixed effects. In specifications (1)-(5), I include fiscal year fixed effects.

model. Included variables are the number of pages in the contract, related expenditures[41] in the same 9-digit ZIP, and an indicator for whether the contract falls under the simplified acquisition protocol. These three variables enter with the expected sign, though simplified acquisition seems to be explained by other observable characteristics.

High-expenditure locations are associated with higher transaction costs. Economic theory could rationalize a sign in either direction, as economies of scale lead to a positive association and capacity constraints produce a negative one.

5.4 Counterfactual Analysis

In this section, I consider counterfactual analysis in the presence of transaction costs for the estimated model. First, I examine an existing policy that caps the maximum contract duration at five years. The cap is exogenous to janitor services and applies broadly to federal government procurement. I provide suggestive evidence that the five-year cap is close to the optimal (capless) policy.

Second, I examine counterfactual analysis more broadly when transaction costs are omitted

---

[41]Spending on other housekeeping, maintenance, and office furniture.

Table 10: Policy Counterfactual: Changing the Cap

| New Cap | Percent Change in Costs | | Affected Contracts | | |
	Aggregate Effect	Marginal Effect	Count	Percent of Costs	Aggregate Effect
$\overline{T} = 4$	0.4	0.4	238	69.8	0.6
$\overline{T} = 3$	2.2	1.8	271	74.6	3.3
$\overline{T} = 2$	8.1	5.7	353	88.9	9.2
$\overline{T} = 1$	30.9	21.1	383	93.8	33.1

Notes: This table displays the results of a policy counterfactual of lowering the cap. The first column reports the aggregate percent change in costs across all contracts, and the second column reports the marginal effect lowering the cap by one year. In the final column, the aggregate effect is reported for the affected contracts only.

from the model. I consider an example policy change that reduces transaction costs by 20 percent. This example highlights the importance of quantifying transaction costs. Total costs fall by 2.3 percent, but prices fall by only 1.3 percent. An event study estimate based on prices alone would capture only 56 percent of the true impact.

5.4.1 Policy Counterfactual: Changing the Cap

Government-wide budgeting regulations lead to a five-year cap on contracts. This cap applies to all agencies, though an exception can be requested and obtained for large projects (e.g. long-term weapons contracts). In my sample, 43 percent of contracts are for five years. A natural policy counterfactual is to evaluate how duration and prices respond to an increase in the cap.

A challenge in performing this counterfactual is that estimates on the upper tail of transaction costs cannot be obtained without placing further assumptions on the distribution of transaction costs. By constructing conservative, distribution-free bounds earlier, the result is that the observed five-year contracts are optimal at the expected transaction cost, so an increase in the cap would have little impact. As stated earlier, these estimates likely understate the latent value of transaction costs.[42]

One approach is to estimate the effect of reducing the cap. This will still understate the impact for five-year contracts, but will accurately incorporate the effect on the contracts under the cap. Table 10 displays the result of the counterfactual. The aggregate effect in the second column is the percent change in total costs form moving from a five-year cap to the new lower cap. The marginal effect in the third column is the percent change in costs from moving the cap down by one year. That is, the compounded marginal effects are equal to the aggregate effect.

The results from this analysis suggest that the five-year cap generates little excess cost. The marginal effect of lowering the cap from five years to four years increases total cost by only 0.4 percent, which is small and much less than the 1.8 percent marginal effect of moving from $\overline{T} = 4$ to $\overline{T} = 3$. The intuition behind this result is that the marginal benefit of amortizing the transaction

---

[42]This seems especially likely given the high proportion of contracts at the cap. Any estimate that accounts for potential censoring requires a stance on the tail behavior of the distribution of transaction costs. For this reason, I take the conservative approach described.

Table 11: Effects of a 20 Percent Reduction in Transaction Costs

Variable	Symbol	Percent Change	95 Percent C.I.
Total Costs	$P + \frac{\delta}{T}$	-2.3	[-3.8, -0.9]
Prices	P	-1.3	[-2.3, -0.5]
Duration	T	-10.4	[-10.6, -10.2]
Percent Captured by Prices		55.6	[51.1, 60.6]

Notes: This table reports the equilibrium changes to prices, duration, and total costs when transaction costs fall by 20 percent. The last row reports the percent of the welfare impact that is captured by prices, which represents the treatment effect estimated when price changes are observed but transaction costs are not quantified. Transaction costs are evaluated at the mean transaction cost using a uniform prior over previously estimated distribution-free bounds.

costs over a longer period is declining rapidly, at rate $\frac{1}{T}$. For other markets, it is worth examining policies on maximum duration as well as rule-of-thumb practices that have a fixed duration for a particular product,[43] as caps can be costly. As stated earlier, these estimates could greatly understate the impact of the cap if the true transaction costs at the cap are much higher.

5.4.2 Welfare Analysis with Transaction Costs

Transaction costs are important to welfare analysis as they can constitute a substantial portion of total costs and affect how equilibrium prices respond to a change in the economic environment. When transaction costs are unaffected by a policy change, a welfare analysis that omits transaction costs will misstate the impact for two reasons. First, the measured impact on prices should be weighted by the share of total costs attributable to prices. That is, the impact should be discounted toward zero by the share attributable to (unaffected) transaction costs. Second, market participants adjust equilibrium behavior in response to the change. The choice of duration provides an additional margin of adjustment, improving welfare compared to an analysis that takes duration as fixed.

When transaction costs are affected by a policy change, the above two forces also affect welfare estimates. Changes to transaction costs should be directly accounted for in the welfare calculation, and any such changes allow for new duration and price choices that may improve welfare. Consider the two leading forms of counterfactual analysis: a structural model and an event study. For both approaches, assume that transaction costs are ignored. Any counterfactual evaluation in the structural model will find no impact due to transaction costs, as they are outside of the model. The event study, which makes use of observable responses to a policy change, will capture the impact of transaction costs on prices and duration. However, the direct impact on transaction costs will be left out.

To illustrate the impact of transaction costs on welfare analysis, I consider a hypothetical policy that reduces transaction costs by 20 percent. As Table 11 indicates, procurers respond to lower

---

[43]For example, a three-year contract is the industry standard for office supplies.

transaction costs by issuing shorter contracts, reducing prices by 1.3 percent. Total costs, accounting for the direct effect on transaction costs, fall by 2.3 percent. The change in total costs is nearly 80 percent greater than the estimate obtained by an analysis of prices that omits transaction costs.

5.5 Examining Model Assumptions

For the structural model, I have made two strong assumptions: 1) I have ruled out auction-to-auction effects on procurer and supplier behavior, and 2) I have assumed that entry is exogenous conditional on observables. In Section 4.2, I presented evidence using instrumental variables for the validity of the exogenous entry assumption. [Structural results that relax this assumption are in progress].

One common source of asymmetries in practice is the presence of an incumbent bidder, who may have advantages over the competitors through familiarity with the buyer. Here, I present some evidence to suggest that the assumption of symmetry (i.e. no incumbent advantage) is reasonable in this setting.

5.5.1 Dynamic Considerations

In a dynamic setting, the procurement process might result in asymmetries between bidders that would invalidate the assumptions of symmetry that are maintained in the paper. Winning bidders may be at an advantage in subsequent auctions due to learning by doing or lowered transaction costs of retaining the same supplier. Additionally, competing bidders may retain some information about competitors if costs are correlated over time.[44]

In the janitorial data, dynamic considerations do not appear to be meaningful. I implement two tests for asymmetries. In the first test, I identify 147 follow-on contracts in my sample. Follow-on contracts are contracts that begin within a year of another contract in the same 9-digit zip, where the prior contract may be any of the approximately 30,000 janitor contracts in the procurement data. I also identify 49 contracts as those won by an incumbent supplier, where the incumbent was a supplier of janitorial activities at the same 9-digit zip in the previous year. Under symmetry, an incumbent is expected to win 24 percent of the 147 contracts.[45] The incumbents win $\frac{49}{147} = 33$ percent of the contracts.[46] One reason why the incumbents may win a larger proportion than the naive symmetry estimator is that the government disqualifies bids from vendors with no prior experience. If the government disqualifies one bidder per auction, then the incumbents' expected win percentage rises to 32 percent under symmetry. I do not have measurements of disqualification in my data. Incumbents may have an advantage, but it does not appear to be overwhelming.

---

[44]Saini (2012) discusses the literature on endogenous asymmetries and evaluates a model in which capacity constraints hurt the winning bidder.

[45]As I only observe winning bidders, I am unable to adjust for when a supplier does not bid on a follow-on to the supplier's current contract. I construct the expected winning bid rate as $\min(\frac{I}{n}, 1)$ where I is the number of incumbent bidders, as for two observations the number of incumbents exceeds the number of bids n.

[46]When looking at contracts where there is only one qualified incumbent, the difference increases to 34 percent compared to 20 percent under symmetry, though the sample is smaller.

The second test is to include a dummy for follow-on contracts in the descriptive regressions to determine if variation in prices and entry are explained by the presence of an incumbent bidder. None of the coefficients on the dummy are significant, and its inclusion does not meaningfully change any of the coefficients of interest. For these regressions, see Appendix F. Based on these tests, I perform my analysis maintaining the assumptions of symmetry.

6 Conclusion

In this paper, I develop a model of optimal contract duration arising from underlying supply costs and transaction costs. I show how latent transaction costs may be recovered from the duration decision of the procurer. Transaction costs may be a significant portion of total costs, and the method that I provide in this paper should prove useful for welfare analysis, especially in markets where long-term relationships are prevalent.

For much of the paper, I omit a discussion of efficiency. In the data, I find that deviations from the efficient contracts are small. For a treatment of efficiency, including the optimal tax to incentive the socially efficient contract, see Appendix A.

In many settings, the tradeoff of this paper here may complement other first-order concerns, such as ex post incentive problems (e.g. arising from asset specificity) and risk considerations. An appropriate model should be tailored to the industry in question. Future work of this author includes demonstrating how observed structural cost features predict variation in relationship duration across industries.

The identification results presented here are practical, as they allow for flexible models to be estimated with readily available data. In the auction setting, I allow for unobserved heterogeneity and selection on unobservables, while only observing the winning bid. As many auction datasets only contain the winning bid, the results can be used for a more flexible and accurate economic analysis. Further, empirical progress can be made for competitive supplier selection mechanisms that do not map to the auction model.

The analysis presented here offers, albeit indirectly, one novel prediction regarding the theory of the firm. Supply relationships lie in between independent firms and full vertical integration. We should expect that conditions favorable for long-term contracts also give rise to integration, as the end is similar and integration may result in additional benefits. Long-term relationships arise when competition is sufficiently low, and also when it very intense. Likewise, integration with an upstream firm may be least likely when the upstream industry is moderately competitive, as downstream firm realizes a large benefit by switching among suppliers.

References

ADAMS, W. J. AND J. L. YELLEN (1976): "Commodity Bundling and the Burden of Monopoly," *The Quarterly Journal of Economics*, 90, 475.

ARADILLAS-LÓPEZ, A., A. GANDHI, AND D. QUINT (2013): "Identification and Inference in Ascending Auctions With Correlated Private Values," *Econometrica*, 81, 489–534.

ARMSTRONG, M. (2000): "Optimal multi-object auctions," *The Review of Economic Studies*, 67, 455–481.

AVERY, C. AND T. HENDERSHOTT (2000): "Bundling and optimal auctions of multiple products," *The Review of Economic Studies*, 67, 483–497.

BAKOS, Y. AND E. BRYNJOLFSSON (1999): "Bundling information goods: Pricing, profits, and efficiency," *Management science*, 45, 1613–1630.

BEGGS, A. AND P. KLEMPERER (1992): "Multi-Period Competition with Switching Costs," *Econometrica*, 60, 651.

CABRAL, L. (2015): "Dynamic pricing in customer markets with switching costs," Working paper.

CANTILLON, E. AND M. PESENDORFER (2006): "Combination bidding in multi-unit auctions," Unpublished manuscript.

CARLTON, D. W. AND B. KEATING (2015): "Antitrust, Transaction Costs, and Merger Simulation with Nonlinear Pricing," *Journal of Law and Economics*, 58, 269–289.

COASE, R. (1960): "The Problem of Social Cost," .

DUBÉ, J.-P., G. J. HITSCH, AND P. E. ROSSI (2009): "Do switching costs make markets less competitive?" *Journal of Marketing research*, 46, 435–445.

DYE, R. A. (1985): "Optimal Length of Labor Contracts," *International Economic Review*, 26, 251.

ELZINGA, K. G. AND D. E. MILLS (1998): "Switching Costs in the Wholesale Distribution of Cigarettes," *Southern Economic Journal*, 65, 282.

FARRELL, J. AND P. KLEMPERER (2007): "Chapter 31 Coordination and Lock-In: Competition with Switching Costs and Network Effects," in *Handbook of Industrial Organization*, Elsevier, vol. 3, 1967–2072.

GRAY, J. A. (1978): "On indexation and contract length," *The Journal of Political Economy*, 1–18.

GURIEV, S. AND D. KVASOV (2005): "Contracting on time," *American Economic Review*, 1369–1385.

HOLMSTROM, B. (1983): "Equilibrium Long-Term Labor Contracts," *The Quarterly Journal of Economics*, 98, 23.

HU, Y., D. MCADAMS, AND M. SHUM (2013): "Identification of first-price auctions with non-separable unobserved heterogeneity," *Journal of Econometrics*, 174, 186–193.

JOSKOW, P. L. (1987): "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets," *The American Economic Review*.

KLEMPERER, P. (1995): "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *The Review of Economic Studies*, 62, 515–539.

KRASNOKUTSKAYA, E. (2011): "Identification and Estimation of Auction Models with Unobserved Heterogeneity," *The Review of Economic Studies*, 78, 293–327.

LAFONTAINE, F. AND M. SLADE (2007): "Vertical integration and firm boundaries: the evidence," *Journal of Economic Literature*, 629–685.

MASTEN, S. E. AND K. J. CROCKER (1985): "Efficient adaptation in long-term contracts: take-or-pay provisions for natural gas," *The American Economic Review*, 1083–1093.

MCAFEE, R. P., J. MCMILLAN, AND M. D. WHINSTON (1989): "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *The Quarterly Journal of Economics*, 104, 371.

MONTEVERDE, K. AND D. J. TEECE (1982): "Supplier Switching Costs and Vertical Integration in the Automobile Industry," *The Bell Journal of Economics*, 13, 206.

PALFREY, T. R. (1983): "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information," *Econometrica*, 51, 463.

REY, P. AND B. SALANIE (1990): "Long-term, Short-term and Renegotiation: On the Value of Commitment in Contracting," *Econometrica*, 58, 597.

RHODES, A. (2014): "Re-examining the effects of switching costs," *Economic Theory*, 57, 161–194.

ROBERTS, J. W. (2013): "Unobserved heterogeneity and reserve prices in auctions," *The RAND Journal of Economics*, 44, 712–732.

SAINI, V. (2012): "Endogenous asymmetry in a dynamic procurement auction," *The RAND Journal of Economics*, 43, 726–760.

SALINGER, M. A. (1995): "A graphical analysis of bundling," *Journal of Business*, 85–98.

SCHMALENSEE, R. (1982): "Commodity bundling by single-product monopolies," *Journal of Law and Economics*, 67–71.

STIGLER, G. J. (1963): "United States v. Loew's Inc.: A note on block-booking," *Sup. Ct. Rev.*, 152.

STIGLER, G. J. AND J. K. KINDAHL (1970): *The behavior of industrial prices*, vol. 90, National Bureau of Economic Research New York.

WALKER, G. AND D. WEBER (1984): "A transaction cost approach to make-or-buy decisions," *Administrative science quarterly*, 373–391.

A Efficiency

This paper provides a theory of buyer-seller relationships governed by fixed-term contracts. In this section, I show how market-determined contracts may differ from efficient contracts. Contracts that are determined by market participants (buyers and sellers) may be too long or too short, resulting in wasteful social costs. Counterintuitively, these extra costs may increase as a market becomes more competitive. Therefore, from a policy standpoint, highly competitive markets may be of more concern for regulators than those that are more concentrated. This result occurs because market participants care about price rather than cost, and the price responds more quickly to a change in contract length than the cost when the number of bidders is large.[47]

Imperfect competition drives a wedge between the revenue-optimal contracts and the contract size that maximizes social surplus. I demonstrate that the direction of the wedge is tied to whether the buyer surplus is increasing or decreasing with the length of the contract.

A.1 A Framework Relating Optimal and Efficient Bundle Size

In Section 3, I presented the buyer-optimal solution to the duration-setting problem. What about efficiency? The social planner's concern is expected costs, rather than expected price.[48] Thus, the efficient contract solves

$$\min_{T} E[C(N,T)] + \frac{\delta}{T}$$

with the first-order condition

$$\frac{dE[C(N,T)]}{dT}\big|_{T=\ddot{T}} = \frac{\delta}{\ddot{T}^2} \tag{5}$$

In general, $E[P(N,T)] \neq E[C(N,T)]$, which will result in an inefficiency when the contract is determined by the buyer. As long as interior solutions exist (see Proposition (10)), this result gives the simple relationship that the efficient bundle size $\ddot{T}$ will be larger than the equilibrium bundle size when $\frac{dE[P(N,T)]}{dT} > \frac{dE[C(N,T)]}{dT}$. Defining the expected seller surplus as $E[\pi(N,T)] = E[P(N,T)] - E[C(N,T)]$, we have the following result:

Proposition 9. *When interior solutions exists, the efficient contract will be longer than the equilibrium*

---

[47]If we think of expected price as the expected second-order statistic, and the cost as the first-order statistic, then we have some intuition for why this could be true. The second-order statistic responds more strongly to a change in variance (or mean) than the first-order statistic when the number of draws is large and the cost distribution is bounded from below. The buyer (or seller) internalizes the contract length's effect on the second-order statistic rather than its effect on the first-order statistic.

[48]In this setting, I assume the social planner is limited by information constraints; in this setting the social planner cannot observe the private information about sellers' costs. This reflects the idea that the mechanism (and the associated transaction costs) are important to the truthful revelation of information. A third party with full information would solve a different problem, awarding the contract to the lowest-cost seller at every instant and switching when the net savings outweigh the transaction cost.

contract if and only if the expected seller surplus is increasing at $\dot{T}$:

$$\ddot{T} > \dot{T} \iff \left(\frac{dE[P(N,T)]}{dT} - \frac{dE[C(N,T)]}{dT} \right) |_{T=\dot{T}} > 0$$

$$\iff \frac{dE[\pi(N,T)]}{dT} |_{T=\dot{T}} > 0$$

The existence of interior solutions depends on the concavity of the expected price function.

Proposition 10. *Interior solutions to the buyer's problem and social planner's problem exist as long as the first-order conditions (2) and (5) can be satisfied and $E[P(N,T)]$ and $E[C(N,T)]$ are not too concave. In particular, $\frac{d^2 E[P(N,T)]}{dT^2}|_{T=\dot{T}} > -\frac{2}{\dot{T}} \frac{dE[P(N,T)]}{dT}|_{T=\dot{T}}$ and $\frac{d^2 E[C(N,T)]}{dT^2}|_{T=\ddot{T}} > -\frac{2}{\ddot{T}} \frac{dE[C(N,T)]}{dT}|_{T=\ddot{T}}$. These are the second-order conditions to ensure that first-order conditions above achieve a minimum.*

A.2 Allocation of Term-Setting Rights

Given the general model, we can identify settings in which inefficiency arising from market power over contract length may be of first-order importance. In this section, I provide some intuition and a heuristic guide to the assignment of term-setting rights to limit such inefficiencies.

The buyer's problem can be written in the following form:

$$\min_{T} E[P(N,T) - C(N,T)] + E[C(N,T)] + \frac{\delta}{T}$$

$$= \min_{T} E[\pi(N,T)] + E[C(N,T)] + \frac{\delta}{T}$$

Notice that when $\frac{dE[\pi(N,T)]}{dT} = 0$, this problem is equivalent to the social planner problem. Therefore, when the buyer sets the duration of the contract, these contracts will be efficient when the seller surplus does not change with the length of the contract. The more sensitive buyer surplus is to the duration of the contract, the greater the potential for inefficiency.

What about assigning contract term-setting power to the sellers? Sellers solve the problem:

$$\max_{T} E[P(N,T) - C(N,T)] - \frac{\delta}{T}$$

$$= \min_{T} -E[P(N,T)] + E[C(N,T)] + \frac{\delta}{T}$$

Sellers solve the social planner problem when $\frac{dE[P(N,T)]}{dT} = 0$. Therefore, if price is not sensitive to contract duration, it is efficient to let the sellers determine the length of the contract.[49]

If either price or buyer surplus changes with the duration of the contract, there is potential for inefficiency arising from market power. A simple heuristic to mitigate efficiency loss is to let sellers

---

[49]Sellers have an equivalent rule to Proposition 9: $\dot{T}_S > \ddot{T} \iff \frac{dE[P(N,T)]}{dT}|_{T=\dot{T}} > 0$. This means that either 1) $\dot{T}_S > \ddot{T} > \dot{T}_B$, 2) $\dot{T}_B > \ddot{T} > \dot{T}_S$, or 3) $\dot{T}_S > \ddot{T} \cap \dot{T}_B > \ddot{T}$. The case where both the buyer-optimal and seller-optimal contract are shorter than the efficient contract is ruled out by the fact that per-period costs must be increasing at the efficient contract for an interior solution.

determine contract duration when the duration affects price more than buyer surplus, and to let buyers determine contract duration otherwise.

These heuristics, combined with Proposition 9, provide insight into which settings may allow for substantive inefficiencies and whether the efficient contract is longer or shorter. Below, I provide a simple example to illustrate how changing the allocation of rights over duration may lead to vastly different outcomes.

Example: Markup Pricing Suppose sellers in equilibrium follow a simple markup pricing rule, $P = \mu C$. Then the buyer's problem is

$$\min_{T} \mu E[C(N,T)] + \frac{\delta}{T}$$

and the seller's problem is

$$\min_{T}(1 - \mu)E[C(N,T)] + \frac{\delta}{T}$$

As $\mu \geq 1$ in equilibrium, the seller's problem reverses the sign that expected costs enter in the objective function. By increasing costs, sellers increase total profits. In this setting, the buyer should determine the duration. The greater the markup, the more that the equilibrium contract will diverge from the efficient contract.

A.3 Achieving Efficiency with a Tax

The efficient contract can be achieved with a per-transaction tax (or subsidy) when either side of the transaction holds the term-setting rights. When the buyer determines the length of the contract, the efficient per-transaction tax τ_B solves

$$\tau_B = T^2 \frac{dE[\pi(N,T)]}{dT}\Big|_{T=\ddot{T}}$$

This tax equates the buyer's problem with the social planner's problem. Note below how the tax causes the externality on the seller to drop out at the efficient contract.

$$
\begin{aligned}
\ddot{T} &= \arg\min_{T} E[\pi(N,T)] + E[C(N,T)] + \frac{\delta + \tau_B}{T} \\
&= \arg\min_{T} E[\pi(N,T)] + \frac{\tau_B}{T} + E[C(N,T)] + \frac{\delta}{T} \\
&= \arg\min_{T} E[C(N,T)] + \frac{\delta}{T}
\end{aligned}
$$

Analogously, the efficient tax on the seller (when the seller has term-setting rights) is given by

$$\tau_S = -T^2 \frac{dE[P(N,T)]}{dT}\Big|_{T=\ddot{T}}$$

44

In general, $\tau_S \neq \tau_B$. A policymaker has a choice between two efficient taxes, with different effects on tax revenue.

B Model Proofs

B.1 Proof of Proposition 1

Rearranging the FOC and taking the total derivative, we obtain

$$2T\frac{d\overline{P}(T,x,m)}{dT} + T^2\frac{d^2\overline{P}(T,x,m)}{dT^2} = \frac{d\delta}{dT}$$

Proof. From the second-order condition for a minimum,

$$\frac{d^2\overline{P}(T,x,m)}{dT^2} > -2\frac{\delta}{T^3}$$

.

Therefore

$$\frac{d\delta}{dT} > 2T\frac{d\overline{P}(T,x,m)}{dT} - 2\frac{\delta}{T}$$

As the RHS is equal to zero, $\frac{dT}{d\delta} > 0$. □

B.2 Proof of Proposition 2

Proof. Taking the total derivative of the first-order condition with respect to M and solving for $\frac{dT}{dM}$ produces

$$\frac{dT}{dM} = \frac{-\frac{d^2\overline{P}(T,X,M)}{\partial T\partial M}}{\frac{\partial^2 d\overline{P}(T,X,M)}{\partial T^2} + \frac{2\delta}{T^3}}. \tag{6}$$

The denominator is positive, as it is the second-order condition to ensure a minimum. Therefore, we have the simple relation

$$\mathrm{sgn}\left(\frac{dT}{dM}\right) = \mathrm{sgn}\left(-\frac{d^2\overline{P}(T,X,M)}{\partial T\partial M}\right).$$

Likewise, this also holds for X. □

C Identification Proofs

To demonstrate this proof, it will be useful to first introduce several lemmas.

Lemma 1. *For symmetric auctions with independent private values, $E[b_{1:N}] = E[c_{2:N}]$.*

This is a standard result and can be obtained directly by taking the expectation given the equilibrium bid function. I omit the proof here.

Lemma 2. $\min b_{1:N} = E[c_{1:(N-1)}]$ *for the IPV model when the support of c is bounded from below by $\underline{c} > -\infty$.*

Proof. The equilibrium bid function is given by

$$\beta(c; N) = c + \frac{\int_c^\infty [1 - F(\xi)]^{N-1} d\xi}{[1 - F(c)]^{N-1}}$$

Then the minimum bid is

$$
\begin{aligned}
\beta(\underline{c}; N) &= \underline{c} + \frac{\int_{\underline{c}}^\infty [1 - F(\xi)]^{N-1} d\xi}{[1 - F(\underline{c})]^{N-1}} \\
&= \underline{c} + \int_{\underline{c}}^\infty [1 - F(\xi)]^{N-1} d\xi \\
&= \underline{c} + \xi[1 - F(\xi)]^{N-1}\big|_{\underline{c}}^\infty + \int_{\underline{c}}^\infty \xi(N-1)f(\xi)[1 - F(\xi)]^{N-2} d\xi \\
&= \underline{c} + (0 - \underline{c}) + \int_{\underline{c}}^\infty \xi(N-1)f(\xi)[1 - F(\xi)]^{N-2} d\xi \\
&= E[c_{1:(N-1)}]
\end{aligned}
$$

Where the third line comes from integration by parts. Here we require the assumption that $\lim_{\xi \to \infty} f(\xi)[1 - F(\xi)]^N = 0$, so that

$$
\begin{aligned}
\xi[1 - F(\xi)]^{N-1}\big|_{\underline{c}}^\infty &= \lim_{\gamma \to 0} \frac{[1 - F(\frac{1}{\gamma})]^{N-1}}{\gamma} - \underline{c}[1 - F(\underline{c})]^{N-1} \\
&= \lim_{\gamma \to 0} \frac{-(N-1)f(\frac{1}{\gamma})[1 - F(\frac{1}{\gamma})]^{N-2}}{1} - \underline{c} \\
&= 0 - \underline{c}
\end{aligned}
$$

$\square$

Lemma 3. *The expected k-th order statistic of N draws can be written in terms of the expected k-th and $(k+1) - th$ order statistics from $N + 1$ draws: $E[c_{k:N}] = \frac{k}{N+1} E[c_{(k+1):(N+1)}] + \frac{N+1-k}{N+1} E[c_{k:(N+1)}]$*

Proof. First, examining the difference between the k-th order statistics of N and $N + 1$ draws. Expressing $E[c_{k:N}] - E[c_{k:(N+1)}]$ and rearranging terms gives:

$$
\begin{aligned}
&E[c_{k:N}] - E[c_{k:(N+1)}] \\
&= \int \frac{N!}{(k-1)!(N-k)!} cf(c)F(c)^{k-1}[1 - F(c)]^{N-k} dc - \int \frac{(N+1)!}{(k-1)!(N+1-k)!} cf(c)F(c)^{k-1}[1 - F(c)]^{N+1-k} dc \\
&= \int \left(\frac{N!(N+1-k)}{(k-1)!(N+1-k)!} - \frac{(N+1)!}{(k-1)!(N+1-k)!}[1 - F(c)] \right) cf(c)F(c)^{k-1}[1 - F(c)]^{N-k} dc \\
&= \int \frac{(N+1)!}{(k-1)!(N+1-k)!} cf(c)F(c)^k [1 - F(c)]^{N-k} dc - \int \frac{kN!}{(k-1)!(N+1-k)!} cf(c)F(c)^{k-1}[1 - F(c)]^{N-k} dc \\
&= \frac{k}{(N+1-k)} \left(E[c_{(k+1):(N+1)}] - E[c_{k:N}] \right)
\end{aligned}
$$

Rearranging, we obtain

47

$$E[c_{k:N}] = \frac{k}{N+1} E[c_{(k+1):(N+1)}] + \frac{N+1-k}{N+1} E[c_{k:(N+1)}]. \hspace{2cm} \square$$

C.1 Proof of Proposition 3

Consider the entry equation

$$E[\pi_n|n,t] \cdot h(x) \cdot U - k(m) \cdot \varepsilon > 0 \iff N \geq n$$

$$E[\pi_n|n,t] \cdot \frac{h(x)}{k(m)} > \frac{\varepsilon}{U} \iff N \geq n$$

For any realization (t, x, m), $\exists (t, x', m')$ such that $U|(N, t, x, m) = U|(N, t, x', m')$.[50] Using this fact, we can identify $h(X)$ by finding (x', m') such that $\Pr(N \geq n|t, x, m) = \Pr(N \geq n|t, x', m')$ for all N, then calculating

$$\frac{E[B \cdot U \cdot h(x)|N,t,x,m]}{E[B \cdot U \cdot h(x')|N,t,x',m']} = \frac{E[B|N,t] \cdot E[U|N,t,x,m] \cdot h(x)}{E[B|N,t] \cdot E[U|N,t,x',m'] \cdot h(x')} = \frac{h(x)}{h(x')}$$

$k(M)$ is identified by finding

$$\frac{h(x)}{k(m)} = \frac{h(x')}{k(m')}$$

For a particular realization (n_0, t_0, x_0, m_0), normalize $h(x_0) = 1$, and $k(m_0) = 1$, and $E[U|n_0, t_0, x_0, m_0] = 1$. This pins down the scale of $E[B|n_0, t_0]$ from the observed transaction price, and the scale of ε is identified from the participation equation. Once h and k are identified, the distribution of $\frac{\varepsilon}{U}$ is identified directly from the participation equation and continuous variation in either X or M.

Now that h, k, and the distribution of $\frac{\varepsilon}{U}$ are identified, I turn to the identification of the unobserved components of transaction price.

C.1.1 Identification of Offers and Unobserved Heterogeneity

For any realization (n, t), the expected offer can be identified by finding (x, m) such that $U|(n, t, x, m) = U|(n_0, t_0, m_0, x_0)$. Again, the pair (x, m) is found by setting $\Pr(N \geq n|t, m, x) = \Pr(N \geq n_0|t_0, m_0, x_0)$. At (x, m), the mean transaction price is equal to the expected offer scaled by $h(X)$, which is now known:

$$E[P|n, t, m, x] = E[B|n, t] \cdot E[U|n, t, x, m] \cdot h(X)$$

$$= E[B|n, t] \cdot h(X)$$

As $E[B|N, T]$ is identified for any (n, t), $E[U|N, T, X, M]$ is identified from the mean transaction price at any realization of (N, T, X, M).

---

[50]Here, and once more in the proof, I rely on either $h(\cdot)$ or $k(\cdot)$ having broad support.

To identify surplus, consider the entry condition:

$$E[\pi_n|n,T] \cdot \frac{h(X)}{k(M)} > \frac{\varepsilon}{U}$$

For every (n,x,m) and $n' \neq n$, $\exists(x',m')$ such that

$$E[\pi_n|n,T] \cdot \frac{h(x)}{k(m)} = E[\pi_{n'}|n',T] \cdot \frac{h(x')}{k(m')}$$

As $h(\cdot)$ and $k(\cdot)$ are identified, $\frac{E[\pi_n|n,T]}{E[\pi_{n'}|n',T]} = R$ is identified. Likewise, relative profits $\frac{E[\pi_n|n,t]}{E[\pi_n|n,t']}$ are identified.

C.2   Proof of Proposition 6

The ratio of profits is given by

$$R = \frac{E[\pi_n|n,T]}{E[\pi_{n'}|n',T]} = \frac{\frac{1}{n}\left(E[B|n,T] - E[C|n,T]\right)}{\frac{1}{n'}\left(E[B|n',T] - E[C|n',T]\right)} \tag{7}$$

When the selection mechanism is a symmetric auction. $E[B|n,T] = E[C_{2:n}|T]$ and $E[C|n,T] = E[C_{1:n}|T]$. From here on I suppress notation indicating that costs are conditional on T. From Lemma (3), we have $E[C_{1:n}] = \frac{1}{n+1}E[C_{2:(n+1)}] + \frac{n}{n+1}E[C_{1:(n+1)}]$. Plugging this into the equation for R obtains

$$R\left(E[C_{2:(n+1)}] - E[C_{1:(n+1)}]\right) = E[C_{2:n}] - \frac{1}{n+1}E[C_{2:(n+1)}] - \frac{n}{n+1}E[C_{1:(n+1)}]$$

$$\left(R + \frac{n}{n+1}\right)E[C_{1:(n+1)}] = E[C_{2:n}] - \left(R + \frac{1}{n+1}\right)E[C_{2:(n+1)}]$$

Therefore, $E[C_{1:(n+1)}]$ is identified. $E[C_{1:n}]$ is obtained from equation (7).

C.3   Proof of Proposition 7

For each observed sequential value of $N \in \{\underline{N}, ..., \overline{N}\}$, the first-order and second-order statistics of N draws from the cost distribution are identified. Using the recursive relationship of order statistics shown in Lemma 3, these are equivalent to identifying the first $\overline{N} - \underline{N} + 2$ expected order statistics from $\overline{N}$ draws of C.

C.4 Proof of Proposition 8: Identification with No Instrument

The ratio of second-order statistics is identified by comparing winning bids for different values of n and n'.

$$\frac{E[Y|n,T,X,M]}{E[Y|n',T,,X,M]} = \frac{E[B_n|T,X,M] \cdot E[U|n,T,X,M]}{E[B_{n'}|T,X,M] \cdot E[U|n',T,X,M]} = \frac{E[C_{2:n}|T,X,M]}{E[C_{2:n'}|T,X,M]}$$

Where $E[U|n,T,X,M] = E[U|n',T,X,M] = E[U|T,X,M]$ by independence and no selection on unobservables.

From here on, C_i and U may be conditional on (T,X,M). I suppress this in my notation for clarity. Normalizing $E[U] = 1$ pins down the scale of $E[C_{2:n}]$. [51]

Suppose that another $(\hat{F}, \hat{G})$ rationalizes the data. Then

$$B_n \cdot U \overset{d}{=} \hat{B}_n \cdot \hat{U}$$
$$B_{n'} \cdot U \overset{d}{=} \hat{B}_{n'} \cdot \hat{U}$$

Construct $\tilde{b}_{n'}$, $\hat{\tilde{b}}_{n'}$, $\tilde{U}$, and $\hat{\tilde{U}}$ as random variables that are independent of and have the same conditional distributions as their tilde-free counterparts. Then it follows that

$$(B_n \cdot U) \cdot \left(\hat{\tilde{B}}_{n'} \cdot \hat{\tilde{U}}\right) \overset{d}{=} \left(\hat{B}_n \cdot \hat{U}\right) \cdot \left(\tilde{B}_{n'} \cdot \tilde{U}\right)$$
$$\implies B_n \cdot \hat{\tilde{B}}_{n'} \overset{d}{=} \hat{B}_n \cdot \tilde{B}_{n'}$$

From this relation, we may take the minimum on both sides. By independence and Lemma 2, I obtain

$$E[C_{1:(n-1)}] \cdot E[\hat{C}_{1:(n'-1)}] = E[\hat{C}_{1:(n-1)}] \cdot E[C_{1:(n'-1)}]$$
$$\frac{E[C_{1:(n-1)}]}{E[C_{1:(n'-1)}]} = \frac{E[\hat{C}_{1:(n-1)}]}{E[\hat{C}_{1:(n'-1)}]}$$

That is, any $(\hat{F}, \hat{G})$ that rationalizes the data has a private cost distribution with the same ratio of first order statistics.

Finally, using the fact that $E[C_{1:(n-1)}] = \frac{1}{n}E[C_{2:n}] + \frac{n-1}{n}E[C_{1:n}]$, we can link together these ratios when $n' = n + 1$.

$$\frac{\frac{1}{n}E[C_{2:n}] + \frac{n-1}{n}E[C_{1:n}]}{E[C_{1:n}]} = \frac{\frac{1}{n}E[\hat{C}_{2:n}] + \frac{n-1}{n}E[\hat{C}_{1:n}]}{E[\hat{C}_{1:n}]}$$
$$\implies \frac{E[C_{2:n}]}{E[C_{1:n}]} = \frac{E[\hat{C}_{2:n}]}{E[\hat{C}_{1:n}]}$$

---

[51]Note that, in practice, we may normalize $E[U|t,x,m] = 1$ for all (t,x,m) realizations. How the mean of $C_{2:n} \cdot U$ changes is captures in changes to the mean of C.

As we have identified $E[C_{2:n}]$, $E[C_{1:n}]$ and $E[C_{1:(n-1)}]$ is also identified. Seller surplus is obtained. With sequential values of $N \in \{\underline{N}, ..., \overline{N}\}$, the recursive relationship between order statistics from Lemma 3 gives the first $\overline{N} - \underline{N} + 2$ expected order statistics from $\overline{N}$ draws of C from the identified first-order and second-order statistics.

C.4.1 Identification of Entry Costs

Now that $\{E[\pi_n(X)]\}$ and $E[U|N,T,X,M]$ are identified, we can identify entry costs $k(M)$ from the entry equation.

$$E[\pi_n|T,X,M] \cdot E[U|n,T,X,M] > k(M) \cdot \varepsilon \iff N \geq n$$

Additional assumptions:

1. $\varepsilon \perp\!\!\!\perp (X, M)$

2. $E[\pi_n|T,X,M]$ varies continuously in X conditional on M.

Normalize $k(m_0) = 1$. To identify $k(M)$, find (x, m) and (x', m_0) for any m so that $\Pr(N = n|t,x,m) = \Pr(N = n, |t, x', m_0)$. Then

$$\frac{k(m)}{k(m_0)} = \frac{E[\pi_n|t,x,m] \cdot E[U|n,t,x,m]}{E[\pi_n|t,x',m_0] \cdot E[U|n,t,x',m_0]}$$

Finally, the distribution of ε is identified once $k(M)$ is identified.

C.5 Non-symmetric Entrants

Assume that the winning type is not observed. The expected winning bid across types is identified.

$$\frac{E[Y|n,T,X,M]}{E[Y|n',T,,X,M]} = \frac{E[B_n|T,X,M] \cdot E[U|n,T,X,M]}{E[B_{n'}|T,X,M] \cdot E[U|n',T,X,M]} = \frac{E[B_n|T,X,M]}{E[B_{n'}|T,X,M]}$$

Again, normalize $E[U|T,X,M] = 1$ for all realizations of (T,X,M).

Additional assumptions:

1. $\varepsilon \perp\!\!\!\perp (X, M)$

2. $E[\pi_n|T,X,M]$ varies continuously in X conditional on M.

3. $(C_i, U) \perp\!\!\!\perp (X, M)$. [Independence, Exclusion restriction]

4. (C_i, U) and X are multiplicatively separable.

Then,

$$\frac{E[Y|n,T,x,M]}{E[Y|n,T,x',M]} = \frac{E[B_n|n,T] \cdot E[U|n,T,M,x] \cdot h(x)}{E[B_{n'}|n,T] \cdot E[U|n,T,M,x'] \cdot h(x')} = \frac{h(x)}{h(x')}$$

Normalize $h(x_0) = 1$. Then $h(x)$ is identified and the scale of $E[B_n|N,T]$ is pinned down.

To identify $k(M)$, find (n,t,x,m) and (n,t,x',m_0) for any m so that $\Pr(N = n|t,x,m) = \Pr(N = n, |t,x',m_0)$. Then

$$\frac{k(m)}{k(m_0)} = \frac{E[\pi_n|t] \cdot h(x)}{E[\pi_n|t] \cdot h(x')} = \frac{h(x)}{h(x')}$$

After normalizing $k(m_0) = 1$, $k(M)$ is identified. Once h and k are identified, the distribution of ε is identified directly from the participation equation and continuous variation in either X or M.

Relative profits are identified. For any (n,t,x,m), find (n',t',x',m') so that $\Pr(N = n|t,x,m) = \Pr(N = n'|t',x',m')$. Then

$$\frac{E[\pi_n|t]}{E[\pi_{n'}|t']} = \frac{h(x')}{k(m')} \frac{k(m)}{h(x)}$$

To pin down profits, one observation of actual profits would be sufficient. Another approach is to impose additional structure (e.g. symmetry as above).

C.6 Selective Entry and Non-Symmetry

1. $\varepsilon \perp\!\!\!\perp (X,M)$

2. $E[\pi_n|T,X,M]$ varies continuously in X conditional on M.

3. $(C_i, U) \perp\!\!\!\perp (X,M)$. [Independence, Exclusion restriction]

4. (C_i, U) and X are multiplicatively separable.

Then, find data with (n,t,x,m) and (n,t,x',m') where $\Pr(N = n|t,x,m) = \Pr(N = n, |t,x',m')$. By independence of U and ε, $E[U|n,T,x,m] = E[U|n,T,x',m']$ and:

$$\frac{E[Y|n,T,x,m]}{E[Y|n,T,x',m']} = \frac{E[B_n|n,T] \cdot E[U|n,T,x,m] \cdot h(x)}{E[B_n|n,T] \cdot E[U|n,T,x',m'] \cdot h(x')} = \frac{h(x)}{h(x')}$$

Normalize $h(x_0) = 1$. Then $h(x)$ is identified and the scale of $E[B_n|N,T]$ is pinned down.

To identify $k(M)$, find (n,t,x,m) and (n,t,x',m_0) for any m so that $\Pr(N = n|t,x,m) = \Pr(N = n, |t,x',m_0)$. Then

$$\Pr(E[\pi_n|t] \cdot h(x) \cdot U > k(m) \cdot \varepsilon) = \Pr(E[\pi_{n'}|t] \cdot h(x') \cdot U > k(m') \cdot \varepsilon)$$

$$F_{\frac{\varepsilon}{U}|t}\left(E[\pi_n|t] \cdot \frac{h(x)}{k(m)}\right) = F_{\frac{\varepsilon}{U}|t}\left(E[\pi_n|t] \cdot \frac{h(x')}{k(m')}\right)$$

By independence,

$$\frac{k(m)}{k(m_0)} = \frac{E[\pi_n|t] \cdot h(x)}{E[\pi_n|t] \cdot h(x')} = \frac{h(x)}{h(x')}$$

After normalizing $k(m_0) = 1$, $k(M)$ is identified. Once h and k are identified, the distribution of $\frac{\varepsilon}{U}$ is identified directly from the participation equation and continuous variation in either X or M.

Relative profits are identified. For any (n, t, x, m), find (n', t', x', m') so that $\Pr(N = n|t, x, m) = \Pr(N = n'|t', x', m')$. Then

$$\frac{E[\pi_n|t]}{E[\pi_{n'}|t']} = \frac{h(x')}{k(m')} \frac{k(m)}{h(x)}$$

To pin down profits, one observation of actual profits would be sufficient. Another approach is to impose additional structure (e.g. symmetry as above).

D A Model with Microfoundations

In the empirical application of this paper, I employ a "reduced-form" approach to a capturing how the distribution of private costs changes with T. Here, I will provide a model of underlying costs that generates both the distribution of costs and how the size of the bundle shapes the distribution. Suppose that instantaneous costs follow an Ornstein-Uhlenbeck diffusion process. The continuous-time cost process X_t is governed by the differential equation

$$dx_t = \theta(\mu - x_t) + \sigma dW_t$$

where W_t is a Wiener process. This process is stationary over t. That is, any bundle of size T will have the same unconditional distribution as any other bundle of size T. Define the average cost over time T as

$$c_T = \frac{1}{T} \int X_t dt$$

Then c_T is Gaussian with mean μ and variance $\frac{1}{T^2} \frac{\sigma^2}{\theta^3} \left(\theta T + e^{-\theta T} - 1 \right)$. When costs are Gaussian, $E[c_{1:N}(\sigma)] = E[z_{1:N}]\sigma + \mu$, where z is a standard normal. Define $\xi : T \to \sigma$. The efficient bundle T solves

$$\min_T E[z_{1:N}]\xi(T) + \mu + \frac{\delta}{T}$$

This results in the first-order condition

$$
\begin{aligned}
E[z_{1:N}]\xi'(T) &= \frac{\delta}{T^2} \\
-\xi'(T)T^2 &= -\frac{\delta}{E[z_{1:N}]}
\end{aligned}
\tag{8}
$$

Lemma 4. $\frac{d}{dT}\left(-\xi'(T)T^2\right) > 0.$

As shown in the appendix, $\frac{d}{dT}\left(-\xi'(T)T^2\right) > 0$. Therefore,

Proposition 11. *The efficient bundle size is decreasing in the number of bidders.*

Proof. $\frac{d}{dT}\left(-\xi'(T)T^2\right) = -2T \cdot \xi'(T) - T^2 \xi''(T)$. Combining the second-order conditions and first-order conditions, we obtain.

$$
\begin{aligned}
E[z_{1:N}]\xi''(T) &> -\frac{2}{T}E[z_{1:N}]\xi'(T) \\
\implies T^2 \cdot \xi''(T) &< -2T\xi'(T)
\end{aligned}
$$

An increase in N increases the RHS of equation 8. As $\frac{d}{dT}\left(-\xi'(T)T^2\right) < 0$, the optimal T falls. $\qquad \square$

Unlike the U-shape models, the microfounded model here does not have a lower bound on costs. Further, the analysis above holds for the second-order statistic when $N > 3$, so we can extend the

results to the buyer-optimal bundle:[52]

Proposition 12. *The buyer-optimal bundle size is decreasing in the number of bidders. It is optimal to sell everything in a single bundle for $N \in \{2, 3\}$.*

Additionally, in we have that $E[z_{1:N}] < E[z_{2:N}]$. Therefore,

Proposition 13. *The efficient bundle size is smaller than the buyer-optimal bundle.*

---

[52]For $N \in \{2, 3\}$, $E[z_{2:N}] > 0$.

E   Likelihood Function

For estimation, we obtain the likelihoods for Y_n and N given by

$$f_{Y_n|N,X,T,M} = \int f_{B_n|T,N}\left(\frac{y}{U}\frac{1}{h(X)}\right)\frac{1}{U}\frac{1}{h(X)}f_{U|N,T,X,M}(U)dU$$

$$\Pr(N = n|T, X, M) = \int \Pr(N = n|U, T, X, M)f_{U|T,X,M}(U)dU$$

For estimation, I make the assumption that $U \perp\!\!\!\perp (X, M)$. As U is not observed by the procurer when setting T, $U \perp\!\!\!\perp (T, X, M)$. This simplifies the problem so that $f_{U|T,X,M}(U) = f_U(U)$. The conditional distribution of U used in the likelihood of Y_N is given by $f_{U|N,T,X,M}(u) = \frac{\Pr(N=n|U,T,X,M)f_U(u)}{\Pr(N=n|T,X,M)}$. This simplifies so that the joint contribution is given by

$$
\begin{aligned}
f_{Y_n|N,X,T,M}(y_n) \cdot \Pr(N = n|T, X, M) &= \left(\int f_{B_n|T,N}\left(\frac{y}{u}\frac{1}{h(X)}\right)\frac{1}{u}\frac{1}{h(X)}f_{U|N,T,X,M}(u)du\right)\Pr(N = n|T, X, M) \\
&= \left(\int f_{B_n|T,N}\left(\frac{y}{u}\frac{1}{h(X)}\right)\frac{1}{u}\frac{1}{h(X)}\frac{\Pr(N = n|u, T, X, M)f_U(u)}{\Pr(N = n|T, X, M)}du\right)\Pr(N = n|T, X, M) \\
&= \int f_{B_n|T,N}\left(\frac{y}{u}\frac{1}{h(X)}\right)\frac{1}{u}\frac{1}{h(X)}\Pr(N = n|u, T, X, M)f_U(u)du
\end{aligned}
$$

With the assumption that the shock ε is independent of (U, T, X, M), we have the following expression for conditional probability of N.

$$
\begin{aligned}
\Pr(N = n|U, T, X, M) &= F_{\ln \varepsilon}\left(\ln E[\pi_n|T] + \ln h(X) + \ln U - \ln k(M)\right) \\
&\quad - F_{\ln \varepsilon}\left(\ln E[\pi_{n+1}|T] + \ln h(X) + \ln U - \ln k(M)\right)
\end{aligned}
$$

I use the joint likelihood of Y_n and N to obtain estimates for cost and entry parameters.

E.1   A Computation Innovation

In this setting, there is a symmetric equilibrium in which each bidder has a monotone bid function $\beta(\cdot; n)$ mapping private costs to the submitted bid. The density of an observed bid is given by

$$f_{B_n}(b) = f_c(\beta^{-1}(b; n))\frac{1}{\beta'(\beta^{-1}(b; n))}$$

In maximum likelihood estimation of the cost distribution, it is necessary to invert the bid function to calculate the density. This can be computationally intensive when β does not have a closed-form solution.

In the presence of unobserved heterogeneity, the density of the observed bid $\tilde{B} = B \cdot U$ is given

by the convolution when $B \perp U$.

$$
\begin{aligned}
f_{\tilde{B}}(\tilde{b}) &= \int_{\underline{u}}^{\bar{u}} f_B\left(\frac{\tilde{b}}{u}\right) \frac{1}{u} f_u(u) du \\
&= \int_{\underline{u}}^{\bar{u}} f_c\left(\beta^{-1}\left(\frac{\tilde{b}}{u}; n\right)\right) \frac{1}{\beta'\left(\beta^{-1}\left(\frac{\tilde{b}}{u}; n\right)\right)} \frac{1}{u} f_u(u) du
\end{aligned}
$$

Here, the computational burden increases greatly. Integrating out the unobserved heterogeneity means that the bid function must be inverted for each value of u within the integral, in order to calculate $\beta^{-1}\left(\frac{\tilde{b}}{u}; n\right)$. As the inverse bid function has an analytic solution for only a few specialized cases, in practice this computation relies on a non-linear equation solver or an approximation. Thus, the calculations are constrained by the efficiency and accuracy of such an approach.

One easy-to-implement solution that makes maximum likelihood significantly more tractable is to use a change-of-variables to calculate the density. Instead of integrating out the unobserved heterogeneity by integrating over u, replace u with $u = \frac{\tilde{b}}{\beta(c)}$ and integrate over c. The density then becomes:

$$
\begin{aligned}
f_{\tilde{B}}(\tilde{b}) &= \int_{\underline{u}}^{\bar{u}} f_C\left(\beta^{-1}\left(\frac{\tilde{b}}{u}\right)\right) \frac{1}{\beta'\left(\beta^{-1}\left(\frac{\tilde{b}}{u}\right)\right)} \frac{1}{u} f_u(u) du \\
&= \int_{\psi^{-1}(\underline{u})}^{\psi^{-1}(\bar{u})} f_C\left(\beta^{-1}\left(\beta(c)\right)\right) \frac{1}{\beta'\left(\beta^{-1}\left(\beta(c)\right)\right)} \frac{\beta(c)}{\tilde{b}} f_u\left(\frac{\tilde{b}}{\beta(c)}\right)\left(-\frac{\tilde{b}}{\beta(c)^2}\beta'(c)\right) dc \\
&= \int_{\underline{c}}^{\bar{c}} f_C\left(c\right) f_u\left(\frac{\tilde{b}}{\beta(c)}\right) \frac{1}{\beta(c)} dc
\end{aligned}
$$

Note that in this form, there is no need to invert the bid function. As the general form for the symmetric equilibrium bid function is

$$
\beta(c) = c + \frac{\int_c^{\infty}[1 - F(z)]^{n-1}}{[1 - F(c)]^{n-1}},
$$

the primary computational cost is a numerical integration routine. Therefore, the model is computationally tractable for a vast class of parametric distributions of C and U, as well as non-parametric approximations such as B-splines.

F   Incumbency and Asymmetries

In this section, I present regressions for the dependent variables of price and the number of bids, including an indicator for whether or not a single incumbent bidder was identified from a previous contract. That is, the indicator equals one if janitorial services for the same agency and 9-digit ZIP were performed by a single supplier in the previous year. The coefficient on this variable is not significant, and its inclusion does not meaningfully impact the estimated coefficients.

Table 12: Descriptive Regressions: ln(Price per Month)

	(1)	(2)	(3)	(4)	(5)
Number of Bids	−0.008	−0.030***	−0.025***	−0.030***	−0.025***
	(0.010)	(0.007)	(0.007)	(0.007)	(0.007)
Duration (Years)	0.059*	0.025	−0.011	0.025	−0.011
	(0.032)	(0.021)	(0.021)	(0.021)	(0.021)
ln(Square Footage)		0.530***	0.590***	0.531***	0.589***
		(0.031)	(0.033)	(0.031)	(0.033)
Weekly Frequency		0.179***	0.115***	0.179***	0.116***
		(0.018)	(0.022)	(0.018)	(0.022)
Incumbent Ind.				−0.012	0.058
				(0.064)	(0.063)
Site Type FEs			X		X
Agency FEs			X		X
Fiscal Year FEs			X		X
Observations	420	420	420	420	420
R^2	0.01	0.57	0.66	0.57	0.66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays regression results for regressions of log monthly price on auction characteristics. The variables from specification (2) are included in the structural model. These regressions show that square footage and cleaning frequency are important in explaining variation in prices. Once square footage and cleaning frequency are accounted for, fixed effects for fiscal year, agency, and site type add little explanatory power. In (4) and (5), I include an indicator for the presence of a single incumbent.

Table 13: Descriptive Regressions: Number of Bids

	(1)	(2)	(3)	(4)	(5)	(6)
Duration (Years)	0.313**	0.347**	0.234*	0.117	0.235*	0.117
	(0.154)	(0.147)	(0.139)	(0.144)	(0.139)	(0.144)
ln(Square Footage)	0.727***	0.695***	0.794***	0.846***	0.781***	0.840***
	(0.221)	(0.211)	(0.199)	(0.218)	(0.200)	(0.218)
Weekly Frequency	−0.064	−0.119	−0.291**	−0.079	−0.282**	−0.077
	(0.130)	(0.124)	(0.119)	(0.147)	(0.120)	(0.147)
ln(Unemp./(2004 Unemp.))		4.198***				
		(0.657)				
ln(Unemployment)			4.019***	2.822***	3.995***	2.827***
			(0.617)	(0.761)	(0.618)	(0.762)
ln(2004 Unemp.)			−3.119***	−2.070***	−3.077***	−2.066***
			(0.633)	(0.781)	(0.636)	(0.782)
Incumbent Ind.					0.296	0.194
					(0.424)	(0.428)
Site Type FEs				X		
Agency FEs				X		
Fiscal Year FEs				X		
Observations	420	420	420	420	420	420
R^2	0.04	0.12	0.23	0.34	0.23	0.34
F-statistic	5.1	14.4	24.4	7.4	20.4	7.1

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays regression results for regressions of the number of bids on auction characteristics and local labor market variables. Specification (2) is equivalent to the first-stage regression of IV-1 in Table 12. In (3) and (4), I split the instrument into the current component and the baseline (2004) component. In (5) and (6), I include an indicator for the presence of a single incumbent.

G   Count of Sites by Government Agency

Table 14: Count of Sites by Contracting Agency

Agency	Count	Percent
Defense	151	36.0
Agriculture	148	35.2
Commerce	56	13.3
Veterans Affairs	19	4.5
Homeland Security	19	4.5
Interior	11	2.6
GSA	9	1.2
Energy	3	0.7
EPA	2	0.5
Transportation	2	0.5
Labor	1	0.2
Railroad Retirement Board	1	0.2
Health And Human Services	1	0.2
National Archives	1	0.2
Total	420	100.0

H Contract Documents

The following page is an example first page from a janitorial service contract. The subsequent pages contain an example description of the required services and their respective frequencies.

CONTRACT DOCUMENTS, EXHIBITS OR ATTACHMENTS

C.1 SCOPE OF CONTRACT

Description of Work: The intent of this contract is to secure services (inclusive of supplies) for normal custodial (janitorial) and routine maintenance service at the Georgetown Ranger District of the Eldorado National Forest.

2 Project Location & Description

Location: The project is located on the Georgetown Ranger District, 7600 Wentworth Springs Road, Georgetown, CA 95634.

Description: The headquarters office of the Georgetown Ranger District is located at 7600 Wentworth Springs Road, Georgetown, California. Winter working hours are 6:00 a.m. through 5:30 p.m. Monday through Friday from November through May. Summer hours are 7:00 a.m. through 6:00 p.m. Sunday through Saturday.

The office building contains approximately 6,376 gross square feet of space. The office is carpeted throughout, expect for restrooms and front reception area. There are 6 restrooms in the building.

Any prospective contractor desiring an explanation or interpretation of the solicitation, drawings, specifications, etc., must request it in writing from the Contracting Officer soon enough to allow a reply to reach all prospective contractors before the solicitation closing date. Oral explanations or instructions given before the award of a contract will not be binding.

3 Estimated Start Date & Contract Time

Start: January 1, 2010
Time: 9 Months

4 Cleaning Schedule

Work Days and Hours. Work shall be performed during Monday through Friday, provided that no work is performed between 7 a.m. and 4:30 p.m. on normal Federal workdays. Regularly scheduled twice weekly work will not be on consecutive days. The contractor may work in the building on weekends and Federal holidays without restrictions to hours.

Quarterly cleaning items will be performed the first week (preferably on Friday) of December, March, June, and September. Annual cleaning shall be performed during the first 2 weeks of May.

5 Licenses and Insurance

Contractor shall provide proof of Workman's Compensation. If the contractor is working alone, with no employees, no Workman's Compensation is required.
.

6 Contractor-Furnished Materials and Services

6-1. The Contractor shall provide everything--including, but not limited to, all equipment, supplies (listed below), transportation, labor, and supervision--necessary to complete the project, except for that which the contract clearly states is to be furnished by the Government.

18. TECHNICAL SPECIFICATIONS

The janitorial services shall be performed in accordance with the following specifications at the frequencies prescribed.

1. Services Performed Daily - Bid Item #0001

    a. Restrooms
       - Clean and sanitize all surfaces including sinks, counters, toilet bowls, toilet seats, urinals, etc.
       - Clean and sanitize tile walls adjacent to and behind urinals and water closets.
       - Clean and sanitize sanitary napkin receptacles and replace liners.
       - Sweep, mop and sanitize tile floors.
       - Clean and polish mirrors, dispensers and chrome fixtures
       - Empty, clean and sanitize all wastebaskets.
       - Spot clean all other surfaces and dust horizontal surfaces including tops of partitions and mirrors.
       - Re-stock restroom supplies.

    b. Front Foyer and Doors
       - Wash inside and outside of all glass surfaces on entrance doors. Remove dust and soil from metal frames surrounding entrance glass doors.
       - Vacuum rugs.
       - Sweep and mop tile floors and clean baseboards.

    c. Reception Area
       - Vacuum all reception carpeted areas and rugs including edges.
       - Clean and polish all counter surfaces.

    d. Drinking Fountains
       - Clean and sanitize drinking fountains.

    e. Breakroom Waste Receptacles
       - Empty all waste receptacles, wash if needed with a sanitizing cleaner.

2. Services Performed Weekly – Bid Item #0002

    a. Waste Receptacles
       - Empty all waste receptacles unless needed more frequently. Wash if needed with a sanitizing cleaner. Change liners only if needed.

    b. Breakroom
       - Sweep and mop, use a cleaner that doesn't require rinsing and is a sanitizer and will not damage the wax. Mop under table, chairs, coffeemaker cabinet, trash can and wheeled carts.
       - Clean Formica countertops.

- Spot clean walls and doors.

c. Back Door Foyers
   - Sweep and mop, use a cleaner that doesn't require rinsing and is a sanitizer and will not damage the wad. Vacuum rug and clean baseboards.
   - Spot clean walls and doors.

d. Hallways
   - Vacuum all carpeted areas, including wall edges.
   - Spot clean anytime a stain or soiled area needs cleaning.
   - Tile floors sweep and mop, use a cleaner that doesn't require rinsing and is a sanitizer and will not damage the wax.
   - Spot clean walls, doors and partitions that appears to be soiled.

e. Outdoor Waste Receptacles
   - Empty all outdoor waste receptacles and ash trays at the front entrance and two back entrances. Wash if needed with a sanitizing cleaner. Change liners if needed.

f. Conference Room
   - Clean and polish conference room tables.
   - Vacuum all carpeted areas, including wall edges and around the edges of all furniture which is not easily moveable, this includes under desks, tables, chairs etc. All light weight furniture must be moved and vacuumed under. All electrical cords must be picked up and vacuumed under.
   - Spot clean anytime a stain or soiled area needs cleaning.
   - Vacuum chalk dust out of chalk tray. Wash chalkboard only if it has been erased by the Forest Service.

g. Copy Machine and Mail room area
   - Vacuum all carpeted areas, including wall edges and around the edges of all furniture which is not easily moveable, this includes under desks, tables, chairs etc. All light weight furniture must be moved and vacuumed under. All electrical cords must be picked up and vacuumed under.
   - Spot clean anytime a stain or soiled area needs cleaning.
   - Clean and polish table and counter tops.

3. <u>Services Performed Monthly - Bid Item #0003</u>

a. Dusting
   - Dust below a 5 foot level. Dust all horizontal and vertical surfaces including but not limited to furniture, baseboards, wood molding, windowsills, bookcases, ledges, signs, wall hangings, photographs, fire alarm boxes, exhibits, top edge of privacy partitions, excluding desktops and computers.

b. Offices
   - Vacuum all carpeted areas, including wall edges and around the edges of all furniture which is not easily moveable, this includes under desks, tables,

chairs etc. All light weight furniture must be moved and vacuumed under. All electrical cords must be picked up and vacuumed under.
- Spot clean anytime a stain or soiled area needs cleaning.
- Tile floors sweep and mop, use a cleaner that doesn't require rinsing and is a sanitizer and will not damage the wax.

   c. Outside Foyer and Adjacent Areas
- Sweep outside area around all outside doors and adjacent area.
- Pick up any trash laying within 100 feet on the outside of the office building and parking area. This includes all the bushes and trees.

4. <u>Services Performed Annually - - Bid Item #0004</u>

   a. Dusting above 5 feet
- All horizontal and vertical dust catching surfaces shall be kept free of obvious dust, dirt, and cobwebs. Dust furniture in all offices above the 5 foot level, including, but not limited to tops of high bookcases and top edge of privacy partitions.

   b. Windows
- Clean all windows and screens inside and outside of building, with an appropriate glass cleaner. Removing screens on windows that have screens for cleaning.

   c. Blinds
- Dust, clean and/or vacuum all window blinds. Vinyl blinds may require a liquid cleaner and blinds with fabric may require vacuuming. Clean in accordance with manufacturer's recommendations by type of fabric or material.

   d. Chairs
- Vacuum all upholstered chairs.
- Clean all vinyl covered chairs with an appropriate cleaner for vinyl.
- Clean chair legs and/or pedestal bases on all the chairs in the office.
- Wood chairs use an oil, such as lemon oil.

   e. Door and Door Frames
- Clean with appropriate wood/metal cleaner and apply a good penetrating oil to the wood doors.