Contract Duration and the Costs of Market Transactions*

Alexander MacKay
Harvard University†

August 1, 2018

Abstract

The duration of a vertical relationship depends on two types of costs: (i) the transaction costs of re-selecting a supplier and (ii) the cost of being matched to an inefficient supplier when the relationship lasts too long. For commodified goods and services, this tradeoff can be the primary determinant of the duration of supply contracts. I develop a model of optimal contract duration that captures this tradeoff and quantifies the costs of market transactions. These (unobserved) costs are identified even when the exact market mechanism is unknown. I estimate the model using federal supply contracts and find that transaction costs are a significant portion of total buyer costs. I use the structural model to estimate the value of the right to determine duration to the buyer, compared to a standard duration. Finally, a counterfactual analysis illustrates why quantifying transaction costs is important for the accurate analysis of welfare.

*An earlier version of this paper circulated with the title "The Structure of Costs and the Duration of Supply Relationships." This version reflects an expanded data collection effort and a larger dataset. I am especially grateful for the helpful comments and suggestions of Ali Hortaçsu, Brent Hickman, Casey Mulligan, Chad Syverson, and Stephane Bonhomme. This paper has benefited from conversations with Ramon Casadesus-Masanell, Scott Kominers, Maciej Kotowski, Dennis Yao, and my Ph.D. classmates at the University of Chicago, among others. I also thank seminar and conference participants at the University of Chicago, the CEPR-JIE Conference, the Econometric Society World Congress, Northwestern (Kellogg), Harvard Business School, Carnegie Mellon, UCLA (Anderson), Rice, and the International Industrial Organization Conference.

†Harvard University, Harvard Business School. Email: amackay@hbs.edu.

1 Introduction

When buyers select sellers, they select not only who but also how long. The duration of a buyer-seller relationship may have a significant impact on total costs, especially in settings where the costs of "going to the market" to re-select a supplier are meaningful. These transaction costs include the costs of identifying sellers, discovering the relevant prices, and drawing up and concluding contracts (Coase, 1937, 1960). As noted by Coase (1937), a key motivation for a longer contract is to avoid these costs. Choosing a single contract for two years instead of sequential one-year contracts can cut market transaction costs in half.

In the exchange of intermediate goods, market transactions can be especially costly. Compared to the retail sector, spot markets for intermediate goods are typically not well-established, and the buyer often bears a cost for *creating* the market. The buyer may be tasked with soliciting sellers, implementing a price mechanism, and designing a contract. Even for standardized commodities, these costs may be large, as the desired terms vary from one buyer to another. Though creating a market may be costly, there are typically gains from implementing a price mechanism. Auctions, for example, often allow the buyer to select the efficient (i.e., low cost) supplier, reducing total cost to the buyer and increasing overall welfare. The frequency with which the buyer goes to the market impacts supply costs, as the buyer forgoes the benefit of switching to a lower-cost supplier during the duration of a contract.

In this paper, I quantify the costs of creating a market, and I examine the tradeoff faced by a buyer when weighing these *market transaction costs* against the benefit of selecting more efficient suppliers. I focus on a setting in which goods and services are standardized, there is little uncertainty, and relationship-specific investments are negligible. Consider supply contracts for commodified goods and services, such as raw materials, electricity, paper products, accounting services, and some information technology services, including content delivery networks. One might expect these typical inputs to be sold via spot markets, but they are usually governed by fixed-price, fixed-duration contracts.[1] Even in these straightforward economic environments, the duration decision is non-trivial, depending on (1) the magnitude of the transaction costs, (2) the degree of competition, and (3) the stochastic properties of the underlying supply costs.

To illustrate these effects, suppose that a buyer has unit demand in an environment with perfect information and idiosyncratic variation in supply costs over time.[2] When going to the market is costless, a spot market would be optimal, as it allows the buyer to select the low-cost supplier in each period. If the buyer bears a cost for each market transaction, then the buyer can reduce these costs by signing a supplier to a multi-period contract. The resulting multi-period

[1] In their seminal NBER survey, Stigler and Kindahl (1970) found that about half of the commodities in their sample were purchased with fixed-term contracts. A more recent comprehensive survey has not been conducted (and would be welcome). Fixed-price contracts constitute the vast majority of U.S. federal government contracts, and, anecdotally, remain predominant in the private sector.

[2] The lowest-cost supplier may change over time due to capacity constraints, heterogeneity in outside options, innovation, and other factors.

contract has a supply cost premium, which depends on the stochastic properties of the cost process and, importantly, competition. For example, when the supply side is a monopoly, there is no supply cost premium, as the monopolist will be the lowest-cost supplier in each period. In this case, a permanent contract will be efficient regardless of the properties of supply costs and the magnitude of transaction costs. The equilibrium contract therefore reflects both the degree of competition and underlying costs.

Thus, where we observe fixed-duration contracts for simple goods and services,[3] and not solely spot markets, we may infer that market transaction costs are meaningful. Through equilibrium contracts, these costs mediate how prices respond to changes in the economic environment. Empirically, the impact of contract duration on prices and welfare depends on the magnitude of market transaction costs. This prompts the question: How large are these costs? And, relatedly, how valuable are the rights to determine duration? The second question is motivated by the fact that, across many industries, contracts typically have a standard duration (e.g., three years). In light of the analysis of this paper, rigidity of this form may be costly, especially if there is variation in competition or underlying costs across markets or over time.

To address these questions, I develop an empirical model where a buyer selects a seller from an imperfectly competitive market. The buyer chooses the duration of the contract, and the price and the supplier are selected through a market mechanism (e.g., an auction). In equilibrium, the buyer chooses duration to minimize expected buyer costs, which include both the price paid to the supplier and the (amortized) costs to the buyer of implementing a market mechanism. The model, presented in Section 3, reflects a typical real-world contracting problem. Indeed, the tradeoff between reducing transaction costs on the one hand and more frequently re-selecting a supplier on the other is intuitive.[4]

Of course, there are many real-world factors that affect contract duration, including the classic problems relating to ex post misaligned incentives. For this reason, I select an empirical application that allows me to isolate the tradeoff described above. I estimate the model using a unique dataset of 1,046 contracts for building cleaning services for the U.S. federal government, which is described in Section 4. Building cleaning services are standardized, key cost characteristics (such as square footage) are quantifiable, and demand is inelastic, rendering it a nice setting to examine the tradeoff of this paper. The government is required, by regulation, to choose the low-price offer among qualified bidders at the expiration of the previous contract. Thus, the government mechanism approximates the auction mechanism of the model.[5] Consistent with the model, duration is determined ex ante by the local government agency. Further,

[3] Services and products may also have a quality dimension, though for many commodities suppliers compete for the lowest price that meets a minimum quality or particular specifications.

[4] This tradeoff coincides with how several procurement and purchasing personnel described the duration decision to the author. Re-selecting a supplier may include re-negotiation with an incumbent supplier.

[5] In other settings, buyers may have other means to secure a seller, such as a direct negotiation with the incumbent supplier. This would complicate the model.

the supply-side conditions are relatively stable, and relationship-specific investments are small.[6] This suggests that abstracting away from other contracting concerns may be reasonable, and it allows me to focus on identifying the direct (Coasian) costs of going to the market.

One of the challenges in evaluating the impact of transaction costs is that they are often unobserved or difficult to quantify. The empirical strategy of this paper, based on the observation that the equilibrium duration reflects the underlying cost structure, provides for the identification of latent transaction costs with minimal structure on the market mechanism.[7] In the application, I estimate that the transaction costs of creating markets are large and economically meaningful. The median estimate is $10,300, and, in aggregate, they comprise 10.8 percent of total buyer costs. I find that these costs are correlated with other observables in ways that align with our intuition. For example, the costs are positively correlated with the number of words in the contract, which is a proxy for the complexity of the service. These empirical results are presented in Section 5.

After quantifying the magnitude of market transaction costs, I consider the impact of the duration margin in Section 6 using the structural model. By choosing duration, the buyer can optimize contract by contract, as opposed to a regime in which contract terms are fixed. To calculate the value to the buyer, I consider an alternative policy where all contracts are issued with a standard duration. Standard terms could be costly: issuing one-year contracts for the data in my sample would increase total costs by 36 percent, as the frequency of market transactions would rise sharply. Of the full-year durations, the four-year standard term has the lowest impact, increasing total costs by 1.4 percent. We might expect standardization to reduce transaction costs, e.g., through reduced effort from the buyer to determine the optimal duration. I find that relatively modest declines in transaction cost (9.3 percent) would offset the costs of moving away from optimal contract-specific durations. Thus, a poorly chosen standard could substantially increase costs, but an informed standard may be cost effective with moderate reductions in transaction costs.

As a second counterfactual, I consider the impact of endogenous duration and transaction costs on the estimation of welfare effects. Relative to a structural model that takes duration as exogenous, the estimated model allows buyers to adjust on an additional margin (duration) in response to changes to the state, improving buyer surplus relative to environment in which they are passive. I then compare the estimated effects from the structural model to a difference-in-differences strategy that measures the change in prices after a policy that affects transaction costs. The price changes capture only 63 percent of the change in total costs, implying that a reduced-form estimate would substantially underestimate the welfare effects compared to a

[6]Ex post incentive problems, which are a large focus of the contract literature, are not a first-order concern. Performance is observable, contracts are rarely canceled, and the suppliers in this market are generally well-established firms.

[7]Thus, this result may be useful for analyzing contracts where the underlying selection process may be obscure, such as contract negotiations that occur privately.

model that explicitly accounted for transaction costs. As we expect that many policy changes would affect transaction costs, this analysis suggests that it would be important to take such costs into account.

As a general contribution, Section 3 provides results for nonparametric identification when the market mechanism is an auction, while allowing for unobserved heterogeneity. I show how the entry cost shifters can be used to identify the joint distribution of costs without information on other bids or a reserve price, even in the presence of a third endogenous outcome (duration). Further, when entry is exogenous, I show that the distribution of private costs and unobserved heterogeneity is identified through variation in the number of bidders alone.[8] 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. The identification results are complementary to the work of Krasnokutskaya (2011) and Aradillas-López et al. (2013), among others, and make use of data that is more broadly available. Thus, in settings that may be motivated by the independent private values assumption, it is also possible to estimate a conditional independent private values model with unobserved heterogeneity. Researchers who wish to apply auction techniques to transaction price data may use this result to test for and quantify the presence of unobserved heterogeneity.

Before introducing the more general model, I explore the tradeoff between transaction costs and supply costs using a simple example. In Section 2, I present a two-period model where the buyer picks either two one-period contracts or a single two-period contract. The simple model provides some intuition about the tradeoff, the relation to underlying supply costs, and the degree of competition. For example, the model shows that higher autocorrelation in supply costs leads to longer contracts. Additionally, the optimal duration is U-shaped in the number of suppliers competing for the contract. With low levels of competition, long-term contracts are optimal, as the benefit of re-selecting a supplier is small. This benefit increases as competition increases, resulting in short-term contracts at moderate levels of competition. When competition is intense, the buyer can secure a low-enough price for an extended period that long-term contracts are once again optimal.

Related Literature

The essential connection between market transaction costs and contract duration was identified at least as early as Coase (1937). Economists studying the effect of transaction costs on vertical relationships have primarily pursued the testable implications of these costs, rather than their direct estimation.[9] Thus, a central contribution of this paper is a structural empirical model that allows for the direct estimation of these costs.[10] Likewise, the empirical literature on contract

[8] Previous approaches relied on observing either multiple bids per auction or a reservation price.

[9] 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 uncertainty.

[10] Conceptually related is recent work by Atalay et al. (2017), who construct a measurement of external transaction costs by examining input flows between integrated and non-integrated firms across sectors.

duration has also focused on testable implications, so providing an empirical model of contract duration is also a contribution of the paper.

The theoretical and empirical analysis of the costs of market transactions has typically been cast in light of the decision to vertically integrate (see Lafontaine and Slade (2007) for a summary). As integration often induces additional benefits from the alignment ex post incentives, disentangling the direct cost of transacting on the market from, for example, the cost of hold-up can be challenging. By focusing on the question of contract duration, I isolate the direct costs of transacting on the market while abstracting away from other features of vertical integration. Of course, an increase in these costs should also increase the propensity to vertically integrate, as noted by Coase (1960).

To the best of the author's knowledge, this is the first paper to focus on a general ex ante cost of longer contracts, which arises from an inefficient supplier match over the duration. The previous literature on contract duration has focused on ex post coordination problems, primarily through costly renegotiation (Masten and Crocker, 1985) and relationship-specific investments (Joskow, 1987). To focus on the novel aspect of contracting, I select an empirical setting that allows me to abstract away from ex post incentive problems, including risk sharing, principal-agent relationships, the holdup problem, and incomplete contracting. Recent empirical work on these features (e.g., Decarolis (2014) and Bajari et al. (2014)) focus on one-time projects and therefore do not model repeated demand.[11] 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. I am also able to test for and abstract away from incumbency advantage, which is often a concern in settings with repeated contracts (see, e.g., Greenstein (1993)). My work is complementary to models with these features.

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. In markets with vertical contracts, transaction costs may be a sizable portion of total costs and should be accounted for in addition to any price effects.[12]

There is a parallel empirical literature on switching costs in consumer markets (e.g., Dubé et al. (2010); Handel (2013); Honka (2014); Illanes (2017)), 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 (Klemperer, 1995).[13] When buyers and sellers agree on future prices, as in this

[11]The theoretical literature linking these features to contract duration includes Holmstrom (1983), Guriev and Kvasov (2005), and Rey and Salanie (1990).

[12]Carlton and Keating (2015) emphasize the role of transaction costs in welfare analysis when the affected variable is not simply the price level, through the effect on a firm's ability to implement nonlinear pricing.

[13]For recent papers on this subject, see Cabral (2016) and Rhodes (2014).

paper, these effects are competed away. Further, switching costs in previous studies can be inferred from posted prices,[14] whereas contract prices for intermediate goods are idiosyncratic to the buyer-seller match. Finally, the switching costs literature tends to take supply costs as fixed (e.g., Beggs and Klemperer (1992)), whereas variation in supply costs is a key factor in the decision to switch suppliers in my setting.

My contribution to the auction identification literature is closely related to Krasnokutskaya (2011), who solves the problem of disentangling private costs from auction-specific heterogeneity by relying on two bids per auction. Concurrent work by Quint (2015) exploits variation in the number of bidders in a model with additively separable unobserved heterogeneity. That identification strategy does not translate to the multiplicative structure examined here. 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.

The contract duration decision is theoretically linked to a simultaneous bundling problem, where the contract bundles demand over time. In the bundling literature, Zhou (2017) and Palfrey (1983) provide the most closely related analogs and generate complementary insights.[15] In this paper, I demonstrate that the smaller variance induced by bundling reduces total surplus when there are no transaction costs.[16]

2 A Simple Model

In this section, I develop a simple model to illustrate how competition, supply costs, and transaction costs interact to determine the duration of contracts. I focus on the buyer-optimal contract, as it reflects the theoretical outcome for the empirical setting of this paper. The qualitative predictions developed for the buyer-optimal contract will also pertain to the efficient contract.[17] I defer a discussion of efficiency to Section 7.

Suppose that a risk-neutral buyer is seeking the supply of a good for two periods. There are N symmetric suppliers in the market, and the set of suppliers stays the same across both

[14] See, for example, Dubé et al. (2010) for orange juice and margarine or Elzinga and Mills (1998) for wholesale cigarettes. The wholesale market in the analysis of Elzinga and Mills (1998) mirrors a consumer market in that pricing, though nonlinear, is uniformly applied.

[15] Compared to their analysis, I allow for intermediate degrees of bundling and introduce transaction costs.

[16] Salinger (1995) and Bakos and Brynjolfsson (1999) note that bundling affects prices by reducing the variance of average valuations. Cantillon and Pesendorfer (2006) share this insight in their analysis of combination bidding for multi-unit auctions.

[17] Intuitively, the predictions depend on the covariance structure and the properties of order statistics. The expected cost and the expected price in this simple setting are the first-order and second-order statistics, and thus display similar properties when the number of draws is greater than three.

periods. The buyer can issue either a single two-period contract or sequential single-period contracts. The transaction cost of going to the market is δ, which captures implementation and contracting costs, as well as a complementarity of supplying for two periods. The game proceeds in three steps. First, the buyer selects a duration of either one or two periods. Second, suppliers realize cost draws for both periods. Third, suppliers participate in an auction for each contract, i.e., an efficient mechanism.[18]

The per-period cost to each supplier is the random variable c. When the buyer issues single-period contracts, the per-period cost of the winning supplier is $c_{1:N}$, which is the minimum of N draws of c. When the buyer issues a two-period contract, the average per-period costs for each supplier is average of two draws, $\tilde{c} = \frac{1}{2}(c^{(1)} + c^{(2)})$, and the cost to the winning supplier is $\tilde{c}_{1:N}$.

Remark 1 As long as the per-period costs c are not perfectly correlated across periods, $\tilde{c} \neq c$ and $Var(\tilde{c}) < Var(c)$.

Thus, the buyer changes the effective cost structure faced by suppliers when changing the contract duration. When the distribution of supply costs is stable over time, this serves to reduce the variance of cost draws. The cost of a longer contract is that the low-cost supplier may not be selected in each period. In the absence of transaction costs, short-term contracts would be optimal.

If we further assume that the buyer is risk-neutral, symmetry in this setting generates the standard auction result that the expected winning bid is equal to the second-order statistic from the cost draws. Thus, the buyer-optimal contract solves

$$\min\{\underbrace{2E[c_{2:N}] + 2\delta}_{\textit{short-term}}, \underbrace{2E[\tilde{c}_{2:N}] + \delta}_{\textit{long-term}}\}$$

The buyer will pick the long-term contract if the increase in expected supply costs is less than the reduction in (amortized) transaction costs:

$$E[\tilde{c}_{2:N}] - E[c_{2:N}] < \frac{\delta}{2}. \tag{1}$$

This simple decision rule generates a number of comparative statics, which I present to build intuition about the model. From equation (1), it follows that:

Remark 2 Higher marginal costs lead to shorter contracts, and higher transaction costs lead to longer contracts.

[18] 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.

Further, we can employ the properties of order statistics to generate comparative statics related to competition and the stochastic properties of costs. For simplicity of exposition, these are presented here with a brief discussion, and a numerical illustration follows. For further details on each, see the Appendix.

Remark 3 The optimal duration is increasing with autocorrelation in supply costs.

This prediction is intuitive. As the autocorrelation in marginal costs increases, there is less of a benefit from switching suppliers, and longer-term contracts are preferred.

Remark 4 The optimal duration is decreasing in the variance of costs across suppliers, provided there is sufficient competition ($N > 3$).

Remark 4' When costs are bounded from below, the optimal duration is U-shaped in the variance in costs, provided there is sufficient competition ($N > 3$).

From a starting point of zero variance across suppliers, increasing the variance of marginal costs leads to shorter contracts, as there is more to gain from selecting the low-cost supplier in each period. This holds for the buyer-optimal contract as long as there are more than three suppliers, in which case the expected second-order statistic falls below the median. When costs are bounded from below, eventually both $E[\tilde{c}_{2:N}]$ and $E[c_{2:N}]$ approach zero, and the cost of a longer duration falls with respect to transaction costs. After a certain threshold, contract duration increases.

Remark 5 When costs are bounded from below, the number of suppliers has an inverse U-shape effect on the marginal cost of longer contracts. Therefore, duration may be decreasing, increasing, or U-shaped with N.

As this simple model illustrates, an increase in the intensity of competition, i.e., the number of suppliers, has an ambiguous effect on equilibrium contract duration, depending on underlying market conditions and parameters. For low levels of competition, the benefit of switching suppliers is low, and long-term contracts are preferred. At moderate levels of competition, there is an increased benefit of switching among suppliers more frequently. When competition is intense, the expected costs of both long-term and short-term contracts approach the lower bound of costs, and therefore long-term contracts, which minimize transaction costs, are optimal.

2.1 A Numerical Example

To illustrate the above predictions, I present a simple case in which per-period costs are drawn from a beta distribution with parameters $(\alpha, \beta) = (0.5, 0.5)$. Recall that the beta distribution has support $[0, 1]$. With the parameters $(\alpha, \beta) = (1, 1)$ it is equivalent to a uniform distribution, and as α and β approach zero it approaches a Bernoulli distribution.

Figure 1: Competition, Costs, and Contract Duration: A Numerical Example

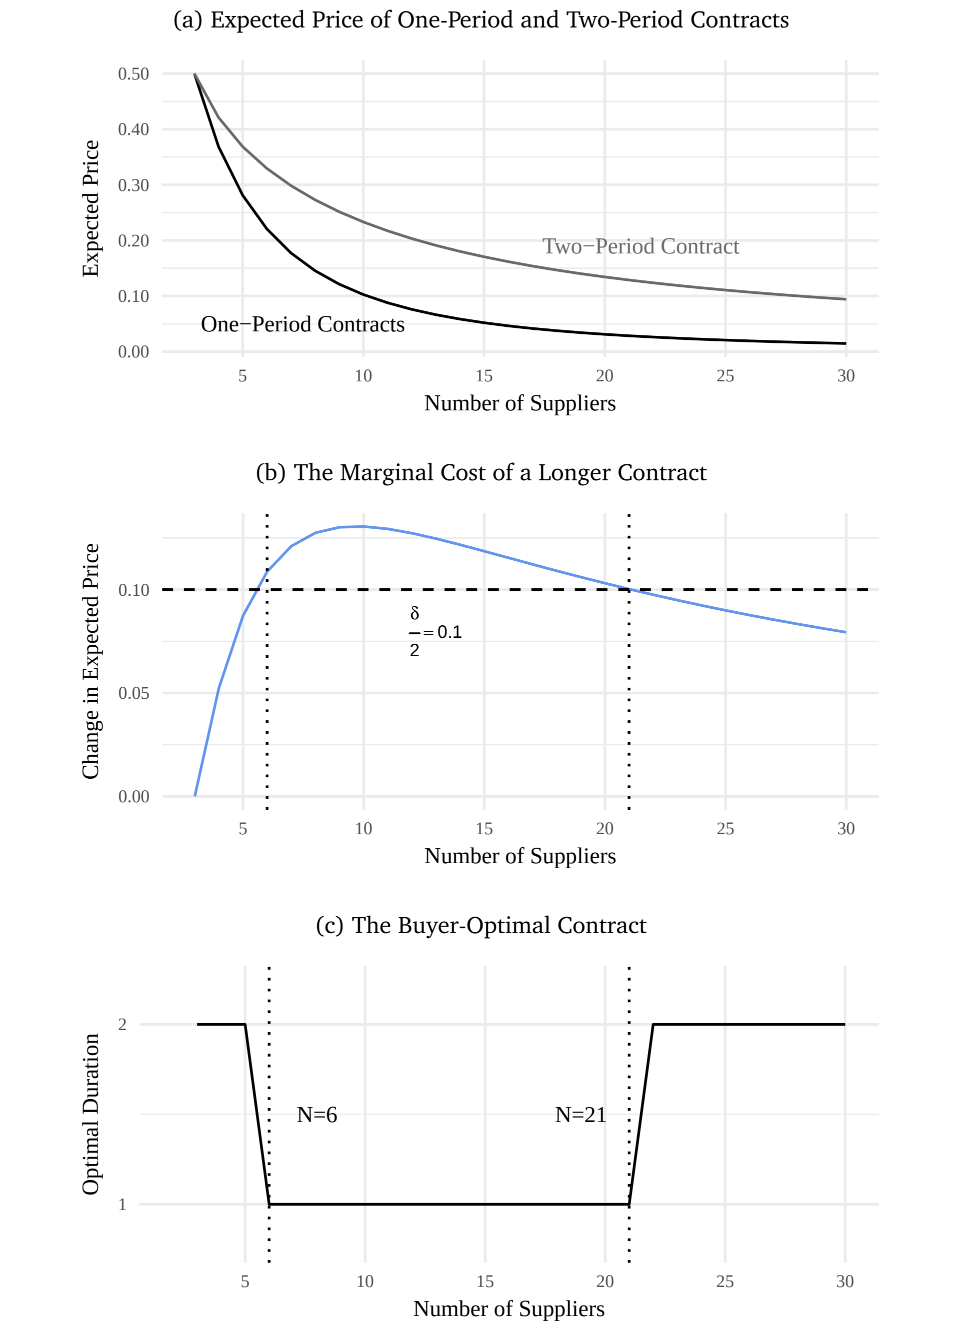

Notes: Panel (a) plots the expected per-period costs for separate one-period contracts and a bundled two-period contract, as a function of the number of bids. The blue line in panel (b) is the difference between the two, which is the expected price increase to the buyer. The dashed line in panel (b) reflects a transaction cost of 0.2 amortized over two periods, which is the amount saved by issuing a two-period bundled contract. For values of N where the blue line is above the dashed line ($N \in \{6, ..., 21\}$), short-term contracts are optimal, as the increase in supply costs from the long-term contract is greater than the savings in transaction costs. Panel (c) plots the buyer-optimal contract duration.

Figure 1 illustrates how the marginal cost of a longer contract varies with the competitive conditions in the marketplace. Panel (a) plots the expected supply price for one-period contracts and two-period contract. For $N = 3$, the expected prices are the same, and for $N > 3$ the single-period contracts always have a lower expected price. The blue line in panel (b) plots the difference between these two lines. The dashed line indicates a transaction cost of 0.20, which is amortized by two periods. When the blue line falls above this dashed line, the increase in the expected supply price exceeds the savings in transaction costs, and one-period contracts are optimal. Panel (c) plots the U-shaped buyer-optimal duration as a function of N. Short-term contracts are optimal for moderate level of competition; in this case, when $N \in \{6, ..., 21\}$.

For the sake of brevity, I omit an extended exposition of the results from the numerical model, which can be used to illustrate the predictions outlined earlier. For an illustration of the impact of increased variance, see Figure 6 in the Appendix.

3 An Empirical Model of Contract Duration

In this section, I first introduce a more general purchasing problem facing a buyer. The buyer can affect the outcome of the transaction by changing the duration of the contract, but the buyer takes transaction costs, contract characteristics, and supply conditions as given. I show how the problem simplifies when the distribution of costs 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. The model is the basis for the empirical approach of Section 5.

3.1 The Buyer's Problem

I now generalize the simple model of the previous section into an infinite-horizon problem. Suppose the buyer has inelastic demand for a good. The buyer selects a single seller of the good and commits to buy from that seller for T periods. The seller is selected via a market mechanism, which costs the buyer δ to implement and also determines price.

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 $\delta > 0$. 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.[19]

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

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

tral, and that the buyer expects contract characteristics and entry cost shifters to remain the same in future periods.[20] This allows us to suppress (X, M) and write $\overline{P}(T)$ in the exposition below.

The value function for the buyer in period τ who has not yet chosen a seller can be expressed as

$$V(\tau) = \min_{T} \delta + \sum_{k=1}^{T} \beta^{k-1}\overline{P}(T) + \beta^{T} V(\tau + T). \tag{2}$$

After incurring the transaction cost to determine the seller and the price, the buyer pays $\overline{P}(T)$ in each period and returns to the decision problem in period $\tau + T$. The buyer discounts future periods at rate β.

For an optimal T, it must be that, for any other duration S:

$$\delta + \sum_{l=1}^{T} \beta^{T-1}\overline{P}(T) + \beta^{T} V(\tau + T) \leq \delta + \sum_{l=1}^{S} \beta^{S-1}\overline{P}(S) + \beta^{S} V(\tau + S). \tag{3}$$

To derive a more intuitive comparison, we expand the component of the continuation value to period $\tau + T \cdot S$. As the buyer expects market conditions to persist, if T is optimal in period τ, the buyer expects T to be optimal at the expiration of a contract in a future period, e.g., in period $\tau + T$. Plugging in a sequence of contracts of duration T and S, we obtain

$$\begin{aligned} &\sum_{l=1}^{S} \beta^{T(l-1)} \left(\delta + \sum_{k=1}^{T} \beta^{k-1}\overline{P}(T) \right) + \beta^{T \cdot S} V(\tau + T \cdot S) \\ \leq &\sum_{l=1}^{T} \beta^{S(l-1)} \left(\delta + \sum_{k=1}^{S} \beta^{k-1}\overline{P}(S) \right) + \beta^{T \cdot S} V(\tau + T \cdot S). \end{aligned}$$

Rearranging,[21] we obtain the simpler expression

$$\overline{P}(T) - \overline{P}(S) \leq \frac{\delta}{\sum_{k=1}^{S} \beta^{k-1}} - \frac{\delta}{\sum_{k=1}^{T} \beta^{k-1}}. \tag{4}$$

This expression corresponds to equation (1) in the simple model. At the optimal contract, potential savings in the per period price by selecting a different (shorter) duration are less than

[20] Thus, we assume that the buyer's expectation is that future periods look like today. We can therefore omit the expectation operator in the buyer's problem.

[21] We can cancel out the continuation value component from the inequality, and, with the substitution $\delta = \sum_{k=1}^{T} \beta^{k-1} \frac{\delta}{\sum_{k=1}^{T} \beta^{k-1}}$, we can factor out the common aggregate discount factor $\sum_{k=1}^{TS} \beta^{k-1} = \sum_{l=1}^{S} \beta^{T(l-1)} \sum_{k=1}^{T} \beta^{k-1} = \sum_{k=1}^{TS} \beta^{k-1} = \sum_{l=1}^{T} \beta^{S(l-1)} \sum_{k=1}^{S} \beta^{k-1}$. This results in

$$\left(\sum_{k=1}^{TS} \beta^{k-1} \right) \left(\frac{\delta}{\sum_{k=1}^{T} \beta^{k-1}} + \overline{P}(T) \right) \leq \left(\sum_{k=1}^{TS} \beta^{k-1} \right) \left(\frac{\delta}{\sum_{k=1}^{S} \beta^{k-1}} + \overline{P}(S) \right),$$

and equation (4) follows.

increased transaction costs from using the market mechanism more frequently.

Given realizations for contract and market characteristics x and m, the optimal duration, t, is therefore given by

$$t = \arg\min_{T} \overline{P}(T, x, m) + \frac{\delta}{\sum_{k=1}^{T} \beta^{k-1}} \tag{5}$$

Intuitively, this expression shows the buyer's objective is to minimize the sum of the per-period supply price and amortized transactions cost.

Comparative Statics

When contract duration can be chosen in arbitrary increments, the length of periods approaches zero and we obtain a continuous-time problem, with $\beta \to 1$. When $\beta = 1$, the objective function may be written

$$\min_{T} \overline{P}(T, x, m) + \frac{\delta}{T} \tag{6}$$

resulting in the first-order condition

$$\frac{d\overline{P}(T, x, m)}{dT}|_{T=t} = \frac{\delta}{t^2}. \tag{7}$$

As $\delta > 0$, it must be that $\frac{d\overline{P}(T,x,m)}{dT}|_{T=t} > 0$ for any interior solution t. Thus, when contracts are finite, the expected supply price is increasing with duration (at the equilibrium). For the rest of this section, I assume that such interior solutions exist. As illustrated in Section 2, $\frac{d\overline{P}(T,x,m)}{dT}$ will tend to be positive when the market is sufficiently competitive, as an increase in T causes suppliers to average cost draws across multiple periods. This shrinks the variance of the cost distribution of the duration of the contract, which increases the expected minimum cost.[22] This increase reflects the fact that the buyer will not be matched to the low-cost supplier in each period.

As a check of the model, we have the intuitive result that higher transaction costs lead to longer contracts.

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

Proof. See Appendix C. □

Additionally, the 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 lower entry costs, for example. Whether or not the equilibrium contract is increasing with respect to these characteristics depends only on the cross-partial of the expected price function, which can be estimated without modeling the buyer's decision or observing transaction costs.

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

Proposition 2. *The optimal duration t is increasing in M if $\frac{\partial^2 \overline{P}(T,X,M)}{\partial T \partial M}$ is negative and decreasing if the cross-partial is positive. Likewise for X.*

Proof. See Appendix C. □

3.2 Supplier Participation and Equilibrium

In this section, I develop a three-stage empirical model, where the first stage reflects the buyer's problem, the second stage is the participation decision of suppliers, and the third stage is the supplier selection mechanism. 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.

The equilibrium is characterized by the buyer choosing duration to minimize expected buyer costs, potential participants entering if expected profits exceed entry costs, and the supplier selection mechanism generating a proportional price B according to the equilibrium strategies of the suppliers. The model is summarized in Figure 2.

1st Stage: Duration Decision The buyer 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 buyer. Thus, we make the usual assumption that the log price can be decomposed into additive (logged) components. I assume that, for the relevant window, $\beta = 1$. The buyer's objective function is:

$$\begin{aligned}
& \min_T \overline{P}(T, x, m) + \frac{\delta}{T} && (8) \\
= \; & \min_T E[B \cdot U \cdot h(X) | T, x, m] + \frac{\delta}{T} && (9) \\
= \; & \min_T \left(\sum_{n=1}^{\mathbb{N}} E[B \cdot U | n, T, x, m] \cdot \Pr(N = n | T, x, m) \right) h(x) + \frac{\delta}{T} && (10)
\end{aligned}$$

2nd Stage: Participation Potential entrants observe (T, X, M), common entry costs $k(M)$, and a multiplicative 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

$$E[\pi_n \cdot U | n, t] \cdot h(x) - k(m) \cdot \varepsilon > 0 \iff N \geq n. \qquad (11)$$

Figure 2: Summary of Model

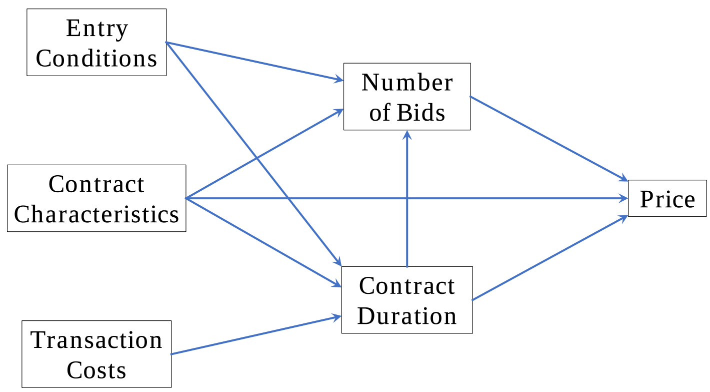

Notes: The figure summarizes the causal assumptions embedded in the empirical model. The three sets of variables on the left: entry conditions, contract characteristics, and transaction costs, are taken as given. Price, number of bids, and contract duration are jointly determined in the model. Arrows indicate the direction of causality.

3rd Stage: Supplier Selection After the participation decision, each participant i realizes a private (proportional) cost C_i, along with the common cost U, and engages in the supplier selection mechanism.[23] Total costs to the supplier are given by $C_i \cdot U \cdot h(x)$, resulting in the stochastic price $B \cdot U \cdot h(x)$. The proportional offer B is an equilibrium outcome when suppliers are risk neutral. 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 of the 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.[24] Furthermore, the identification results allow for suppliers that observe U prior to making the participation 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 buyer and the econometrician, but their

[23]The model is identified if U is observed prior to entry. The results generalize to other information structures, including affiliated values.

[24]For example, the model could be applied to a challenger and incumbent game with alternating offers and asymmetry between the supplier types.

distributions are common knowledge. Assume

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

2. *Independence of Common Shocks and Observables:* $(\varepsilon, U) \perp\!\!\!\perp (T,X,M)$.

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

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. $E[U|N,T,X,M]$.
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 D. □

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.

3.3.2 Identification of Transaction Costs

Once the key components of costs are identified, transaction costs may be obtained via revealed preference. Recall the buyer's objective function:

$$\begin{aligned} &\min_T \left(\sum_{n=1}^{\mathbb{N}} E[B \cdot U|N=n,T,x,m] \cdot \Pr(N=n|T,x,m) \right) h(X) + \frac{\delta}{T} \\ &= \min_T \left(\sum_{n=1}^{\mathbb{N}} E[B|N=n,T] \cdot E[U|N=n,T,x,m] \cdot \Pr(N=n|T,x,m) \right) h(X) + \frac{\delta}{T} \end{aligned} \tag{12}$$

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 (12). 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:

$$
\begin{aligned}
(t'-t)\delta \quad \leq \quad & t \cdot t' \Bigg(\sum_{n=1}^{N} E[B|n,t'] \cdot E[U|n,t',x,m] \cdot \Pr(N=n|t',x,m) - \\
& \sum_{n=1}^{N} E[B|n,t] \cdot E[U|n,t,x,m] \cdot \Pr(N=n|t,x,m) \Bigg) h(x)
\end{aligned} \tag{13}
$$

for all $t' \in \mathbb{T}$. 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. □

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 (13) vary continuously with X, the cumulative distribution function of δ is identified. □

3.3.3 Identification of Seller Surplus

With the above assumptions, relative profits and transaction costs are identified. To conduct an efficiency analysis, it is necessary to also obtain seller surplus. One approach would be to employ supplementary data on profits for one (n, t) pair. This would identify the expected profit function and, therefore, the expected supply cost $E[C|N, T]$. When no such 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, along with partial identification of the joint distribution of outcomes. We proceed with the auction model, which 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. *Symmetry*: $C_i \sim F_i$, with $F_i = F$ for all i.
3. F is continuous with positive support. $U \sim G$, where G has positive support.

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

Symmetry is a typical assumption in auction models of unobserved heterogeneity.[25] To relax symmetry, one could consider alternative restrictions to pin down costs and the joint distribution of outcomes.

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

Proof. See Appendix D. □

Briefly, variation in N, combined with identification of relative profits, allows for identification of seller surplus in the auction model. As I show in the Appendix, we can further build on this identification result to pin down properties of the private cost distribution.

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 D. □

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 have additional power 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 obtained after* F *is identified.*

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. □

3.4 A Special Case: The Empirical Setting of Independent Private Values

In this section, I have outlined the nonparametric identification results for a transaction problem when the buyer chooses the duration of the contract, and the data include the transaction price, the duration of the contract, a measure of competition, and contract and market characteristics. The model provides an empirical approach to modeling prices, estimating transaction costs, and constructing some counterfactuals when the exact supplier selection mechanism is unknown. Additional restrictions on the selection mechanism allow for an efficiency analysis and (partial) identification of the private cost distribution.

It is worth considering another approach, which provides identification for the auction model in the absence of a valid instrument M and with no selection on unobservables. The result is independent of the presence of transaction costs or the duration-setting problem.

[25] See, for example, Aradillas-López et al. (2013) and Krasnokutskaya (2011).

Proposition 8. *First-price, symmetric auctions with unobserved heterogeneity and conditionally independent private values are identified when only the winning bid and the number of bidders is observed. 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 D. □

This third identification result may prove practical, particularly for researchers who are interested in employing auction concepts to study phenomena but lack the detailed data required for richer models. When the data have only transaction prices and a measure of competition (e.g., the number of bids), estimation is often motivated by the independent private values (IPV) assumption. The gist of this identification result is that in any setting where estimation is motivated by IPV, one could also estimate a conditional independent private values model with unobserved heterogeneity. One might expect that unobserved heterogeneity is present, and this provides a theoretical background to test for its importance. In Appendix F, I detail a computational approach that greatly speeds up the maximum likelihood estimation of these models.

4 Empirical Application: Data and Reduced-Form Analysis

4.1 Data

To estimate the cost of market transactions, I construct a dataset of 1,046 competitive contracts for building cleaning services for the United States federal government. By regulation, much of federal procurement is competitive, where the buyer is forced to, in good faith, solicit bids and choose the best offer. As the buyer is compelled to go to the market at the expiration of a contract, this provides a relatively clean case study to analyze the duration decision and estimate the costs of market transactions. Furthermore, regulation mitigates concerns about the impact of relationship-specific investments and the scope for collusion.

Another empirical challenge is that procured goods and services may have heterogeneity that is multi-dimensional and hard to quantify. Thus, I select commodity-like goods and services where cost factors can be readily quantified.[26] Indeed, products of this sort are numerous in procurement and make up a significant portion of all transactions. Of all competitive contracts for commodity-like products, building cleaning services were chosen because they are numerous, cost factors are easily quantified, and there is a lot of variation in contract duration. Finally, demand is inelastic, as there are no significant substitutes during this period. The market for such services is sizable; the federal government spent $1.2 billion annually on such services during the sample period.

[26] A counter-example of the ideal setting for this sort of analysis might be a customized, large-scale computer software system for an agency.

Table 1: Construction of Sample

	Criterion	Observations	Portion
(1)	FedBizOpps Solicitation IDs	7,984	
(2)	FPDS Solicitation IDs	11,210	
	Matched (1) and (2)	4,119	
(3)	In United States	3,818	0.93
(4)	Competitive Procurement	3,584	0.94
(5)	Non-Zero FPDS Value	4,064	0.99
(6)	Square Footage Indicators	1,654	0.40
	Intersection of (3)-(6)	1,427	0.35
(7)	US, Excluding Territories	1,409	0.99
(8)	Regular Cleaning Service	1,405	0.98
(9)	Measurable Square Footage	1,301	0.91
(10)	No Economic Disadvantage Preference	1,289	0.90
(11)	Single Auction, More Than 1 Bid	1,339	0.94
(12)	Annual Price Less Than $1,000,000	1,338	0.94
	Estimation Sample Intersection of (7)-(12)	1,046	0.73

Notes: The table describes the construction of the estimation sample from two data sources for facility cleaning contracts for the U.S. federal government. The relevant range is from October 1, 2003 through May 1, 2017 for the Federal Procurement Data System and though February 3, 2017 for FedBizOpps. After cleaning identification variables, 4,119 of the solicitations were matched. Of these, 1,046 met the criteria needed for analysis, including the availability of square footage data, which is a key cost indicator, non-zero value, and receiving more than one bid from the solicitation.

Key outcomes of the model developed earlier are price, duration, and competition (entry). To the best of the author's knowledge, this is the first large dataset to combine observations on these three outcomes. To construct this dataset, I combined detailed location, price, and vendor information maintained in the Federal Procurement Data System (FPDS)[27] with contract-specific documents downloaded from the Federal Business Opportunities (FedBizOpps) website. By law, the FPDS keeps public records of all contracts for the U.S. federal government.[28] The FedBizOpps website is the most common posting location for competitive contracts, which must be posted publicly. From October 2003 through May 2017, I identified 11,210 unique solicitations in the FPDS data and 7,984 unique solicitations in the FedBizOpps data. I was able to merge 4,119 of these contracts.

From the solicitations found in both systems, I selected competitive, non-zero value contracts in the United States that had documents with relevant cost information (i.e., square footage).[29] I obtained the relevant contract documents (request for proposal, cleaning fre-

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

[28]Its data have been used in recent empirical work by Kang and Miller (2017), Bhattacharya (2018), and Decarolis et al. (2018), among others.

[29]The candidate solicitations were identified with a computational text analysis of documents from all matched

quency charts, maps, etc.), and constructed detailed contract information directly from the documents. The resulting 1,427 contracts were further processed by hand to construct key variables, including the square footage of the site to be cleaned, the frequency of service, and the facility type. Contracts that were restricted to economically disadvantaged businesses were removed from the sample. After identifying contracts for regular cleaning service, I restricted the sample to contracts that received more than one bid and had an annual price of less than $1 million. Table 1 summarizes the construction of the dataset.

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. Local labor market conditions include county-level unemployment from the Local Area Unemployment Statistics and the number of NAICS-code level establishments in the same 3-digit ZIP code from the County Business Patterns data.

4.1.1 Data Cleaning

Though the FPDS data are broadly appealing for research, there is a great deal of measurement error in the data, likely due to user (input) error. As most contracts have multiple entries and multiple indicators of duration and value within each entry, different assumptions about data quality could lead to widely different measures of price. As I obtained high-quality measures of price and duration from a second data source, FedBizOpps, I was able to cross-validate the data and construct preferred measures from the FPDS.

In supplemental work, I detail the steps to cross-check the data and different candidate measures for price and duration. These comparisons result in the following recommendations:

Duration *The maximum observed date in the contract, minus the start date in the first entry within a contract.*

Price *The price is the value of obligated dollars if it is the same (or within 10 percent) in consecutive years. If this is not observable, use the maximum value of the three (summed) measures of dollar amounts for the total value of the contract. Divide this by the duration measure above to obtain the price.*

Any missing values of price or duration in the FedBizOpps data are imputed with the above values constructed from FPDS. Researchers interested working with the FPDS data may contact the author for a short paper that details the measurement error in the data and the accuracy of variables constructed under alternative assumptions.

Table 2: Count of Contracts by Location Type

Category	Sub-Category	Count
Office (424)		
	Office	221
	Recruiting Office	203
Field Office (270)		
	Ranger District Office	171
	Field Office	46
	Ranger Station	43
	Work Center	7
	Reserve Fleet	3
Research (111)		
	Weather Station	43
	Laboratory	28
	Research Center	28
	Plant Materials Center	12
Medical (61)		
	Clinic	36
	Medical Center	25
Services (59)		
	Service Center	38
	Vet Center	21
Visitors (41)		
	Recreation Area	18
	Cemetery	9
	Visitor Center	7
	Restroom	4
	Museum	3
Airport (30)		
	Airport	30
Technical (19)		
	Power Plant	14
	Surveillance Center	4
	Data Center	1
Accommodations (18)		
	Housing	14
	Dormitories	4
Industrial (13)		
	Equipment Center	6
	Warehouse	6
	Gym	2
Total		1,046

Notes: The table lists the count of contracts in the estimation sample by facility type. Types were hand-coded after reading the contract documents.

4.1.2 Institutional Details

Competitive contracts are contracts that are posted publicly and allow open competition from registered vendors.[30] 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.[31]

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 buyer 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.

Contracts include specifications for the tasks to be done and their frequencies. For building cleaning, tasks include mopping, vacuuming carpets, picking up debris, dusting, and emptying trash cans. For an example list of specifications, see Section H in the Appendix.

The majority of the contracts (694) 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 2 lists the frequency of each type of site, which are grouped into ten major categories.[32]

4.1.3 Summary Statistics

Summary statistics for the contracts are displayed in Table 3. Contracts vary in price, duration, and competition. As shown later in this section, much of the variation in price can be captured by the square footage of the building and the cleaning frequency.[33] The median contract is relatively inexpensive, as is typical for many commodity-like goods and support services. For

contracts.

[30]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.

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

[32]For a breakdown of contracts by the issuing department or agency, see Appendix G.1.

[33]Cleaning frequency is encoded as the maximum required weekly frequency in the contract specifications.

Table 3: Summary Statistics

	Mean	Min	p25	Median	p75	Max
Contract Value ($1000s)	190.2	2.91	28.5	50.5	102.0	4882.7
Price (Annual, $1000s)	43.9	1.11	7.3	13.2	26.7	976.5
Duration (Years)	4.2	0.42	3.0	5.0	5.0	6.5
Square Feet (1000s)	25.7	0.14	3.7	7.0	14.5	2031.8
Price per Square Foot	2.9	0.16	1.3	2.0	3.1	33.0
Number of Bids	6.5	2.00	4.0	5.0	8.0	40.0
Weekly Frequency	3.5	0.11	2.0	3.0	5.0	7.0
Winner: Num. Employees	61.5	1.00	3.0	14.0	75.0	650.0
Observations	1046					

Notes: The table displays summary statistics for key variables in the contract data. Included are outcomes (price, duration, and number of bids), as well as cost characteristics such as the number of square feet and the frequency of cleaning. The last variable is the size of the winning firm, in terms of number of employees.

the sample, which removes contracts greater than $1 million per year, the mean contract is for $44,000 annually. The sample contains 76 contracts with an annual price greater than $100,000.

One important source of variation in the analysis is in the number of bids received. The median is 5 bids, and the maximum is 40. Thus, there is a good deal of competition for these contracts. The variation in the number of bids will help to 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 typically small, with a median of 14 employees. Over 25 percent of the winning suppliers have 3 or fewer employees.

Figure 3 plots the logged values of the winning bids on the y-axis against the number of bidders on the x-axis. The second panel displays residualized values for the (log) winning bids. The residuals were constructed from a regression of price on duration, square footage, cleaning frequency, baseline unemployment, and fixed effects for facility type. Even after controlling for observable characteristics, there is large variation in prices for auctions with many bidders. The pattern observed in the figure – large variation in prices with clustering at the median price, rather than the minimum – motivates the assumption of unobserved auction-specific heterogeneity used in the model. 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 costs fits far better.

The contracts in the dataset have a good deal of variation in duration. Figure 4 provides a histogram of duration in three-month intervals. There is a good deal of variation in duration, ranging from 5 months to 6.5 years, though contracts tend to cluster at yearly increments. Additionally, 53 percent of contracts are for 5 years, which is the typical maximum contract duration imposed by federal budgeting regulations. Longer durations require the contracting

Figure 3: Price versus Number of Bids

(a) Annual Price

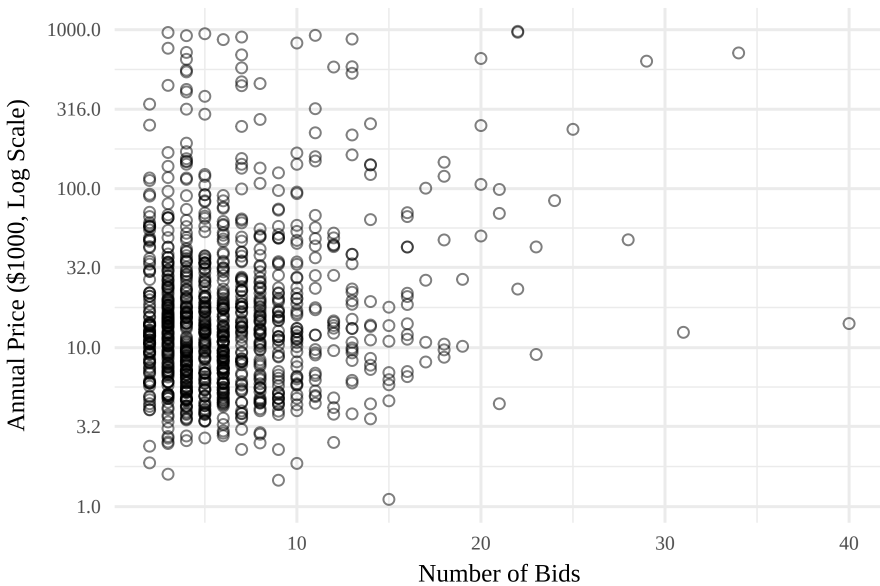

(b) Residualized Annual Price

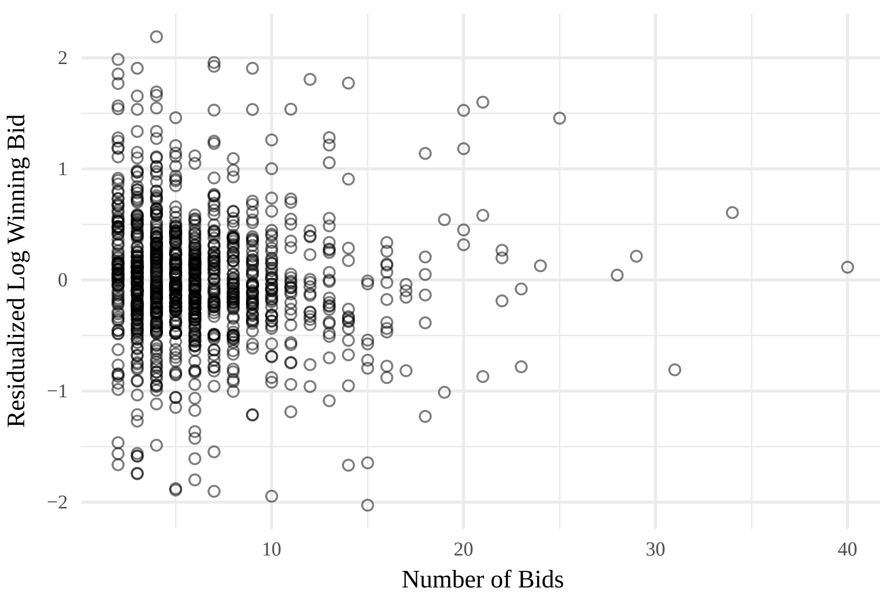

Notes: The figure plots the log annual price against the number of bids received for each contract. There is a great deal of variation in the annual price, much of which cannot be explained by observable variables. This is illustrated by the residualized bids in the lower panel. The R^2 of the regression used to construct the residuals, which includes duration, square footage, frequency, baseline unemployment, and fixed effects for facility type, is 0.74. It is notable that some of the highest and lowest prices are realized with few bidders.

Figure 4: Contract Duration

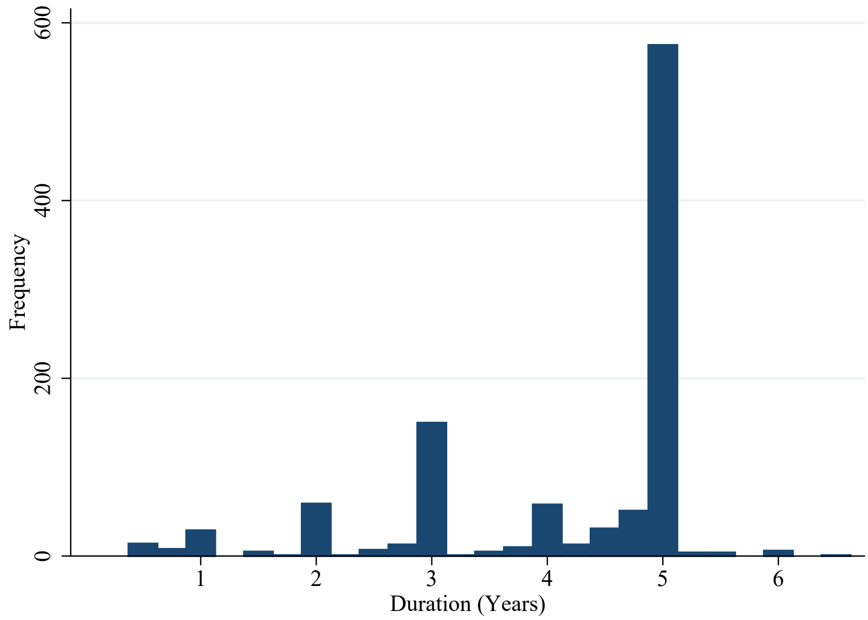

Notes: The figure displays a histogram of contract duration in 3-month bins. Over half of the contracts have a five-year duration, which is the maximum duration (by regulation) without specifically requesting an extension. Contracts are clustered in yearly intervals, though the support in between full years is relatively well-covered.

officer to request and justify an extension. The observed variation in duration, combined with the presence of a five-year cap on contract duration, help motivate the counterfactual analysis I perform in Section 6, where I consider the value of a flexible-term policy compared to one in which duration is standardized.

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 4 provides regressions of the log annual price on the number of bids, duration, and controls. The first three columns display the results from ordinary least squares regressions. Square footage alone, as reported in the first specification, captures 62 percent of the variation in prices.

To account for endogenous entry, I instrument for the number of bidders using time-series and cross-sectional variation in local labor market conditions, as well as variation in the type of bidders permitted to compete for the contract. The first instrument is the (log) ratio of unemployment, relative to the 2004 baseline, in the county, which generates a time-varying county-specific unemployment shock. The second instrument is the number of establishments for NAICS code 561720 (corresponding to building cleaning services) in the same 3-digit ZIP code.[34] It is plausible that an increase in unemployment or the presence of more firms in the

[34] I add 1 to the raw value to use the logged value in estimation, as a few contracts have zero in the raw value.

Table 4: Descriptive Regressions: ln(Annual Price)

	OLS-1	OLS-2	OLS-3	IV-1	IV-2
ln(Square Footage)	0.730***	0.658***	0.658***	0.689***	0.687***
	(0.018)	(0.017)	(0.017)	(0.024)	(0.024)
Number of Bids		−0.014***	−0.009*	−0.053**	−0.047**
		(0.005)	(0.005)	(0.022)	(0.022)
Duration (Years)		0.041***	0.032**	0.043***	0.033**
		(0.015)	(0.015)	(0.016)	(0.015)
ln(Weekly Frequency)		0.459***	0.394***	0.467***	0.407***
		(0.039)	(0.038)	(0.041)	(0.040)
ln(2004 Unemp.)		0.054***	0.037***	0.080***	0.060***
		(0.012)	(0.012)	(0.019)	(0.018)
High-Intensity Cleaning		0.586***		0.559***	
		(0.071)		(0.075)	
Building Type FEs			X		X
Observations	1046	1046	1046	1046	1046
R^2	0.62	0.71	0.74	0.69	0.73

Standard errors in parentheses

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

Notes: The table displays estimated coefficients from regressions of log annual price on auction characteristics. The variables from specification IV-1 are included in the structural model. These regressions show that square footage, cleaning frequency, and market characteristics explain much of the variation in prices. Once square footage, cleaning frequency, and market characteristics are accounted for, fixed effects for location type add little explanatory power. Specifications IV-1 and IV-2 are two-stage least squares regressions, where the instruments for the number of bids are monthly (log) county-level unemployment relative to 2004, the (log) number of NAICS code 561720 establishments in the same 3-digit ZIP code in 2004, and an indicator for whether the set-aside was for generic small businesses.

broader geographic area are not driven by unobservable characteristics of these contracts, yet they are likely to generate increased entry.

A third instrument is developed from the federal government practice of "setting aside" certain contracts for firms with particular types of owners. Specialized set asides include women-owned and veteran-owned small businesses. As we have removed economically disadvantaged set-asides (e.g., for Economically Disadvantaged Women-Owned Small Business) from the sample, it is plausible that the ownership type is uncorrelated with the underlying cost structure of the participating firms. If the cost structure is independent of ownership for these firms, then the type of set aside is a valid instrument for price (by affecting entry). This instrument is implemented as a binary variable with the value of 1 if the set aside is for generic small businesses.

The last two columns report the estimated coefficients from instrument variables regressions. Consistent with endogenous entry, I find a larger negative effect of the number of bidders

Table 5: Descriptive Regressions: Number of Bids

	(1)	(2)	(3)	(4)
Duration (Years)	0.104	−0.017	−0.002	−0.002
	(0.104)	(0.099)	(0.099)	(0.100)
ln(Square Footage)	0.760***	0.779***	0.834***	0.825***
	(0.111)	(0.106)	(0.106)	(0.112)
ln(Weekly Frequency)	0.487*	−0.081	0.009	0.137
	(0.254)	(0.247)	(0.253)	(0.257)
ln(2004 Unemp.)		−0.832***	−0.794***	−0.793***
		(0.239)	(0.238)	(0.238)
ln(Unemployment)		1.415***	1.420***	1.356***
		(0.232)	(0.231)	(0.231)
ln(Num. Firms in Zip3)		0.241	0.257*	0.276*
		(0.148)	(0.148)	(0.147)
Generic Set-Aside			1.134***	0.987***
			(0.350)	(0.361)
High-Intensity Cleaning			−0.294	
			(0.475)	
Building Type FEs				X
Observations	1046	1046	1046	1046
R^2	0.06	0.16	0.17	0.19
F-statistic	22.2	32.0	25.9	14.7

Standard errors in parentheses

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

Notes: The table displays estimated coefficients from regressions of the number of bids on auction characteristics and local labor market variables. Specification (3) is equivalent to the first-stage regression of IV-1 in Table 4. Specification (4) includes fixed effects for each building type.

on price compared to the corresponding OLS specifications. In the structural model of Section 5, I explicitly model entry to account for this endogeneity. The main motivating specification is IV-1, which uses square footage, weekly cleaning frequency, and baseline (2004) unemployment as controls. To capture variation in the types of buildings and cleaning required, IV-1 includes an indicator for "high intensity" cleaning of airports and medical buildings.[35] IV-2 includes indicators for all building types. The inclusion of fixed effects for all types have low in specifications OLS-3 and IV-2 have a low per-variable impact on R^2 and do not have a substantial effect on the estimated coefficients. Therefore, I omit them from the structural estimation and proceed with the variables used in IV-1.

[35] If separate indicators are estimated for medical buildings and airports, the coefficients on the indicators are very similar and the coefficients on the other variables are unchanged.

Though the linear model does not account for the offsetting effects of duration on price and (via profits) on entry, the regressions capture a positive relationship between price and duration. In the structural estimation, I also find a positive and significant direct relationship between duration and price.

In Table 5, I display regressions of the number of bids on auction characteristics and local measures of unemployment. Specification (3) is equivalent to the first-stage regression in IV-1, with an F-statistic of 25.9. All three instruments – the unemployment shock, the presence of existing firms, and a generic set-aside – have the expected positive signs. Though current unemployment is associated with more bids, higher baseline levels are associated with fewer bids. I interpret the negative correlation between higher 2004 unemployment and fewer bids as a reflection of local labor market frictions, leading to reduced competition and higher wages.

5 Estimation of the Structural Model and Transaction Costs

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 buyer 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. These estimates are inputs to the policy counterfactuals in Section 6.

5.1 Estimation of Entry and Bidding

For the contracts in my dataset, I estimate the three-stage model of Section 3.3, where bidders are symmetric and participate in a first-price auction. I employ a parametric approach for parsimony. The nonparametric identification results provided earlier, along with robustness checks, suggest that first-order features of the estimated distributions are not entirely driven by functional form. In this application, there is an added complication of estimating a duration-dependent distribution of private costs, which would increase the number of parameters needed for any nonparametric approach.

The parameterizations are given in Table 6. A central consideration of this paper is that the distribution of the average per-period private cost, C_i, changes with the duration of the contract. One approach to estimation would be to estimate a microfounded model where the per-period cost shocks are governed by an autocorrelation parameter. Instead, I estimate the average per-period cost distribution as a primitive, allowing the mean of the average per-period cost and the variance to vary with T. As I am not taking a stand on the underlying cost process, I am in one sense estimating a "reduced-form" primitive for the cost distribution.[36] By picking an appropriately flexible distribution, this approach may better approximate a wider range of per-period distributional families.

[36] For a microfounded model, see Appendix E.

Table 6: Empirical Parameterizations

Cost Component	Notation	Parameterization
Private Costs	C_i	$\sim Weibull(\mu_0 + \mu_1 T, \alpha_0 + \alpha_1 T)$
Unobserved Heterogeneity	U	$\sim \ln \mathcal{N}(-\frac{\sigma_U^2}{2}, \sigma_U^2)$
Entry Shock	ε	$\sim \ln \mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2)$
Observed Heterogeneity	$h(X)$	$= square\_footage^{\gamma_1} \cdot weekly\_frequency^{\gamma_2}$ $\cdot 2004\_unemployment^{\gamma_3} \cdot \gamma_4^{\mathbb{1}[high\text{-}intensity\_cleaning]}$
Entry Costs	$k(M)$	$= T \cdot square\_footage^{\kappa_1} \cdot weekly\_frequency^{\kappa_2}$ $\cdot unemployment\_shock^{\kappa_3} \cdot establishments^{\kappa_4}$ $\cdot \kappa_5^{\mathbb{1}[generic\_set\text{-}aside]}$

For private costs, the Weibull distribution is chosen for tractability and flexibility, as it allows the estimated probability density functions to be either convex or concave. It is governed by the mean parameter $\mu(T) = \mu_0 + \mu_1 T$ and the shape parameter $\alpha(T) = \alpha_0 + \alpha_1 T$. I allow the parameters of the private cost distribution to vary linearly with duration to capture the first-order effects of interest in this model. A finding of $\alpha_1 > 0$ would indicate autocorrelation in cost shocks and a reduced variance in average per-period costs due to longer contract duration. For the distribution of unobserved heterogeneity, the log-normal distribution was chosen because it best fit the model out of several choices.[37]

Using the entry and bidding problems below, I estimate these parameters using maximum likelihood. Entry costs shocks are parameterized as increasing linearly with the duration of the contract.[38] Note that square footage and weekly frequency affect the entry decision by both increasing supply costs (price) and affecting entry costs.[39]

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

One of the computational challenges in maximum likelihood estimation of auction models arises from the need to invert the bid function, which may be computationally costly. I employ a simple innovation to greatly speed up this process in the presence of unobserved heterogeneity, which I provide along with details of the likelihood function in Appendix F.

Table 7 displays the parameter estimates from the first-stage estimation. Square footage,

[37]Other estimated distributions of unobserved heterogeneity were the gamma distribution and the Weibull distribution. Both have the desirable properties of support on $(0, \infty)$ and can be normalized to have a mean of 1.

[38]This has the interpretation that entry costs are borne annually and could reflect the opportunity costs of other contracts. Allowing a free parameter on the entry costs in estimation generates a coefficient close to one.

[39]Since I assume that that U is not observed prior to entry, the information available to potential bidders is captured by X and M. Therefore, $E[U|n, t, x, m] = E[U] = 1$. This assumption may a reasonable approximation because observables explain roughly 70 percent of the variation in prices.

Table 7: Parameter Estimates

Group	Parameter	Variable	Estimate	95 Percent C.I.
Private Costs	μ_0	-	18.521	[16.658, 20.861]
	μ_1	Duration	0.546	[0.085, 0.979]
	α_0	-	4.807	[3.527, 6.969]
	α_1	Duration	0.386	[0.053, 0.674]
Heterogeneity	σ_U	-	0.608	[0.572, 0.646]
	γ_1	Square Footage	0.664	[0.627, 0.701]
	γ_2	Weekly Frequency	0.488	[0.411, 0.556]
	γ_3	2004 Unemployment	0.087	[0.070, 0.106]
	γ_4	High-Intensity Cleaning	0.302	[0.187, 0.437]
Entry	μ_ε	-	-0.459	[-0.749, -0.173]
	σ_ε	-	0.649	[0.608, 0.687]
	κ_1	Square Footage	0.543	[0.488, 0.599]
	κ_2	Weekly Frequency	0.542	[0.420, 0.650]
	κ_3	Unemployment Shock	-0.309	[-0.449, -0.208]
	κ_4	Establishments	-0.066	[-0.108, -0.025]
	κ_5	Generic Set-Aside	-0.262	[-0.390, -0.148]

Notes: The table displays maximum likelihood parameter estimates from the structural model. The first group of coefficients indicate how the mean and shape of the private cost distribution change with the duration of the contract. The second set of coefficients indicate the distribution of unobserved auction-specific heterogeneity and how auction-specific common costs vary with observable cost characteristics. The third set of coefficients pertain to entry costs in the model. 95 percent confidence intervals are displayed in the last column. As minor data cleaning steps (de-meaning) are data-dependent, confidence intervals are constructed via 500 bootstrap samples.

weekly frequency, and 2004 unemployment 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 $18,521 and increases by 2.9 percent per contract year (μ_1/μ_0).[40] Prices increase with duration due to both the increase in mean costs and the reduction in variance, as we would expect if cost shocks are not perfectly correlated over time. Reduced variance is captured by the positive coefficient α_1.

As expected, higher values for square footage and weekly frequency increase costs. Consistent with the findings from the descriptive regressions, baseline unemployment and high-intensity buildings have higher costs. For entry, higher current unemployment and the presence of more local establishments lower entry costs. Generic small business set-asides also have lower entry costs, relative to demographic-specific set-asides. 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 $\gamma_2 < \kappa_2$. This is consistent with capacity constraints, as some firms

[40] 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 8 in the Appendix.

Figure 5: Model Fit: Actual Versus Predicted Annual Price

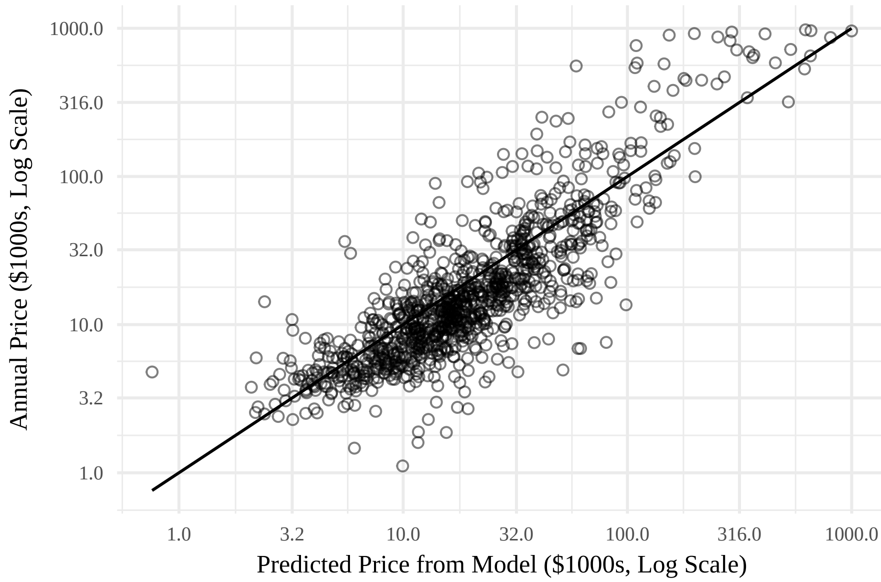

Notes: The figure plots observed prices against predicted prices from the model. The R^2 of the predicted values is 0.71, which compares favorably to the R^2 of 0.69 from the linear instrumental variables model.

may be limited in the days they are available to clean.

The model fits the data well. In Figure 5, I display actual values for annual prices compared to the predicted values. The R^2 for the structural model is 0.71, which compares favorably to the linear model IV-1 in Table 4. Though the correlation is very strong, unobserved heterogeneity is important to match the distribution of prices. Unobserved common costs are economically meaningful, in that they capture approximately 30 percent of the variance of log prices.

5.2 Estimation of Market Transaction Costs

In this section, I develop distribution-free bounds for market transaction costs based on the estimated parameters for entry and bidding. The buyer's decision problem[41] 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, the optimality condition can be used to construct bounds for a contract-specific δ via revealed preference. In my data, contracts are either set to the nearest monthly or nearest yearly increment, providing a set of tight and loose bounds, respectively.

Finally, to construct expected market transaction costs and conduct counterfactuals, I apply

[41] I maintain the assumption that $\beta = 1$, as a very low discount rate seems appropriate for this setting. Procurement offices tend to have a long tenure and are likely to see the renewal of the contract.

Table 8: Estimated Market Transaction Costs ($1000s)

Contract-Specific Measure	Median	95 Percent C.I.	Mean	25th Pct.	75th Pct.
Transaction Costs	10.3	[3.0, 17.2]	24.1	5.1	21.3
Annualized	2.4	[0.7, 4.0]	5.3	1.3	4.7
Contract Value	50.5	[46.9, 54.1]	190.2	28.5	102.0
Price (Annual)	13.2	[12.3, 13.9]	43.9	7.3	26.7
Percent Share of Costs	15.0	[4.9, 21.8]	16.1	9.8	20.5
Aggregate Measure	**Estimate**	**95 Percent C.I.**			
Percent Share of Costs	10.8	[3.8, 20.0]			

Notes: Estimated transaction costs are the expectation taken with a uniform prior over the distribution-free bounds identified from the duration decision of the buyer. For $T = 5$, conservative upper bounds are projected by assuming that the duration is optimally chosen. The median transaction cost in the data is $10,300. Transaction costs are also expressed as a share of total (buyer) costs. The aggregate share of total costs attributable to transaction costs is 10.8 percent, which is calculated by comparing the mean annualized transaction costs to the mean price. Confidence intervals are constructed via the bootstrap.

a uniform prior for the density between the distribution-free bounds.[42] Using the prior, I construct point estimates by taking the expectation. In practice, many of the contracts in my data face a cap on maximum duration of five years, due to federal regulation. For contracts affected by the cap, only a lower bound for δ can be obtained without additional assumptions. I make the conservative assumption that the chosen duration at five years is optimal. This generates a relatively conservative upper bound on transaction costs. Many of the optimal contracts under a higher cap would be likely be longer, implying larger transaction costs.

Market transaction costs are significant in this setting, comprising 10.8 percent of annual costs. Table 8 contains summary statistics for the estimated transaction costs. To obtain the aggregate share of costs attributable to transaction costs, I divide the mean annualized transaction costs by the mean total annual cost (the sum of annualized transaction costs and the price). Also displayed in the table are the median values for transaction costs, the annualized values, and the corresponding medians for contract value and price. The median market transaction cost is estimated to be $10,300, and the median share of costs attributable to transaction costs is 15 percent across contracts. 95 percent confidence intervals are captured via 500 bootstrap samples.

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 cleaning 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 work weeks, this gives a cost range of $4,500 to $5,400, which is roughly in line with the 25th percentile estimate. Larger projects may take

[42] The uniform prior is appealing for its transparency and also for the reason that the observed duration is optimal at the mean transaction cost when buyers 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 contracts observed in yearly increments.

months of work and multiple officers.

The sequential, revealed-preference approach has the benefit of providing testable implications of the model via the unconstrained estimates presented here. A finding of negative transaction costs, which would arise with private costs that fall with duration, would suggest that the tradeoff in this paper is not first-order to contract duration. Instead, the 95 percent confidence intervals of μ_1 and α_1 have positive support, implying positive transaction costs only, which is consistent with the model. As previously, the premium on duration can arise simply from averaging cost draws across multiple periods. In addition to this effect, a seller may charge a premium when she expects better options to arrive at a stochastic rate, as this will increase the opportunity cost over time.

In some cases, the cost of creating a market are quite large as a percent of total costs. The 95th percentile of share of costs attributable to market transaction costs is 32.2 percent. For these estimates, this is driven by moderate transaction costs realized by low-price projects, rather than very high absolute costs. For example, contracts with a portion of transaction costs in the 95th percentile or above (greater than 32.2 percent) have a mean price of $8,800, which is much smaller than the full-sample mean of $43,900.

Verifying the Estimated Transaction Costs

As an exercise to check the plausibility of the estimated market transaction costs, I project the estimates on other variables not used in the structural estimation. First, I calculate the median transaction costs by facility type and by department in Table 9. As expected, the highest transaction costs are among facilities with relatively complicated or technical requirements, such as medical centers, airports, and technical facilities (e.g., power plants). Simpler settings such as office cleaning have the lowest estimated transaction costs. In the second panel, I calculate the median by government department. The Department of Homeland Security has the highest median transaction costs, at $38,000 per contract. This might be expected given the high levels of security required at their facilities and the relative lack of institutional knowledge at the recently-formed department. Conversely, Agriculture and Defense have low median transaction costs, at $9,300 and $7,200, respectively. After controlling for square footage, cleaning frequency, and facility type in a regression, Homeland Security has the highest fixed effect for (log) transaction costs, 91 percent larger than Defense. Agriculture has lowest fixed effect, 16 percent smaller than Defense.[43]

In Table 10, I regress the estimated market transaction costs on variables excluded from the structural model. Included covariates are the number of pages in the contract, related expenditures and contract modifications[44] in the same 9-digit ZIP, and an indicator for whether the contract falls under the simplified acquisition protocol. One would expect that lengthier contracts and busier agencies are reflective of higher opportunity costs, and that the simplified

[43] Results available upon request.

[44] Spending and contract actions for other housekeeping services, maintenance, and office furniture.

Table 9: Estimated Market Transaction Costs ($1000s) by Category

(a) Location Type

Type	Median	Contract Value	Size (1000 SF)	Count
Medical	38.2	206.0	10.0	61
Airport	32.3	254.3	7.8	30
Technical	30.6	84.0	15.6	19
Industrial	28.4	60.0	27.9	13
Accommodations	28.0	121.2	32.0	18
Services	18.3	87.7	8.3	59
Research	13.2	58.5	6.0	111
Visitors	12.8	189.7	6.5	41
Field Office	9.1	48.5	8.4	270
Office	7.2	35.6	4.8	424

(b) Department

Department	Median	Contract Value	Size (1000 SF)	Count
Homeland Security	39.4	269.3	14.6	45
GSA	28.8	223.1	12.8	40
Veterans Affairs	24.0	143.6	8.8	80
Other	20.6	69.2	11.7	24
Commerce	13.9	59.7	5.5	78
Interior	11.9	70.9	8.9	43
Agriculture	9.3	46.7	9.3	347
Defense	7.2	35.9	4.2	389

Notes: The table displays the median estimated transaction cost to the buyer. Also displayed are the median contract value, the median square footage of the facility, and the count of observations in the sample. In panel (a), observations are grouped by location type. In panel (b), observations are grouped by contracting department.

Table 10: Projecting Market Transaction Costs on Variables Outside of the Model

	(1)	(2)	(3)	(4)	(5)	(6)
High-Intensity Cleaning	1.446***	1.188***	1.138***	1.020***	1.156***	0.611***
	(0.144)	(0.133)	(0.134)	(0.138)	(0.138)	(0.108)
ln(Word Count)	0.085***				0.124***	0.042**
	(0.024)				(0.023)	(0.018)
ln(Related Expenditures)		0.096***			0.073***	0.053***
		(0.011)			(0.023)	(0.017)
ln(Related Modifications)			0.283***		0.047	−0.140**
			(0.035)		(0.071)	(0.055)
Simplified Acquisition Ind.				−0.693***	−0.716***	−0.429***
				(0.090)	(0.089)	(0.068)
ln(Square Footage)						0.577***
						(0.026)
ln(Weekly Frequency)						0.510***
						(0.057)
Observations	1046	1046	1046	1046	1046	1046
R^2	0.09	0.14	0.13	0.13	0.20	0.54

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The 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, (iii) the count of contract actions for these expenditures, and (iv) an indicator for whether the contract falls under the federal government's simplified acquisition protocol.

acquisition label would reflect lower market transaction costs. Indeed, all four enter with the expected sign. After controlling for square footage and cleaning frequency, 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. The negative coefficient on contract modifications in the fifth specification may reflect economics of scale or simply that lower transaction costs lead to more contract modifications.

5.3 Robustness

Symmetry

For the structural model, we have proceeded under the assumption that bidders are symmetric with respect to private supply costs. In a dynamic setting, the procurement process might result in asymmetry between bidders that would invalidate this assumption. One common source of asymmetry in procurement is the presence of an incumbent bidder who may have an advantage via a relationship-specific investment (e.g., through learning-by-doing or lowered transaction costs of retaining the same supplier). Additionally, competing bidders may retain

Table 11: Test for Asymmetry: Do Incumbents Have an Advantage?

Follow-On Contracts	Symmetric Win Rate	Incumbent Win Rate	*N*	*t*-Statistic
Estimation Sample	22.4	21.7	175	(0.20)
Extended FPDS Sample	27.8	26.3	845	(1.00)

Notes: The table displays the results of a test for asymmetry in performance by incumbent bidders. The expected win rate (percent) for symmetric bidders, based on the number of bids, is compared to the observed win percentage by incumbent bidders. The *t*-statistics indicate no significant difference in either sample. The first sample is follow-on contracts in the estimation sample, and the second sample uses the same criteria for all FPDS building cleaning contracts. Follow-on contracts are identified as contracts that have a single leading contract for the same agency in the same nine-digit zip code. A leading contract is one that is active in the year prior to the start of the follow-on contract and begins at least thirty days prior to the start of the follow-on contract.

some information about competitors if costs are correlated over time.[45]

I check for the presence of asymmetries by comparing the expected win percentage under symmetry (based on the number of bidders) to the win percentage for incumbent suppliers in follow-on contract. I identify follow-on contracts in the analysis sample by finding contracts that have a single active supplier on another contract in the same 9-digit ZIP code within the prior year (and starting at least thirty days before). The prior contract may be any of the approximately 11,000 cleaning contracts in the FPDS data. I also construct a broader set of follow-on contracts from the extended FPDS sample.

Table 11 compares the expected win percentage for symmetric bidders to the actual win percentage for incumbent bidders in identified follow on contracts. There is no significant difference between the two, suggesting that the incumbency advantage is not first-order in this setting. I obtain similar results for the 175 contract in the estimation sample and the 845 contracts from the broader FPDS sample.[46] There are a priori reasons to believe that the incumbency advantage is not large for competitive federal procurement, as, per regulation, the agencies are mandated to seriously consider all qualified bidders and, in most cases, select the lowest price. The degree of relationship-specific investments in facility cleaning is likely to be low, as the menu of services tend to be standardized.

As a second test, I include a dummy for whether the contract is an identified follow-on contract 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 G.3. The results of these tests are consistent with the maintained assumption of no endogenous asymmetries.

An additional general concern might be that there is heterogeneity in supplier types. The

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

[46] 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.

above tests for endogenous asymmetries are also valid tests for exogenous asymmetries in supplier types. Lower-cost types would be more likely to win the first contract in the identified set of follow-on contracts, generating a correlation in win rates over time. Thus, the above findings are consistent with symmetry across suppliers more generally. In contrast to many other industries, there is no great distinguishing factor that separates types of building-cleaning firms, and it is reasonable to expect that production is roughly constant returns-to-scale. This makes the empirical setting a nice fit for the model.

Buyer Behavior

The empirical model makes two simplifications regarding buyer behavior that merit discussion. The first is that buyers choose fixed-price, fixed-duration contracts. The reader may wonder why alternative contract structures are not considered, such as "evergreen" contracts, which provide indefinite mutual options for renewal. The present empirical study is concerned with the contracts as they appear in the real world, and, to the best of the author's knowledge, these fixed-duration contracts remain the most common form of buyer-seller contract. Furthermore, these simple contracts provide a nice setting to study the cost of market transactions, as the end a contract reflects ex ante information, rather than ex post. There are reasons why fixed-duration contracts may be optimal (e.g., to align incentives within the buying firm and discourage collusion), but an analysis of this kind lies outside the scope of this paper.

Another simplification is that the market transaction costs are exogenous. In other settings, we may expect firms to put forth effort to create a market, and this effort may be correlated with other variables in the model. For example, if the number of bidders is low, the buyer may exert additional effort to solicit more bids. Here, the standardized environment and strict criteria imposed by regulation ensure that the market costs are roughly invariant with respect to duration. When examining other settings, it would be important to consider extensions allow market transaction costs to depend on buyer effort and react to market characteristics.

Sensitivity of Estimated Transaction Costs

In constructing an empirical model, one must often make restrictions to make estimation possible. Most of the assumptions follow standard practice in the literature, such as the multiplicative separability in cost components. As mentioned, I have tested alternative assumptions for the distribution of unobserved heterogeneity and the realization of entry costs.

Two assumptions will tend to underestimate the market transaction costs. The first is the assumption that $\beta = 1$. Under a higher discount rate, the implied transaction costs would be larger. This is because δ is amortized by the factor $\frac{1}{\sum_{k=1}^{T} \beta^{k-1}}$, which increases when β is smaller. As the amortization factor increases, transaction costs get a relatively larger weight in the buyer's objective function. This, in turn, implies that a greater transaction cost is needed to rationalize the chosen duration.

Second, many five-year contracts are chosen because of the cap on contract duration imposed by federal regulation. For the contracts for which this constraint is binding, the transaction costs would be higher than the estimates. The extent to which the estimates are conservative depends on the underlying distribution of optimal contracts for those at the cap.

6 Counterfactuals: The Impact of the Duration Margin

6.1 The Costs of Standardization

When the buyer can adjust the non-price terms of a contract, the buyer can minimize expected costs for each transaction. This flexibility provides a cost-minimizing advantage to the firm. In many settings, contract terms are standardized. For example, a three-year contract is the industry standard for office supplies. The structural model allows us to estimate how costly such standard-duration contracts would be in the empirical setting of this paper, holding fixed the cost of going to the market.

Though moving away from the optimal duration will raise costs, one might expect that, by simplifying the contracting procedure, market transaction costs would fall. Thus, it is worthwhile investigating the reduction in market transaction costs necessary to offset the increased costs arising from a standardized policy. This compensating transaction cost would make the firm (the government) indifferent between the flexible-duration and standard-duration policies.

Table 12 reports the impact on aggregate buyer costs by moving to standardized terms of yearly increments. Total costs would increase substantially, by 36 percent, if all contracts were issued in one-year terms. This is not surprising, as the median duration in the data is 5 years and standardization would result in much more frequent contracting. On the other hand, standard durations of 4 years or 5 years would have a small impact, increasing buyer costs by less than two percent.[47]

In the final column of Table 12, I report the change in transaction costs that would make the standardized term policy equivalent to the flexible term policy. For a four-year standard term, the necessary reduction in transaction costs is modest. If the government could reduce transaction costs by 10 percent by implementing a standardized four-year duration policy for these services, the results indicate that it would be beneficial to do so.

These results suggest that flexible terms may be quite valuable, compared to a poorly-chosen standard (e.g., one year or two years in this setting). Thus, knowledge of the relevant cost structure and transaction costs is important for setting non-price terms. However, an intelligently-chosen standard may be cost-effective, as the required reduction in transaction costs to offset the costs of standardization is modest for a four-year standard.

A related question to standardized terms is that of a cap on maximum duration, similar

[47] Recall that a conservative assumption was used to calculate transaction costs, so the impacts on total costs are also conservative.

Table 12: Effects of Standardized Terms

$\bar{T}$	Total Cost	95 Percent C.I.	Price	Trans. Cost	Affected	Change to δ
1	36.0	[11.9, 54.4]	−11.9	353.9	1018	−60.6
2	9.7	[3.1, 14.6]	−8.0	126.9	992	−32.6
3	3.1	[1.0, 4.7]	−4.2	51.3	907	−15.7
4	1.4	[0.4, 2.3]	−0.4	13.5	992	−9.3
5	1.6	[0.5, 2.8]	3.2	−9.2	496	−13.3
6	2.7	[0.9, 4.4]	6.8	−24.4	1041	−27.4

Notes: The table displays the resulting percent changes in total costs, prices, and annualized transaction costs when all contracts are issued in standardized durations corresponding to $\bar{T}$. For a uniform duration policy of 4 years or less, the average price paid decreases and the amount spent on transaction costs increases. Affected contracts are the count of those that are displaced from the optimal duration. The final column displays the reduction in transaction costs that would render a uniform policy equivalent to the existing policy in terms of buyer costs. Confidence intervals are reported for total costs and are constructed via the bootstrap.

to the five-year cap imposed by government-wide budgeting regulations in my data. This is analogous to the imposition of standard terms on only a subset of contracts. In Appendix G.4, I provide a detailed breakdown of the effects by whether duration is increased or decreased by the standard, which provides insight into the cost of the cap.[48]

6.2 Contract Duration and Welfare Analysis

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 two typical forms of welfare analysis: an event study and a counterfactual simulation with a structural model. For both approaches, assume that transaction costs and the duration margin are ignored. Any counterfactual evaluation in the structural model will find no impact from a policy that changes 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

[48] Table 17 reports averages by contract, rather than in aggregate, which is why the numbers differ slightly from those in Table 12.

Table 13: Effects of a 20 Percent Reduction in Transaction Costs

Variable	Symbol	Percent Change	95 Percent C.I.
Total Costs	$P + \frac{\delta}{T}$	-2.8	[-4.2, -1.0]
Prices	P	-1.7	[-2.7, -0.5]
Duration	T	-11.4	[-13.2, -10.3]
Percent Captured by Prices		62.6	[52.4, 71.2]

Notes: The 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 estimated treatment effect 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. Confidence intervals are constructed via the bootstrap.

impact of transaction costs on prices and, if measured, 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 13 indicates, buyers respond to lower transaction costs by issuing shorter contracts, reducing prices by 1.7 percent. Total costs, accounting for the direct effect on transaction costs, fall by 2.8 percent. Thus, the hypothetical difference-in-differences analysis would capture only 63 percent of the change in total costs, illustrating the importance of accounting for transaction costs.

7 Discussion and Generalizability

The estimated market transaction costs are obtained in a specialized setting which allows for the isolation of market transaction costs. One nice feature about the application is that there is a good deal of variation in observables, which allows us to look at how these costs vary with certain factors. We might expect that transaction costs are somewhat lower for buyers in the private sector, who face similar sets of suppliers but may be able to avoid some of the requirements mandated by regulation. On the other hand, the size of government procurement may provide some economies of scale and reduce these costs. Furthermore, buyers in other contexts (including a portion of government procurement) often have the option to choose what type of competitive procedure to engage in and the type of contractual arrangement, such as choosing to negotiate with the incumbent supplier or signing an evergreen contract.

The nature of supply costs is an important consideration in more general settings. With building cleaning services, supply is stable over time, and the service represents a small fraction of overall expenditures. In other settings, such as technology-dependent firms, the per-period supply costs may be falling over time. Whether or not this feature leads to shorter contracts depends on whether cost-reducing innovations are predictable and whether they are driven by high-cost or low-cost firms. Furthermore, buyers may not be risk-neutral with respect to a primary input that is specialized to their business, which is another factor to consider when taking the model to other settings. Even so, the tradeoff I identify here remains relevant.

Efficiency

For the analysis of the paper, I have focused on the buyer-optimal contracts, because the duration of a fixed-price, fixed-duration contract is typically determined by the buyer. The intuition developed for the buyer-optimal contracts translates to the efficient contracts as well. The response of supply costs to the duration of the contract depends on the properties of order statistics; the price in a symmetric auction reflects the second order statistic, whereas the cost (which is needed for the efficient contract) reflects the first order statistic. In addition, supply-side frictions, including entry costs, should be factored in when estimating welfare effects. Other mechanisms may generate dependencies between marginal costs and contract duration. For example, learning-by-doing would reduce the seller's opportunity cost, generating marginal costs that decline in duration over some range. Regardless, we should still expect that marginal costs are increasing in duration at the equilibrium, the buyer will opt for a longer contract if both the supply price and amortized transaction costs are declining.

These additional factors raise a question about property rights: should the buyer or the sellers be endowed with the ability to decide duration? In the Appendix, I show how market-determined contracts may differ from efficient contracts. Appendix A explores the efficient contract for the simple numerical example, and Appendix B examines the efficient contract in the more general empirical model and the allocation of duration-setting rights. Imperfect competition drives a wedge between the revenue-optimal contracts and the duration 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. 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.[49]

Bundling

There is a direct connection between the model of contract duration and bundling. Fixed-duration contracts can be thought of as bundling demand over time. The analysis here could be reinterpreted to allow T to represent the bundle size and δ represent the transaction cost

[49] 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.

for each bundle. Thus, we obtain predictions relating the underlying variance of costs (or valuations) to the optimal bundle size, as well as the effect of competition on optimal bundling. The model could also be applied to any characteristic that have a “scale” effect with respect to a cost δ, like duration or bundle size.

8 Conclusion

In this paper, I develop a model of optimal contract duration arising from underlying supply costs and market transaction costs. I show how latent transaction costs may be recovered from the duration decision of the buyer. Using a dataset of federal supply contracts, I find that the costs of going to the market can be a significant portion of total costs for intermediate goods. The methods developed in this paper may prove useful for welfare analysis, especially in industries where long-term relationships are prevalent. In many settings, the tradeoff presented in this paper may complement other concerns, such as ex post incentive problems (e.g. arising from asset specificity) and uncertainty. An appropriate model should be tailored to the industry in question.

The analysis presented here offers, albeit indirectly, one novel prediction regarding the theory of the firm. Supply contracts lie in between arms-length transactions and vertical integration. As is known, conditions favorable for long-term contracts are also favorable for vertical integration, as the end is similar and integration may result in additional benefits. I demonstrate here that long-term contracts arise when competition is sufficiently low, and also when competition is very intense. Likewise, vertical integration may be most likely for low levels and high levels of competition. When the industry is moderately competitive, a downstream firm realizes a large benefit by switching among suppliers and may have the smallest incentive to integrate upstream.

References

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

ATALAY, E., A. HORTAÇSU, M. J. LI, AND C. SYVERSON (2017): "How Wide Is the Firm Border?" Working paper.

BAJARI, P., S. HOUGHTON, AND S. TADELIS (2014): "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," *American Economic Review*, 104, 1288–1319.

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.

BHATTACHARYA, V. (2018): "An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program," Working paper.

CABRAL, L. (2016): "Dynamic Pricing in Customer Markets with Switching Costs," *Review of Economic Dynamics*, 20, 43–62.

CANTILLON, E. AND M. PESENDORFER (2006): "Combination Bidding in Multi-Unit Auctions," 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. H. (1937): "The Nature of the Firm," *Economica*, 4, 386.

——— (1960): "The Problem of Social Cost," *Journal of Law and Economics*, 3, 1–44.

DECAROLIS, F. (2014): "Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions," *American Economic Journal: Applied Economics*, 6, 108–132.

DECAROLIS, F., L. M. GIUFFRIDA, E. IOSSA, V. MOLLISI, AND G. SPAGNOLO (2018): "Bureaucratic Competence and Procurement Outcomes," Working paper.

DUBÉ, J.-P., G. J. HITSCH, AND P. E. ROSSI (2010): "State Dependence and Alternative Explanations for Consumer Inertia," *The RAND Journal of Economics*, 41, 417–445.

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.

GRAY, J. A. (1978): "On Indexation and Contract Length," *The Journal of Political Economy*, 1–18.

GREENSTEIN, S. M. (1993): "Did Installed Base Give an Incumbent Any (Measureable) Advantages in Federal Computer Procurement?" *The RAND Journal of Economics*, 24, 19.

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

HANDEL, B. R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *The American Economic Review*, 103, 2643–2682.

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

HONKA, E. (2014): "Quantifying Search and Switching Costs in the US Auto Insurance Industry," *The RAND Journal of Economics*, 45, 847–884.

HU, Y., D. MCADAMS, AND M. SHUM (2013): "Identification of First-Price Auctions with Non-Separable Unobserved Heterogeneity," *Journal of Econometrics*, 174, 186–193.

ILLANES, G. (2017): "Switching Costs in Pension Plan Choice," Working paper.

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

KANG, K. AND R. A. MILLER (2017): "Winning by Default: Why Is There So Little Competition in Government Procurement?" Working paper.

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.

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.

QUINT, D. (2015): "Identification in Symmetric English Auctions with Additively Separable Unobserved Heterogeneity," Working paper.

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.

STIGLER, G. J. AND J. K. KINDAHL (1970): *The Behavior of Industrial Prices*, vol. 90, New York, NY: National Bureau of Economic Research.

WALKER, G. AND D. WEBER (1984): "A Transaction Cost Approach to Make-or-Buy Decisions," *Administrative Science Quarterly*, 373–391.

ZHOU, J. (2017): "Competitive Bundling," *Econometrica*, 85, 145–172.

A Additional Results for the Simple Model

A.1 Mathematical Illustrations of Predictions

Recall the decision rule for the buyer. The buyer will choose a long-term contract if

$$E[\tilde{c}_{2:N}] - E[c_{2:N}] < \frac{\delta}{2}.$$

First, it is straightforward to see that higher marginal costs, by raising the left hand side of the above equation, increase the cost of long-term contracts and the likelihood that it outweighs the savings in transaction costs. An increase in transaction costs raises the RHS and leads to shorter contracts.

To provide a slightly more formal treatment of the predictions, I repeat them below and provide mathematical illustrations.

Remark 3 The optimal duration is increasing with autocorrelation in supply costs.

Suppose that d is a cost process with lower autocorrelation than c, but the same per-period marginal distribution, i.e. $E[d_{2:N}] = E[c_{2:N}]$. Let $\tilde{d}$ denote the average cost across two periods. Then it follows that, for $N > 3$,

$$\begin{aligned} E[\tilde{d}_{2:N}] &> E[\tilde{c}_{2:N}] \\ \implies E[\tilde{d}_{2:N}] - E[d_{2:N}] &> E[\tilde{c}_{2:N}] - E[c_{2:N}]. \end{aligned}$$

The marginal cost of long-term contracts is decreasing with the autocorrelation of the cost process. With greater autocorrelation, long-term contracts are preferred.

Remark 4 The optimal duration is decreasing in the variance of costs across suppliers, provided there is sufficient competition ($N > 3$).

For a simple case, consider location-scale transformations of c, such that $d = a + bc$ and $E[d] = E[c]$. Under the marginal cost structure d, a longer contract is chosen if

$$b \cdot (E[\tilde{c}_{2:N}] - E[c_{2:N}]) < \frac{\delta}{2}.$$

As b increases, shorter contracts become more desirable.

Remark 4' When costs are bounded from below, the optimal duration is U-shaped in the variance in costs, provided there is sufficient competition ($N > 3$).

If we impose the reasonable restriction that costs are bounded from below, greater variance may induce longer contracts by pushing the expected average price and expected per-period

Figure 6: Increased Variance in Cost

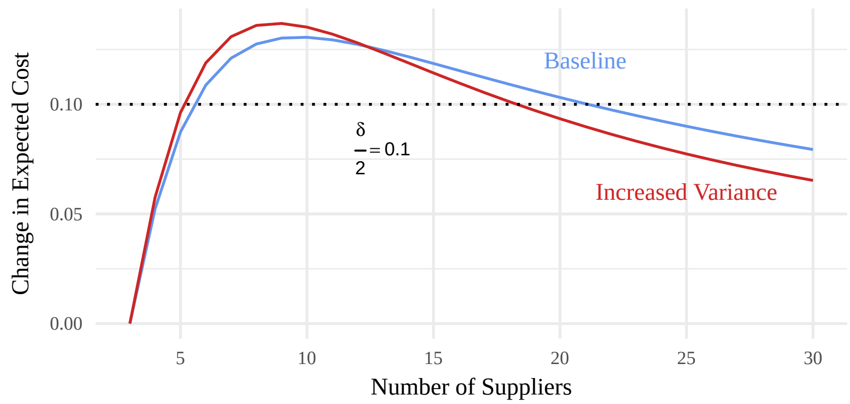

Notes: The blue line shows the marginal cost to the buyer of a two-period contract relative to one-period contracts and is equivalent to the blue line in Figure 1. The red line shows the marginal cost of a longer contract when the costs are draw from the same distributional family (the beta distribution) with 11 percent greater variance.

price close to the lower bound. Let c denote per-period marginal costs with a lower bound at 0, and let σ represent its standard deviation. Then, when $N > 3$,

$$\lim_{\sigma \to \infty} E[\tilde{c}_{2:N}] = \lim_{\sigma \to \infty} E[c_{2:N}] = 0.$$

As $E[\tilde{c}_{2:N}] - E[c_{2:N}] \to 0$, long-term contracts are optimal in the limit. This effect tends to dominate as N gets large, as more draws brings the minimum price closer to the lower bound.

A.2 An Increase in the Variance of Costs

Figure 6 displays the change in marginal costs when variance of the cost distribution increases by 11 percent. The costs are drawn from symmetric beta distributions with $(\alpha, \beta) = (0.4, 0.4)$. For $N < 12$, greater variance increases the cost of long-term contracts. When $N \geq 13$, the winning supplier's price is close enough to the lower bound to reduce the cost.

This graph illustrates how the underlying cost structure leads to Remark 4', as an increase in variance leads to shorter contracts when competition is lower and longer contracts when competition is intense.

A.3 The Efficient Case

Though the analysis heretofore has focused on the buyer-optimal case, coinciding with the typical market outcome, it is worthwhile to consider the efficient case, which minimizes total social costs. The social planner's problem is similar to the buyers problem, except that the social planner will choose a long-term contract if

$$E[\tilde{c}_{1:N}] - E[c_{1:N}] < \frac{\delta}{2}.$$

Figure 7: Comparing Buyer-Optimal and Efficient Contracts

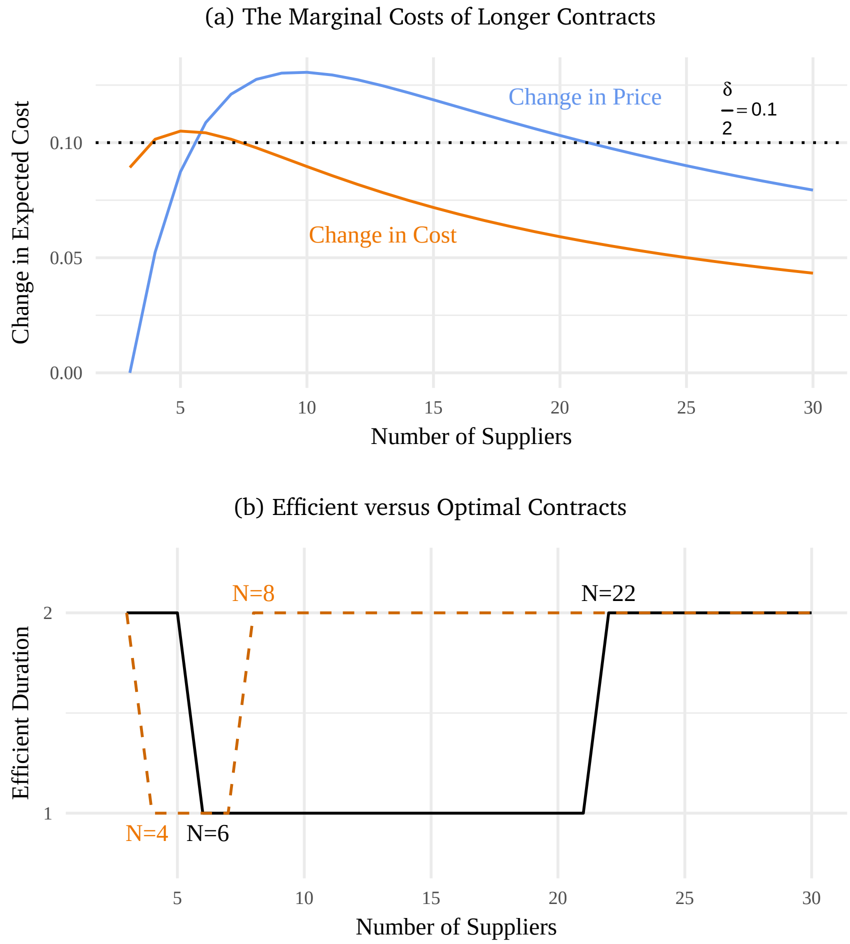

Notes: The figure shows the relationship between competition, the marginal costs of longer contracts, and the effect on buyer-optimal and efficient durations. The blue line in panel (a) shows the marginal cost to the buyer of a two-period contract and is equivalent to the blue line in Figure 1. The orange line shows the increase in marginal social costs of a longer contract. The dash line reflects a transaction cost of 0.20 amortized over two periods. For values of N where the blue line is above the dashed line, $N \in \{6, ..., 21\}$, the buyer would prefer to issue one-period contracts, as the increase in price is greater than the savings in transaction costs. This range does not coincide with the efficient contract, which is plotted with the orange dashed line in panel (b). One-period contracts are efficient for $N \in \{4, .., 7\}$. The buyer will select the efficient contract in this example only if $N \in \{6, 7\}$.

Thus, the social planner's decision depends on the first-order statistic of cost draws, rather than the second-order statistic that generates the expected price. For $N > 3$, these order statistics have similar qualitative behavior, so the comparative static predictions follow for the efficient case. However, the efficient contract will not necessarily coincide with the buyer-optimal contract, raising the question of the allocation of rights to non-price terms of the transaction.

Figure 7 provides a comparison of the buyer-optimal and the efficient contract. Panel (a) displays the marginal impact on the price to the buyer of a longer contract with a blue line. The orange line displays the marginal impact on supply costs. Once N is large enough ($N > 5$ in the example), longer contracts have a greater marginal effect on price than cost. This is intuitive, as the first-order statistic approaches the lower bound faster than the second-order statistic.

Panel (b) plots the efficient contract with an orange dashed line. It has similar qualitative features to the buyer-optimal contract, displaying the inverse U shape. The efficient and buyer-optimal contract coincide only when $N \in \{6, 7\}$. When $N = 4$, the buyer would choose a long-term contract when the short-term contract is efficient, and when $N \in 8, ..., 21$ the buyer would choose a short-term contract when a long-term contract is efficient. Thus, the buyer-optimal contract may be longer or shorter than the efficient contract. Information rents from private costs drive a wedge between the buyer-optimal contract and the efficient contract.

B Extensions of the Model: Efficiency and Allocation of Rights

In this section, I explore the relationship between optimal and efficient contract duration given the more general model of Section 3. It should be noted that the analysis here is not restricted to the special case of the duration-setting problem, rather, any transaction characteristic that has a "scale" effect (as duration does on transaction costs) can be related to this framework. One of the natural extensions is to bundling, where T is the size of a bundle (determined by the buyer or seller) and δ is the transaction cost for the bundle.

B.1 A Framework Relating Optimal and Efficient Contract Duration

Section 3 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.[50] Let $\overline{C}$ denote the ex ante expected price conditional on (T, X, M), so that $\overline{C}(T, X, M) = \sum_{n=1}^{N}(E[C(n, T, X, M)] \cdot \Pr(N = n|T, X, M))$. Thus, the ex ante efficient $\tilde{t}$ contract solves

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

with the first-order condition

$$\frac{d\overline{C}(T, x, m)}{dT}|_{T=\tilde{t}} = \frac{\delta}{\tilde{t}^2}. \tag{14}$$

In general, $\frac{d\overline{C}(T,x,m)}{dT}|_{T=\tilde{t}} \neq \frac{d\overline{P}(T,x,m)}{dT}|_{T=\tilde{t}}$, which will result in an inefficiency when the contract is determined by the buyer. As long as interior solutions exist (see Proposition 10), we have the result that the efficient contract $\tilde{t}$ will be longer than the buyer-optimal contract when $\frac{d\overline{C}(T,x,m)}{dT}|_{T=\tilde{t}} < \frac{d\overline{P}(T,x,m)}{dT}|_{T=\tilde{t}}$

[50] 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.

Defining the expected seller surplus as $E[\pi(T,x,m)] = \overline{P}(T,x,m) - \overline{C}(T,x,m)]$, we have the following result:

Proposition 9. *When interior solutions to the buyer's problem and the social planner's problem exist, the efficient contract will be longer than the equilibrium (buyer-optimal) contract if and only if the expected seller surplus is increasing at $\tilde{t}$:*

$$\begin{aligned} \tilde{t} > t &\iff \left(\frac{d\overline{P}(T,x,m)}{dT}|_{T=\tilde{t}} - \frac{d\overline{C}(T,x,m)}{dT}|_{T=\tilde{t}} \right) > 0 \\ &\iff \frac{dE[\pi(T,x,m)]}{dT}|_{T=\tilde{t}} > 0 \end{aligned}$$

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 (7) and (14) can be satisfied and $\overline{P}(T,x,m)$ and $\overline{C}(T,x,m)$ are not too concave. In particular, $\frac{d^2\overline{P}(T,x,m)}{dT^2}|_{T=\tilde{t}} > -\frac{2}{\tilde{t}}\frac{\overline{P}(T,x,m)}{dT}|_{T=\tilde{t}}$ and $\frac{d^2\overline{C}(T,x,m)}{dT^2}|_{T=\tilde{t}} > -\frac{2}{\tilde{t}}\frac{\overline{C}(T,x,m)}{dT}|_{T=\tilde{t}}$. These are the second-order conditions to ensure that first-order conditions above achieve a minimum.*

B.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:

$$\begin{aligned} &\min_T \overline{P}(T,x,m) - \overline{C}(T,x,m) + \overline{C}(T,x,m) + \frac{\delta}{T} \\ = &\min_T E[\pi(T,x,m)] + \overline{C}(T,x,m) + \frac{\delta}{T} \end{aligned}$$

Notice that when $\frac{dE[\pi(T,x,m)]}{dT} = 0$, this problem is equivalent to the social planner's 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 along with the transaction costs to the sellers? Sellers solve the problem:

$$\begin{aligned} &\max_T \overline{P}(T,x,m) - \overline{C}(T,x,m) - \frac{\delta}{T} \\ = &\min_T -\overline{P}(T,x,m) + \overline{C}(T,x,m) + \frac{\delta}{T} \end{aligned}$$

Sellers solve the social planner problem when $\frac{d\overline{P}(T,x,m)}{dT} = 0$. Therefore, if price is not sensitive to contract duration, it is efficient to let the sellers determine the length of the contract.[51]

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 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 non-price terms, such as 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

$$\min_T \mu\overline{C}(T,x,m) + \frac{\delta}{T}$$

and the seller's problem is

$$\min_T (1-\mu)\overline{C}(T,x,m) + \frac{\delta}{T}$$

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.

B.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

$$\tau_B = \tilde{t}^2 \frac{dE[\pi(T,x,m)]}{dT}|_{T=\tilde{t}}$$

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.

[51] Sellers have an equivalent rule to Proposition 9: $t_S > \tilde{t} \iff \frac{d\overline{P}(T,x,m)}{dT}|_{T=\tilde{t}} > 0$. This means that either 1) $t_S \geq \tilde{t} \geq t$, 2) $t \geq \tilde{t} \geq t_S$, or 3) $t_S \geq \tilde{t} \cap t \geq \tilde{t}$. 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.

$$
\begin{aligned}
\tilde{t} &= \arg\min_T E[\pi(T,x,m)] + \overline{C}(T,x,m) + \frac{\delta + \tau_B}{T} \\
&= \arg\min_T E[\pi(T,x,m)] + \frac{\tau_B}{T} + \overline{C}(T,x,m) + \frac{\delta}{T} \\
&= \arg\min_T \overline{C}(T,x,m) + \frac{\delta}{T}
\end{aligned}
$$

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

$$
\tau_S = -\tilde{t}^2 \frac{d\overline{P}(T,x,m)}{dT}|_{T=\tilde{t}}
$$

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

C Model Proofs

C.1 Proof of Proposition 1

Proof. Rearranging the FOC and taking the total derivative at the optimum (t), we obtain

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

From the second-order condition for a minimum,

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

.

Therefore

$$2T\frac{d\overline{P}(T,x,m)}{dT}|_{T=t} - 2\frac{\delta}{T} < \frac{d\delta}{dt}$$

As the LHS is equal to the first-order condition, it is zero, and therefore $\frac{dt}{d\delta} > 0$. □

C.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

$$\frac{dT}{dM} = \frac{-\frac{d^2\overline{P}(T,X,M)}{\partial T\partial M}}{\frac{\partial^2 d\overline{P}(T,X,M)}{\partial T^2} + \frac{2\delta}{T^3}}. \tag{15}$$

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

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

Likewise, this also holds for X. □

D Identification Proofs

D.1 Some Lemmas

To demonstrate the following proofs, 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

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

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

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

□

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:

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

Rearranging, we obtain
$E[c_{k:N}] = \frac{k}{N+1} E[c_{(k+1):(N+1)}] + \frac{N+1-k}{N+1} E[c_{k:(N+1)}]$. □

D.2 Proof of Proposition 3

Consider the entry equation

$$
\begin{aligned}
E[\pi_n|n,t] \cdot h(x) \cdot U - k(m) \cdot \varepsilon > 0 \quad & \iff \quad N \geq n \\
E[\pi_n|n,t] \cdot \frac{h(x)}{k(m)} > \frac{\varepsilon}{U} \quad & \iff \quad N \geq n
\end{aligned}
$$

For any realization (t, x, m), $\exists (t, x', m')$ such that $U|(N, t, x, m) = U|(N, t, x', m')$.[52] Using this relation, 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, we consider the identification of the unobserved components of transaction price.

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

D.2.1 Identification of Offers and Unobserved Heterogeneity

For any realization (n, t), the expected offer can be identified by finding (x, m) such that $U|(n, t, x, m) = U|(n_0, t_0, x_0, m_0)$. 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:

$$\begin{aligned} E[P|n,t,m,x] &= E[B|n,t] \cdot E[U|n,t,x,m] \cdot h(x) \\ &= E[B|n,t] \cdot h(x) \end{aligned}$$

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).

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.

D.3 Proof of Proposition 6

The ratio of profits is given by

$$R = \frac{E[\pi_n|n,T]}{E[\pi_{n'}|n',T]} = \frac{\frac{1}{n}\left(E[B|n,T] - E[C|n,T]\right)}{\frac{1}{n'}\left(E[B|n',T] - E[C|n',T]\right)} \tag{16}$$

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

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

Therefore, $E[C_{1:(n+1)}]$ is identified. $E[C_{1:n}]$ is obtained from equation (16).

D.4 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.

D.5 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 $E[C_{2:n}]$. [53]

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

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

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

$$\begin{aligned} (B_n\cdot U)\cdot\left(\tilde{\hat{B}}_{n'}\cdot\tilde{\hat{U}}\right) &\overset{d}{=} \left(\hat{B}_n\cdot\hat{U}\right)\cdot\left(\tilde{B}_{n'}\cdot\tilde{U}\right) \\ \Longrightarrow B_n\cdot\tilde{\hat{B}}_{n'} &\overset{d}{=} \hat{B}_n\cdot\tilde{B}_{n'} \end{aligned}$$

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

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

[53] 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.

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 $n' = n + 1$.

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

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.

D.5.1 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.

E 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 provide a model of underlying costs that generates both the distribution of costs and how duration affects 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 contract with duration T will have the same unconditional distribution as any other contract with duration 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 contract T solves

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

This results in the first-order condition

$$\begin{aligned} E[z_{1:N}]\xi'(T) &= \frac{\delta}{T^2} \\ -\xi'(T)T^2 &= -\frac{\delta}{E[z_{1:N}]}. \end{aligned} \tag{17}$$

In this case, we obtain a monotonic relationship between the number of bidders and the optimal duration. Unlike the U-shape models, the microfounded model here does not have a lower bound on costs.

Proposition 11. *The efficient duration 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 17. As $\frac{d}{dT}\left(-\xi'(T)T^2\right) < 0$, the optimal T falls. □

Further, the analysis above holds for the second-order statistic when $N > 3$, so we can extend the results to the buyer-optimal contract:[54]

Proposition 12. *The buyer-optimal duration is decreasing in the number of bidders. It is optimal for the buyer to issue a permanent contract for* $N \in \{2, 3\}$.

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

Proposition 13. *The efficient duration is less than the buyer-optimal duration.*

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

F Likelihood Function

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

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

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

For estimation, I make the assumption that $U \perp\!\!\!\perp (X, M)$. As U is not observed by the buyer 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

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

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

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

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

F.1 A Computational 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 $B \perp\!\!\!\perp U$.

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

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:

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

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 nonparametric approximations such as B-splines.

G Supplemental Tables and Figures

G.1 Count of Sites by Government Agency

Table 14: Count of Sites by Contracting Agency

Agency	Count	Percent
Defense	389	37.2
Agriculture	347	33.2
Veterans Affairs	80	7.7
Commerce	78	7.5
Homeland Security	45	4.3
Interior	43	4.1
GSA	40	3.8
Energy	5	0.5
Labor	5	0.5
Transportation	4	0.4
EPA	2	0.2
State	2	0.2
National Archives	2	0.2
CNCS	1	0.1
Health And Human Services	1	0.1
OPIC	1	0.1
Railroad Retirement Board	1	0.1
Total	1,046	100.0

G.2 Distributions of Bidder Costs

Figure 8: The Distribution of Bidder Costs

(a) Duration-Dependent Private Costs

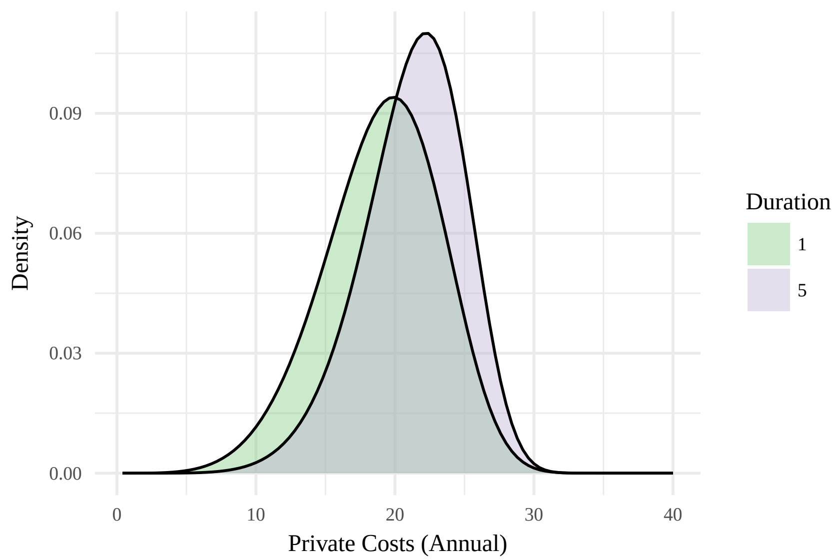

(b) Unobservable Auction-Specific Heterogeneity

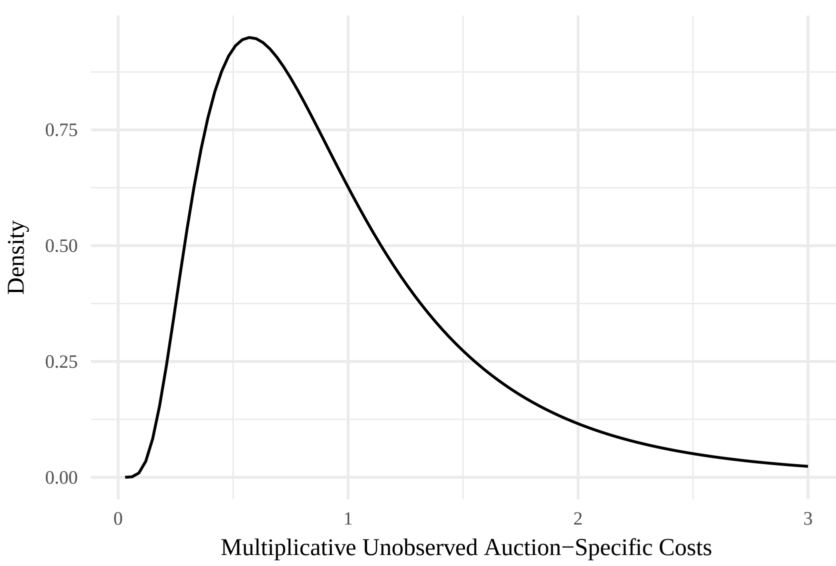

Notes: The figure plots the distributions of the unobservable components of bidder costs. Private costs are displayed in panel (a), and the density of unobserved auction-specific heterogeneity is displayed in panel (b). In panel (a), the density is plotted for a one-year contract and a five-year contract. The estimated parameters indicate an increasing mean and a decreasing variance in private costs with contract duration. The density shifts smoothly between these functions for intermediate values of duration.

G.3 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 building cleaning 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 15: Descriptive Regressions: Incumbency Check - Price

	IV-1 (a)	IV-1 (b)	IV-1 (c)	IV-2 (a)	IV-2 (b)	IV-2 (c)
Number of Bids	-0.053^{**}	-0.052^{**}	-0.052^{**}	-0.047^{**}	-0.046^{**}	-0.046^{**}
	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Duration (Years)	0.043^{***}	0.043^{***}	0.043^{***}	0.033^{**}	0.033^{**}	0.033^{**}
	(0.016)	(0.016)	(0.016)	(0.015)	(0.015)	(0.015)
ln(Square Footage)	0.689^{***}	0.688^{***}	0.688^{***}	0.687^{***}	0.686^{***}	0.686^{***}
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
ln(Weekly Frequency)	0.467^{***}	0.467^{***}	0.467^{***}	0.407^{***}	0.407^{***}	0.407^{***}
	(0.041)	(0.041)	(0.041)	(0.040)	(0.040)	(0.040)
ln(2004 Unemp.)	0.080^{***}	0.080^{***}	0.080^{***}	0.060^{***}	0.060^{***}	0.060^{***}
	(0.019)	(0.019)	(0.019)	(0.018)	(0.018)	(0.018)
High-Intensity Cleaning	0.559^{***}	0.559^{***}	0.559^{***}	-0.076	-0.076	-0.077
	(0.075)	(0.075)	(0.075)	(0.125)	(0.125)	(0.125)
Follow-On Contract		0.018			-0.003	
		(0.053)			(0.050)	
Incumbent Winner			0.012			-0.008
			(0.106)			(0.100)
Site Type FEs				X	X	X
Observations	1046	1046	1046	1046	1046	1046
R^2	0.69	0.69	0.69	0.73	0.73	0.73

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table displays regression results for regressions of log annual price on auction characteristics and local market characteristics. Specifications IV-1 (a) and IV-2 (a) are two-stage least squares regressions and are identical to the descriptive regressions in Table 4. The (b) specifications include an additional regressor indicating whether the contract is a follow-on contract and the (c) specifications include an indicator for whether the contract was won by an incumbent bidder in a follow-on contract.

Table 16: Descriptive Regressions: Incumbency Check - Number of Bids

	(1)	(2)	(3)	(4)	(5)	(6)
Duration (Years)	−0.002	−0.005	−0.009	−0.002	−0.005	−0.009
	(0.099)	(0.099)	(0.099)	(0.100)	(0.100)	(0.100)
ln(Square Footage)	0.834***	0.835***	0.840***	0.825***	0.824***	0.829***
	(0.106)	(0.106)	(0.106)	(0.112)	(0.112)	(0.112)
ln(Weekly Frequency)	0.009	0.014	0.010	0.137	0.146	0.141
	(0.253)	(0.253)	(0.253)	(0.257)	(0.258)	(0.257)
ln(2004 Unemp.)	−0.794***	−0.809***	−0.813***	−0.793***	−0.808***	−0.811***
	(0.238)	(0.239)	(0.239)	(0.238)	(0.238)	(0.238)
ln(Unemployment)	1.420***	1.432***	1.436***	1.356***	1.366***	1.370***
	(0.231)	(0.231)	(0.231)	(0.231)	(0.231)	(0.231)
ln(Num. Firms in Zip3)	0.257*	0.248*	0.250*	0.276*	0.267*	0.269*
	(0.148)	(0.148)	(0.148)	(0.147)	(0.147)	(0.147)
Generic Set-Aside	1.134***	1.125***	1.131***	0.987***	0.982***	0.985***
	(0.350)	(0.350)	(0.350)	(0.361)	(0.361)	(0.361)
High-Intensity Cleaning	−0.294	−0.303	−0.305			
	(0.475)	(0.475)	(0.475)			
Follow-On Contract		−0.351			−0.353	
		(0.326)			(0.326)	
Incumbent Winner			−0.836			−0.814
			(0.650)			(0.646)
Site Type FEs				X	X	X
Observations	1046	1046	1046	1046	1046	1046
R^2	0.17	0.17	0.17	0.19	0.19	0.19

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table displays regression results for regressions of number of bids on auction characteristics and local labor market variables. Specifications (1) and (4) are equivalent to the descriptive regressions (3) and (4) in Table 5. The additional specifications included indicators for whether the contract is a follow-on contract or whether the contract was won by an incumbent bidder in a follow-on contract.

G.4 Detailed Impacts of Standardized Duration

Table 17: Percent Impact of Uniform Term Policies

$\bar{T}$	Affected	Price	Trans. Cost	Total Cost	Count
1	All	−11.2	317.1	33.2	1046
	$T > \bar{T}$	−11.8	334.2	34.9	995
	$T < \bar{T}$	1.5	−36.6	0.6	23
2	All	−7.5	108.5	8.8	1046
	$T > \bar{T}$	−8.7	125.4	9.8	930
	$T < \bar{T}$	4.1	−50.5	2.3	62
3	All	−3.9	39.0	2.8	1046
	$T > \bar{T}$	−6.3	61.8	3.3	761
	$T < \bar{T}$	5.1	−42.6	2.7	146
4	All	−0.4	4.3	1.3	1046
	$T > \bar{T}$	−3.2	24.0	0.7	686
	$T < \bar{T}$	6.0	−39.1	2.9	306
5	All	3.1	−16.6	1.5	1046
	$T > \bar{T}$	−2.0	12.2	0.3	18
	$T < \bar{T}$	6.9	−36.8	3.4	478

Notes: The table displays the average percent changes (by contract, not in aggregate) in total costs, prices, and annualized transaction costs when all contracts are issued in standardized durations corresponding to $\bar{T}$. For a uniform duration policy of 4 years or less, the average price paid decreases and the amount spent on transaction costs increases. The final column lists the count of the affected contracts. The first column indicates the group affected by the policy. Rows corresponding to $T > \bar{T}$ pertain to all contracts that see a reduction in duration, and the reported effects are equivalent to a policy that caps duration at $\bar{T}$.

H Contract Documents

The following page is an example first page from a building cleaning 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.