Conservation Versus Competition? Environmental Objectives in Government Contracting

Evan Herrnstadt *

November 24, 2015

Abstract

Government contracts increasingly incorporate environmental objectives or preferences for sustainable products. I show how such environmental restrictions can influence contract bids through firm costs, strategic bidding behavior, and bidder participation decisions. Using data from Michigan timber auctions, I find that conservation restrictions impose compliance costs of up to 15% of the winning bid. These costs are found to be fully borne by the state. Loggers capture a larger share of auction surplus for restricted contracts, indicating that the policy blunts competitive pressure in the auctions. Optimal reserve prices can partially mitigate the revenue disparity between more- and less-restricted contracts.

*Harvard University Center for the Environment, 24 Oxford St., Cambridge, MA 02139. Email: **eherrnstadt@gmail.com**. Thanks to Ryan Kellogg, Jeremy Fox, Catherine Hausman, and Erich Muehlegger for consistent support and feedback. Thanks also to Dan Ackerberg, Varanya Chaubey, Aaron Flaaen, Joe Golden, Jason Kerwin, Eric Lewis, Karen Stockley, Richard Sweeney, Ophira Vishkin, James Wang, and seminar participants at the University of Michigan, Analysis Group, Charles River Associates, Resources for the Future, the University of Delaware, and the University of Massachusetts for valuable comments; to Mark Anderson, Donovan Asselin, and Doug Heym at the Michigan DNR; and to Phillip Minerick for his insights into the bidding and logging process. The analysis was run using the University of Michigan Flux High Performance Computing Cluster. Funding from the Joseph L. Fisher Dissertation Fellowship from Resources for the Future and from the Michigan Institute for Teaching and Research in Economics is gratefully acknowledged.

1 Introduction

Increasingly, governments are leveraging the scale and scope of their contracts with private firms to reduce the environmental impact of large projects and purchases. At the same time, contracts issued by the government are often awarded to firms through competitive bidding to increase revenues (in the case of a sale) or reduce costs (in the case of a purchase). While the pursuit of environmental objectives will likely impose additional costs on the firms, the economic incidence of these costs will be determined by the intensity of competitive pressure among bidders. Indeed, previous work has recognized that the pursuit of social objectives in government auctions, such as small-business preferences or bids based on estimated contract completion time, can distort competition and affect government revenues, bidder surplus, and efficiency.[1] Unlike these policies that typically modify the rules of the allocation mechanism, environmental objectives often affect the value of the contract itself. However, such objectives can still undermine or bolster competition.

In this paper, I estimate the effect of environmental objectives on competitive pressure in auctions for natural resource extraction contracts. In particular, I analyze competition for timber contracts auctioned by the State of Michigan Department of Natural Resources (DNR) in the presence of varying seasonal operating restrictions. The restrictions are implemented to mitigate the impacts of logging in the state forests on the surrounding ecology and recreational use.

A reduced-form analysis demonstrates that environmental objectives have a negative effect on bids and logger auction participation and suggests that differential competition could play an important role.[2] Specifically, I estimate that the winning bid is 17 percent lower for the most-restricted timber contracts and that these contracts receive 35 percent fewer bids above the reserve price. However, these effects are highly nonlinear; if the contract is restricted for fewer than 5 months, the bids are not significantly different from bids for the unrestricted contracts. I exploit the structure of the policy and find that the estimates are robust to controlling for the reasons underlying the restrictions on a given contract; this mitigates concerns about omitted variable bias. Finally, controlling for the number of participating bidders accounts for half of the effect of restrictions on the winning bid, which suggests that the participation margin matters, but is not driving the entire effect

[1] Examples of this literature include Marion (2007); Krasnokutskaya and Seim (2011); Athey, Coey and Levin (2013); and Bajari and Lewis (2011).

[2] Throughout the paper, I refer to all potential bidders as "loggers" for narrative simplicity. In several other papers that model timber auctions (Athey, Levin and Seira, 2011; Roberts and Sweeting, forthcoming), the authors distinguish between loggers and mills. Conversations with DNR employees suggest that there are few large-scale mills operating in this market.

of the restrictions.

Although the reduced-form analysis of equilibrium bids estimates the effect of restrictions on government revenue, it cannot generally reveal how costly the restrictions are or who bears the economic burden of these costs. If the restrictions cause logger valuations for a contract to become more dispersed, loggers will be more insulated from competition and the winning bidder's equilibrium surplus will increase. In contrast, if the restrictions compress the distribution of valuations, loggers will expect more intense competition and the winning bidder's equilibrium surplus will decrease.

To disentangle these effects, I specify a model of the DNR's first-price auctions and analyze the three channels through which restrictions result in lower bids. First, the bidders' values could be lower, directly resulting in reduced bids. Second, the bidders would further depress their bids if they face less competition locally because of increased dispersion in private values. Third, firms may change their decision rule for when to participate in the auction at all.

I structurally estimate the auction model and find that compliance with stringent environmental objectives is costly. These costs are almost completely borne by the government. Compliance costs are very close to zero for contracts that are restricted for less than 4 months. However, restrictions covering 10 months of the year create compliance costs amounting to 15 percent of the government revenue or 54 percent of the firm surplus from an unrestricted sale. Even when the sale is restricted for only 6 months of the year, the compliance costs amount to 5 percent of government revenue or 17 percent of firm surplus. I find that loggers are able to depress their bids enough to fully pass through the compliance costs to the state. The change in average firm surplus is precisely estimated and very close to zero for most levels of restrictions.

The full pass-through finding is driven by two factors. First, I assume that the DNR is perfectly inelastic to expected bids in supplying timber contracts. This assumption is supported by the timing of and institutional criteria driving the timber harvest process. Second, my estimates suggest that compliance costs do not affect the dispersion of contract values. Thus, firms face a similar "local" distribution of opponents. These mechanisms are related to those discussed by Fabra and Reguant (2014), who estimate full pass-through of carbon permit costs in the wholesale Spanish electricity market.

A different way to evaluate the effect of restrictions on auction competition is to calculate the share of total surplus (government revenue plus firm surplus) captured by loggers. I find that for an average contract, loggers capture a 5.2 to 6.6 percent larger share of surplus for more restricted sales relative to unrestricted sales; the difference is statistically significantly different from zero. This indicates that the restrictions do undermine the competitive performance of the timber auctions, even though the *level* of firm surplus falls slightly.

To better understand the relative importance of various mechanisms, I decompose the effect of restrictions on bids. Holding bidding strategies fixed for a logger with a given valuation, I find that lower valuations due to restrictions directly account for 73 to 82 percent of the decrease in bids. Allowing loggers to revise their bidding strategies and participation decisions in response to their opponents' now-lower valuations accounts for the remaining 18 to 27 percent. This decomposition shows that the change in winning bids reflects a substantial adjustment by loggers to the expected compliance costs of their competitors.

Heavier restrictions reduce the revenue captured by the DNR. However, I find that optimal reserve prices that account for the restrictions can close some of the revenue gap between more- and less-restricted sales. Further, optimal reserve prices can recover forgone revenues without any changes to conservation policy: a heavily restricted sale paired with an optimal reserve price receives similar expected revenue to a lightly-restricted sale using the observed reserve price.

Although these results apply to a particular policy, seasonally-differentiated regulations are used in a variety of settings. For example, drilling for oil and gas on government-issued leases is seasonally restricted, both onshore and offshore, for a variety of environmental reasons. Past and existing ozone regulations have often had seasonal components, such as the NOx Budget Program/State Implementation Plan call and various gasoline content requirements. Finally, seasonal restrictions also arise frequently in the context of regulated fisheries to prevent adverse effects on non-target species.

Furthermore, environmental objectives are increasingly embedded in a wide variety of government contracting settings. While the Competition in Contracting Act of 1984 sets out conditions and exceptions regarding free and open bidding on federal contracts, various statutes allow the federal government to relax the Act's requirements to pursue goals related to the environment and sustainability. For instance, the Obama Administration has issued a series of executive orders (e.g., 13423, 13514, and 13693) and that promote consideration of environmental factors in federal procurement. Recently, the General Services Administration (GSA) included greenhouse gas emissions reporting and reduction strategy requirements in its most recent Domestic Delivery Services contract with FedEx and UPS (U.S. General Services Administration, 2015). Finally, the planning and completion of contracted projects (e.g., highway construction) may be subject to broader environmental regulation, such as the National Environmental Protection Act, the Clean Air Act, and the Clean Water Act.

Governments should be aware that environmental objectives can distort the competitive structure of the contracting process, rendering simple policy predictions and evaluations inaccurate. Strategic firm responses can be an important consideration when evaluating the impacts of various environmental policies (Busse and Keohane, 2007; Brown, Hastings, Mansur and Villas-Boas, 2008; Ryan, 2012). An *ex ante* prediction of future bids based on estimated compliance costs assumes exact one-to-one pass-through, which does not have to be the case. Conversely, an *ex post* regulatory cost calculation from a simple comparison of bids with and without the environmental policy ignores the potential impact of changing firm margins. An understanding of the competitiveness of the market is crucial for an accurate evaluation.

I organize the remainder of the paper as follows: In Section 2, I describe the market for logging contracts in Michigan and outline the role of seasonal operating restrictions. In Section 3, I describe the contract data and discuss my measure of seasonal restrictions. In Section 4, I establish reduced-form effects of the restrictions on equilibrium bid outcomes. In Section 5, I develop and explain the implementation of the structural model. In Section 6, I present the structural parameter estimates, discuss the magnitude of compliance costs and the incidence of the seasonal restrictions, consider an optimal reserve price policy that accounts for the restrictions, and decompose the reduced-form effect to shed light on the importance of various mechanisms at play. Section 7 concludes the paper.

2 Policy and Empirical Setting

To provide context for the empirical analysis, I describe the widespread inclusion of environmental objectives in government contracting. I also outline my specific context: Michigan DNR logging contract auctions. These contracts include seasonal operating restrictions that protect the ecological integrity of the forest and promote multiple uses, but may impose costs on loggers by reducing scheduling flexibility.

2.1 Environmental Objectives and Fostering Competition

Governments rely heavily on goods and services outsourced from private firms. When they contract with such firms, there is an information asymmetry: the firms have better information about their own productivity levels, costs, or values for the contract. The government often uses a competitive bidding process to extract this information; however, firms will still capture some information rents.

Although government contracting is generally carried out with a priority of fostering competition, environmental responsibility is one competing concern. In the federal context, the Competition in Contracting Act of 1984 states that contracts are to be awarded through "full and open competition", with potential exceptions for small business set-asides, an urgent and compelling need, a service with a sole supplier, small purchases, or other reasons authorized in statute. Environmental preferences and objectives are justified under a number of statutes across a wide variety of contracting settings. While the competitive impact of the small business exception has been well-studied (Marion, 2007; Krasnokutskaya and Seim, 2011; Athey et al., 2013), environmental objectives have not.[3]

Such environmental objectives are becoming more pervasive: many federal government agencies have established broad "Green Procurement Programs" to comply with a variety of relevant executive orders and congressional acts (Manuel and Halchin, 2013). Some agencies have expressed concerns that these practices will considerably shorten the list of acceptable contractors or products (United States Department of Defense, 2008). An important, but not easily measured, component of evaluating these programs is whether they affect the competitive performance of the bidding process in terms of the division of surplus. I analyze conservation and multiuse requirements in Michigan state forest logging contracts to illustrate the possible effects of environmental objectives on competition for government contracts.

2.2 Timber Contracts and Seasonal Restrictions

The Michigan DNR is mandated with maintaining the ecological integrity and promoting the recreational use of the state forests, while supporting the timber and timber products industry by auctioning logging contracts.[4] These logging contracts often include clauses that disallow operations during certain times of the year during which the forests are ecologically-sensitive or subject to high recreational demand. The restrictions are known prior to the competitive bidding process. Loggers claim that these restrictions can be quite costly to their operations and affect their bids.

The DNR tracks the condition of the Michigan state forest system on an ongoing basis. Foresters survey each forest compartment (roughly 2000 acres) every 10 years. This

[3] Aral, Beil and Wassenhove (2014) theoretically analyze a company that decides whether to audit possible suppliers for sustainable practices prior to a private procurement auction. Smith, von Haefen and Zhu (1999) compare the cost per mile of highway construction in states with a higher or lower likelihood of triggering federal environmental and cultural preservation review requirements.

[4] This mandate is similar in spirit to the federal Multiple Use-Sustained Yield Act governing the mission of the U.S. Forest Service.

survey includes information about the basic mix, density, and health of the compartment to be used in a statewide timber inventory. Each year, the foresters determine which stands of trees will be contracted for harvest using a combination of inventory and aerial data. According to conversations with DNR officials, the timber is chosen for harvest to pursue forest-management goals. That is, trees are harvested to maintain proper age balance, density, and disease and pest resistance. Once a stand of trees is selected for commercial harvest, the DNR sends a forester out to the stand to obtain more precise measurements of the timber to be harvested. In the process, the forester may determine that there are grounds for seasonal operating restrictions. For instance, if the ground is particularly wet in the summer, operations may not be allowed during that time of year to prevent damage to the forest's root structure.

Once the survey is completed, the DNR holds an auction for the obligation to harvest the timber. A contract is made public, including any seasonal restrictions. There is usually a 4-6 week bidding period before the bid opening date. During the interim, loggers often conduct a "cruise" of the sale to get a first-hand look at the area in which the harvest will take place. The auctions are sealed-bid first-price auctions with public reserve prices.[5] The bids, bidder identities, and number of bids submitted are considered confidential until the results are made fully public at the bid opening. The highest bidder wins the contract, pays a down payment, and is obligated to harvest the specified timber before a contract deadline. Failure to fulfill the contract terms results in a financial penalty and possible exclusion from future sales.[6]

In the DNR auctions, loggers submit a lump-sum bid that is not a function of eventual timber sales. This is an important institutional feature that simplifies the analysis. In some other timber sale settings, the state uses so-called "scale auctions". In a scale auction, each logger submits a per-volume price for each tree species. These prices are multiplied by the vector of predicted species-specific harvests to calculate a logger's total bid and determine the allocation of the contract. However, the logger's eventual payment to the state is determined by the *realized* volume of timber marketed. Athey and Levin (2001) show that this structure creates incentives to skew one's bid based on private information about the relative prevalence of different species. In contrast, the Michigan DNR's sales are standard first-price auctions, and do not create incentives to skew one's bid.

Seasonal operating restrictions are added to timber contracts to help protect the eco-

[5]Reserve prices are a function of past sale prices and the DNR forester-determined cost index, which is described in Subsection 4.1.

[6]Further, fewer than 1 percent of contracts are transferred between firms.

logical integrity and recreational accessibility of the state forests while they are harvested. To this end, the contracts will often specify certain dates during which the loggers cannot operate on the sale. There are a number of reasons that a sale might be restricted in such a way; Table 1 provides the frequency with which the main reasons are cited. Many of these restrictions are related to environmental conservation and resource management. For instance, many sales are restricted in the spring/summer due to "bark slip". From April through July, tree bark tends to loosen from the trunk. Thus, it is easy to damage trees when cutting and hauling nearby timber. An example of such a contract clause is displayed in Figure 1. Another example is the presence of an endangered bird, which would require operations to cease during nesting season. There are also restrictions related to the multiuse mandate of the state forest system: areas with popular snowmobile trails are sometimes restricted during winter months, while an area with a large deer population might be restricted during hunting season.

There is some existing empirical evidence that such restrictions influence a logger's bidding decision for a given contract. Using data from Minnesota state forest auctions, Brown, Kilgore, Coggins and Blinn (2012) find that sales that allow harvesting activity during the summer or fall garner winning bids that are 7 percent higher. Taking a different approach, Brown, Kilgore, Coggins, Blinn and Pfender (2010) surveyed loggers and DNR foresters in Michigan, Minnesota, and Wisconsin. Loggers cited seasonal restrictions as the most important factor for determining their bids, aside from the volume and type of timber included in the contract.

Conversations with Michigan loggers and DNR foresters suggest that these restrictions are costly primarily because they impose scheduling constraints. Loggers attempt to keep their equipment running year-round for three main reasons. First, many loggers have quotas and contracts with sawmills and pulpmills that they need to meet at some frequency. Second, logging can be quite capital-intensive, and consistent revenues are needed to stay up-to-date on loan payments. Third, loggers simply want to provide consistent employment for their workers. This desire to schedule jobs throughout the year leads to a difficult scheduling problem.[7] The scheduling problem becomes more complicated when the sales are seasonally restricted. Essentially, a restricted sale embodies less option value than one that can be cut at any time of year.

These types of restrictions would be less costly if there was a well-functioning short-term equipment rental market. However, a survey of loggers located in the Eastern half of

[7] One DNR employee likened the scheduling problem to "the worst linear programming problem [he] can imagine."

the Upper Peninsula and the Northern Lower Peninsula suggests that rental activity is limited (Srivastava, Abbas, Pan and Saffron, 2011). In 2009, firms used self-owned equipment for an average of 89 percent of their total operations. An average of 19 percent of operations was performed using subcontracted equipment, but this question was only answered by about half of the respondents and the minimum response was 1 percent. Assuming that the non-responding firms did not subcontract at all, 10 percent of total operations used some subcontracted equipment. Although the share is non-negligible, it is small. Furthermore, the cost of the restrictions is likely related to unexpected shocks. Such short-run rentals would be even more difficult to arrange.

3 Features of the DNR Contract Data

In this section, I outline the key outcomes and covariates from the contract data. I also construct a measure of seasonal restrictions; there is considerable variation in the number of months for which a contract is restricted. I will use this variation to estimate a flexible relationship between restriction intensity and bidding behavior, private values, and participation costs.

3.1 Contract Characteristics and Auction Outcomes

I obtained the contract text and auction outcomes for all Michigan state commercial timber sales from April 2004 - March 2013. The data include extensive information about the contract and auction outcomes, such as all bids, bidder identities, reserve prices, DNR volume estimates of each product-species combination in the sale, acreage, DNR cost factor estimates, and precise sale location. To scale bids and reserve prices in a way that makes sales more comparable, I re-express bids and reserve prices in dollars per thousand board feet (MBF).[8] Reserve prices are set using a formula based on recent prices paid for the same species in the same state forest.

Table 2 presents a summary of the sample auctions. Of the 5207 sample auctions, 457 receive zero bids. Conditional on receiving at least one bid, the mean sale receives a winning bid of \$92.2/MBF; in total dollar terms, the DNR earns \$66,000 in revenue from the average contract transaction.[9] The mean reserve price is \$61.6/MBF, or roughly

[8]For reference, 1 MBF of lumber would be a stack of boards that is 10 feet long, 4 feet wide, and just over 2 feet tall. To convert pulpwood, which is measured in cords, to MBF, I use a conversion rate of 2 cords per MBF (Mackes, 2004). I include controls for the composition of the sale in all specifications.

[9]All dollar figures are deflated to 2009 USD.

$43,000. The median number of bidders is 3, and the mean is 3.9, reflecting a long right tail (the maximum number of bidders in a single auction is 19). I measure potential bidders by identifying the set of loggers that are active in similar auctions. Specifically, I define a potential bidder to be any logger who bids for a state forest contract in the same calendar quarter and DNR management unit as the contract of interest.[10] This definition seems reasonable: 87 percent of bidders in a given auction bid in at least one other state timber auction in the same calendar quarter-management unit. There are a mean of 18 potential bidders for the contracts, and the participation rate in a typical auction is approximately 20 percent.

The value of a contract will vary based on the type of timber required to be harvested and the attributes of the harvest site itself. The average sale covers 90 acres (roughly one-eighth of a square mile) and includes an estimated 687 MBF. The cost factor variable captures attributes such as wetness, slope, and distance to a road. It is generated by the forester appraising the sale on the ground, and is used to help inform the appraisal/reserve price.

I restrict the sample slightly to exclude especially unusual sales. I exclude sales with reserve prices less than $20/MBF or greater than $250/MBF or areas less than 20 acres or greater than 640 acres. These values are roughly the 1st and 99th percentiles of these variables. I also drop all salvage sales, which specifically market fire-, wind-, or pest-damaged timber.

3.2 Seasonal Restrictions

The DNR does not maintain a database variable that indicates when harvesting operations are allowed. However, the information is explicitly written into the contract that describes the sale to potential bidders. Thus, I analyzed all relevant contract clauses and constructed such a variable. Specifically, I calculated the number of months that a sale is restricted.[11] This variable captures the first-order driver of lost option value: the number of months during which a sale is inaccessible.[12]

[10] A management unit, of which there are 14 total, is usually a 2 to 3 county area. This general approach is similar to existing work that analyzes entry in timber auctions, such as Roberts and Sweeting (forthcoming) and Athey et al. (2013).

[11] Sales are divided into "payment units", which may be subject to different restrictions. For each calendar month, I determine the fraction of the month that each payment unit is restricted. Then I calculate the average across payment units, weighting them by appraisal value.

[12] One concern with this measure is that if a contract takes a few weeks to fulfill, then a short window of availability is essentially a restriction. Fewer than 1 percent of contracts have any windows between restrictions that last for 15 days or less. Treating these windows as restrictions or omitting such sales from the analysis entirely has no effect on the results. In Appendix A, I also consider seasonality as a possible mechanism.

There is considerable variation in the average number of months for which a sale is restricted. Among sales with any restrictions, the median is 3 months, which reflects that many sales are restricted for a single season. However, there are 536 contracts (10.2 percent of the full sample) that are restricted for 6 or more months. Figure 2 displays the conditional distribution of restrictions. The spike at the 2-3 month bin reflects that the most common restriction (bark slip) generally lasts between 2 and 3 months, from mid-April to mid-July.

One concern with using cross-sectional variation in restrictions to identify a treatment effect is that the distribution of covariates may be very different for more and less restricted sales. Table 3 compares the composition, size, number of potential bidders, and cost factors of more- and less-restricted contracts. More restricted contracts are slightly smaller on average and more likely to be located in the Upper Peninsula. Overall, sales seem fairly similar in terms of observables whether heavily or less heavily restricted.

4 The Effect of Restrictions on Equilibrium Bidding

In this section, I show that seasonal restrictions have a negative effect on equilibrium bidding and participation. First, I estimate the effect of seasonal restrictions on winning bids and the number of bidders participating in an auction, controlling for a rich vector of auction characteristics. Second, I demonstrate robustness of this base specification to omitted variables by exploiting the structure of the seasonal restrictions. Third, I find that the effect on bids is correlated with, but not fully explained by, changes in bidder participation.

4.1 Main Specification

I estimate a large, negative effect of restrictions on bids and participation. When the restrictions are allowed to enter nonlinearly into the regression, I find that the effects are concentrated among the more-restricted sales.

The main specification is given by:

$$Outcome_a = h(MonthsRestricted_a) + \beta X_a + \varepsilon_a$$

where $Outcome_a$ is the number of bidders or the logarithm of the winning bid per MBF in auction a, $h(\cdot)$ is a function of the number of months for which a contract is restricted, and X_a is a vector of controls. These controls contain standard characteristics used in previous work on timber auctions, such as the Herfindal-Hirschman Index of the value of the species in the sale, the size of the sale in acres, and the mix of sawlogs (lumber) versus pulpwood (a paper input).[13] A particularly important and new control variable is a cost index developed by the DNR. This index is meant to capture otherwise difficult-to-capture characteristics, such as the topography of the land, the soil conditions, road and construction requirements, the distance to the nearest road and mill, and an assessment of the timber quality. Importantly, seasonal restrictions are not directly accounted for in these cost factors.[14]

Although seasonal restrictions may be correlated with other determinants of a contract's value, I address much of the omitted variable problem with a comprehensive set of controls and proxies. For example, if a wet stand of timber is more likely to be restricted to preserve the root structure of the stand, but loggers also find working in wet areas more costly due to higher equipment maintenance costs, this would introduce negative bias into the treatment effect. Most of these concerns can be eliminated by controlling for observable auction characteristics. The DNR-calculated cost index is a particularly crucial control variable for this reason. In Section 4.2, I also leverage the structure of the restrictions. In particular, contracts are frequently restricted for multiple reasons, which allows me to control for the underlying basis for the restrictions.

When I specify $h(\cdot)$ as a linear function of the months of restrictions, I find that restrictions have a significant negative effect on the logarithm of winning bids and the number of bidders (see Tables 4 and 5, respectively).[15] The effect of an additional month of restriction is quite robust across different sets of controls and location, quarter-of-year, and year fixed effects. I focus on Column 4 as my preferred specification for both the reduced-form and structural estimates. This specification indicates that that the most-restricted contracts attract winning bids that are 8 percent lower and 0.8 fewer bidders out of an average of 3.9.

Although the linear functional form implies that the additional effect of each month of restrictions is the same, estimates from a more flexible specification suggest that the state should be primarily concerned about losing revenues due to the most stringent restrictions. When I specify $h(\cdot)$ as a restricted cubic spline with knots at 0, 3, 6, and 10 months of restrictions, the effects are strikingly different from the linear effect.[16] Panels A and B of Figure 3 present the cumulative effect of restrictions as the number of months increases from zero to 10 (the maximum in the sample) on the logarithm of the winning bid and the number of bidders, respectively.[17] The first four months of restrictions have zero marginal

[13] The full contents of the control vector can be seen in Table 4.

[14] See Table A4 for a detailed list of the criteria used in developing the cost factors. Note that this variable is defined such that a larger value corresponds to a *less* costly sale.

[15] As contracts in the same area around the same time are likely to be subject to similar shocks, I calculate clustered standard errors at the county-by-year level.

[16] The results are robust to other similar sets of knots.

[17] The underlying regressions are analogous to Column (4) of Tables 4 and 5.

effect on winning bids, while the average marginal effect over restriction months 5-10 is about 4 percent per month. In contrast, the linear specification implies that each additional month of restrictions is associated with a 0.8 percent decline in the winning bid. For a contract that is restricted for 10 months, the effect is quite large: it receives 1.5 fewer bids (the mean is 3.9) and a winning bid that is 17 percent lower relative to an unrestricted contract.[18]

4.2 Robustness: Identification from Multiple Restriction Types

To this point, the identification assumption has been that, conditional on control variables, the restrictions only affect values directly through the scheduling constraint and are also uncorrelated with other unobserved determinants of bids. Given that the restrictions do not require specific procedures during the time that the sale is accessible, this seems reasonable. To further address potential omitted variable bias, I exploit the fact that a single contract could be restricted for multiple unrelated reasons. Specifically, I control for the rationales behind the restrictions with a set of dummy variables.[19] The new weaker identification assumption is that the *interactions* between restriction categories do not directly affect logger valuations and are uncorrelated with other unobserved determinants of bids, conditional on controls. Indeed, conversations with DNR foresters and industry participants suggest that these interaction effects are zero or at most a second-order consideration.

Because multiple regulation types "stack" on top of each other, I can include restriction-type dummies to control for restriction-specific unobservables, leaving only idiosyncratic variation and the (potential) effect of interactions between restriction types. The regression equation is now:

$$Outcome_a = h(MonthsRestricted_a) + \beta X_a + \sum \gamma_a^r I_a^r + \varepsilon_a$$

where I_a^r is an indicator variable equal to one if contract a is restricted for reason r.

In practice, this identification strategy requires that combinations of different restriction categories exist in the data, which is satisfied in this context. There are 32 unique pairwise combinations of restrictions, and 28 percent of the contracts in the sample are restricted for

[18] The effect on the number of bidders suggests that the differences in winning bids might be driven by differences in market thickness. However, I find that participation rate, not market thickness drives the results (see Appendix E for details).

[19] One possible alternative identification strategy would be to exploit the arbitrary assignment of foresters to different sales and use DNR foresters' idiosyncratic tendencies as an instrumental variable. This approach is discussed in Appendix B. Unfortunately, there is not sufficient variation in this instrument to identify the treatment effect.

at least two reasons.[20] This abundance of combined restrictions should allow me to reliably apply the identification strategy.

If unobserved factors correlated with individual restriction types are driving the equilibrium treatment effects, these estimates should shrink toward zero when I control for restriction type. Of course, if the restriction categories are positively correlated with *valuable* unobserved contract characteristics, then the estimates would be larger in magnitude. The linear specifications are presented in Table 6: the treatment effect does increase slightly. In Panel A of Figure 4, I estimate a spline specification with the restriction categories. The magnitude of the treatment effect actually increases a small amount: for the most restricted sales, the effect on winning bids is 19 percent, compared with 17 percent in Panel A of Figure 3. The effect on the number of bidders in Panel B of Figure 4 is also slightly different from the base spline specification in Panel B of Figure 3. In both cases, the difference is well within the 95 percent confidence interval, and I take this as evidence that the reduced-form estimates are robust to omitted variable bias.

4.3 Importance of the Participation Margin

Given the significant effect of restrictions on the number of bidders, I re-estimate the effect of restrictions on the winning bid, but control for the number of participating bidders. In the presence of a binding reserve, a change in the unobserved distribution of values will directly supress participation because fewer bidders will draw values above the threshold necessary to justify bidding. Figure 5 presents the estimated restriction spline: although there is still a significant negative impact, accounting for the number of bidders accounts for roughly half of the restriction treatment effect. Column 6 of Table 4 adds a vector of dummy variables for the number of bids received to the main linear regression specification (i.e., Column 4). Including this measure of participation attenuates the treatment effect by half. These results underline the importance of directly modeling the reserve price and estimating participation costs.

5 A Structural Model of DNR Timber Auctions

The reduced-form effects establish that bids are affected by seasonal restrictions; however, such an approach cannot recover compliance costs, surplus, and incidence of the costs. To that end, I specify a structural model of a first-price auction with costly participation based

[20] See Appendix E for the full matrix of combinations.

on Samuelson (1985). I describe the various channels through which restrictions could affect bidding behavior, parameterize the model such that these channels can be estimated, and outline the actual estimation procedure. The model allows me to estimate the extent to which the equilibrium bid effects are driven by lower valuations versus weakened competition.

5.1 Main Assumptions

To introduce the basic components of the structural model, I specify a model of a first-price auction with endogenous participation and characterize the equilibrium.

The model is a first-price auction with costly participation; the equilibrium is a participation rule combined with a bid function. There are N ex ante identical potential bidders that draw independent private values v_i from a common distribution $F(v)$ with support on $[0, \bar{v}]$. Each bidder observes a private realization of this draw and decides whether to undertake a bid-preparation process, which costs K. Participants then submit bids in a first-price auction with public reserve price R, without first observing the other potential bidders' participation decisions. I restrict my analysis to symmetric perfect Bayesian Nash Equilibria. Given my assumptions, the equilibrium is characterized by a cutoff type $v^*(N,R)$ and equilibrium bidding function $b(v;N,R)$.[21] That is, a potential bidder with valuation v will incur the bid preparation cost and submit a bid $b(v)$ if and only if $v \geq v^*$.

A key informational assumption is that bidders learn their valuations before making the participation decision. This information structure (Samuelson, 1985) implies that the types entering the auction will represent draws from an advantageously selected portion of the value distribution. The main alternative in the literature is a model with no selection (Levin and Smith, 1994), in which firms know only the distribution $F(v)$ when they pay their entry cost. In terms of entry, a marginal firm and an inframarginal firm draw their private values from the same distribution.

I choose the selective entry model for two reasons. First, in my setting loggers tend to bid only on nearby tracts of timber and have often been working in the same small area for years. This suggests that firms probably have a fairly precise signal about their private value prior to incurring any sunk cost. Second, the selective entry model is preferred by Li and Zheng (2012), who formally test the selective and non-selective entry models against one another using Michigan DNR timber auctions and find that that the selective entry model is a much better fit for the data.[22]

[21] I suppress N and R going forward to simplify notation.

[22] A third option is the affiliated signal model introduced empirically by Roberts and Sweeting (forthcoming).

The choice of entry model has important consequences for the model's implications, and the validity of the structural estimates.[23] The difference between the models can be understood by considering a policy that subsidizes entry. In expectation, this policy will induce some marginal firms to bid that would not have otherwise done so. In the selective entry model, these marginal firms will have lower private values than those that would have entered without the subsidy. In contrast, the non-selective entry model implies that the marginal entrant will have the average value of the existing participants in expectation.

The entry model will also affect the structural estimation of the private value distribution. If the non-selective entry model is estimated and there is actually selection in the entry process, firm value estimates will be too high and underdispersed because the bids are assumed to be representative draws of the unconditional (on entry) value distribution. This could lead to misleading estimates of firm surplus.

The model does not allow for dynamic considerations, such as contract backlog. Most studies that have estimated dynamic procurement auctions have done so in the context of highway construction (Jofre-Bonet and Pesendorfer, 2003; Balat, 2013; Groeger, 2014). In their setting, constructing a backlog measure is reasonable: most comparable jobs are observed as state or federal projects and the contracts must generally be completed by the end of the year. In my setting, state forest contracts represent only a quarter of total timber cut in Michigan; private and federal forestland compose the balance. Firms located in the Upper Peninsula may also bid on jobs in Wisconsin. Further, DNR contracts last for 2 to 3 years and I am unable to obtain the true completion date. Thus, any inventory measure I could construct based solely on state forest auctions would be uninformative.

5.2 Potential Effects of Seasonal Restrictions on Bids

I present a closed-form expression for the equilibrium described in the previous subsection and describe in detail the three channels through which restrictions could affect bidding: the value effect, the competition effect, and the participation threshold effect. I will quantify the relative importance of these three channels using the structural model. While I apply this model to high-bid first-price auctions in the Michigan timber market, the basic

In this model, firms receive a noisy signal of their value and decide whether to pay a cost to reveal their true valuation: the S and LS models are limiting cases. Roberts and Sweeting find that participation in U.S. Forest Service auctions is moderately (but not perfectly) selective.

[23] Sweeting and Bhattacharya (2015) provide an extensive comparison of the performance of various mechanisms under a variety of entry models. One of their Monte Carlo findings is that small deviations from the non-selective model of Levin and Smith (1994) can lead to large changes in policy prescriptions, while small deviations from the selective model of Samuelson (1985) lead to somewhat less extreme differences.

intuition can be extended directly to any contract allocated using an auction mechanism. The implications for identifying compliance costs and changes in firm surplus solely from equilibrium transaction prices will still apply as long as the expected firm information rents can be affected by the policy.

To simplify the explanation of the various channels, I start with a model of costless participation. In this model, there are N potential bidders, who will always bid if their valuation is above the reserve price, R. In this case, Holt (1980) and Riley and Samuelson (1981) derived a closed-form solution for the equilibrium bidding function, which I adapt into an expression for the expected winning bid:

$$E_{v^w}[b_i(v^w;F_{-i}(\cdot;r)] = \int_R^{\bar{v}} \left[v^w - \underbrace{\frac{\int_R^{v^w} F_{-i}(u;r)^{N-1}du}{F_{-i}(v^w;r)^{N-1}}}_{\text{markdown}} \right] f_N(v^w;r)dv^w$$

where $F_N(v;r)$ and $f_N(v;r))$ are the distribution and density, respectively, of the highest value draw (v^w) among the N bidders. $F_{-i}(v;r)$ is the distribution from which a bidder expects their competitors to draw. Note that $F_N(v;r)$ and $F_{-i}(v;r)$ are functions of seasonal restrictions r.

This equilibrium assumes bidder symmetry, i.e., that $F_{-i}(v;r)^N = F_N(v;r)$. However, I make the distinction between a bidder's own and opponents' distributions to allow a clear decomposition of the change in the expected winning bid with respect to restrictions. There are two main ways that a change in restrictions could change the equilibrium bid vector: the value effect and the competition effect. These are evident from the derivative with respect to the restrictions:

$$\frac{dE_{v^w}[b_i(v^w;F_{-i}(\cdot;r)]}{dr} = \overbrace{\underbrace{\int_R^{\bar{v}} \left[v^w \frac{d}{dr}[f_N(v^w;r)]dv^w \right]}_{\text{Compliance Cost}} - \int_R^{\bar{v}} \left[\frac{\int_R^{v^w} F_{-i}(u;r)^{N-1}du}{F_{-i}(v^w;r)^{N-1}} \frac{d}{dr}[f_N(v^w;r)]dv^w \right]}^{\text{Value Effect}}$$

$$\underbrace{- \int_R^{\bar{v}} \frac{d}{dr} \left[\frac{\int_R^{v^w} F_{-i}(u;r)^{N-1}du}{F_{-i}(v^w;r)^{N-1}} \right] f_N(v^w;r)dv^w}_{\text{Competition Effect}}$$

Value Effect A change in the distribution of the highest valuation, F_N, due to restrictions will affect the expected winning bid. Even without any change in the markdowns associated with a given valuation, the expected winning bid would be different. This difference is the mechanical effect of changing the mix of private values without allowing firms to re-optimize their bid functions accordingly.

This expression demonstrates that a simple comparison of bids cannot identify compliance costs or pass-through without further assumptions. The compliance costs are the cost to society due to the policy; without costly participation, this is simply the change in the expected highest value draw. In the expression above, this is the first bracketed term. However, even without changes in the bid function, the expected markdown associated with the winning bid will change because different values are associated with different markdowns along the bid function. This is the extent to which compliance costs would be passed through even if loggers did not realize their competitors also face compliance costs.

Thus, the expression for the value effect reveals two facts. First, the relative importance of compliance costs and changes in markdowns cannot be estimated using reduced form relationships between bids and restrictions. Second, the value effect only corresponds exactly to the compliance cost if bidders happen to fully pass costs through along the relevant interval of the bid function.

Competition Effect There is also a direct effect on the expected winning bid due to a change in the distribution of a bidder's competitors. In describing the competition effect, I hold the winning value distribution fixed, and allow the bid strategy to change in response to the change in F_{-i}. The markdown term, conditional on a private value, is affected by a change in the distribution of opposing bidders. The numerator of the competition effect roughly corresponds to the expected margin between a given level of v^w and the second-highest value, conditional on v^w being the highest value. That is, if the dispersion of the distribution changes in the neighborhood of the bidder's value, this margin will change for a given v^w. This is a change in the intensity of competition that is "local" to the bidder within the value distribution. The denominator is the probability that a given value will win the contract. This is less intuitive: the incentive compatibility constraint satisfied in equilibrium means that a high-value firm's markdown is disciplined by the possibility that a lower-valued bidder will want to bid like them.

Altogether, an effect on competition can arise even in the absence of endogenous participation. However, endogenous participation fits the setting and reduced-form evidence more convincingly and allows for an additional mechanism.

Participation Threshold Effect Incorporating a bid preparation cost complicates the equilibrium bidding function and reveals a new channel through which restrictions can affect bidding behavior. Because bidders are symmetric, the expected winning bid can still be expressed in closed form given the marginal type v^*, which is an implicit function of K and the distribution of opposing bidders, $F_{-i}(v;r)$ (Hubbard and Paarsch, 2009):

$$E_{v^w}[b_i(v^w;F_{-i}(\cdot;r)] = \int_{v^*}^{\bar{v}} \left[v^w - \frac{\int_{v^*}^{v^w} F_{-i}(u;r)^{N-1} du}{F_{-i}(v^w;r)^{N-1}} - \frac{F_{-i}(v^*;r)^{N-1}}{F_{-i}(v^w;r)^{N-1}}(v^* - R) \right] f_N(v^w;r) dv^w$$

$$K(r) = (v^* - R)[F_{-i}(v^*;r)]^{N-1}$$

The reserve price R has been replaced in the second term of the bid function by the threshold type v^* in the closed-form bid function. This reflects the fact that the participation cost, K, discourages bidders with valuations very close to the reserve price from participating.

The zero-profit condition in the second equation determines the relationship between restrictions and participation behavior. The marginal type v^* will vary with restrictions because the expected payoffs to a participating bidder with a given value draw will change. These changes will arise because restrictions could affect the distribution of opponents, $F_{-i}(\cdot;r)$, or the cost of participation, $K(r)$.

Endogenous changes in participation through v^* will create a feedback effect that may partially counteract the competition effect. Intuitively, if the expected mix of opposing bidders is weaker than before, some types that barely decided not to participate before will now find it worthwhile to submit a bid. In contrast, if the restrictions increase the cost of participation, then loggers will require a higher value draw to justify bidding.

This change in v^* affects markdowns through the second and third terms. There will be a new group of terms in the derivative of the expected winning bid corresponding to the effect of r on v^*.[24] A different type will now be bidding the reserve price, so the equilibrium bids associated with types above v^* must also change in response.

5.3 Parameterizing the Model

I take a parametric approach to estimation similar to Roberts and Sweeting (forthcoming), which allows me to incorporate rich observed and unobserved auction heterogeneity. The observed heterogeneity is analogous to the controls in the reduced-form section and allows me to isolate the effect of seasonal restrictions, while the unobserved heterogeneity plays

[24] For conciseness, I omit the actual expression.

an important role in obtaining realistic bidder margins.

The objects of interest are closely related to the intensity of competition within an auction. Therefore, it is important that I allow for auction-specific unobserved heterogeneity. To understand this, suppose that all auctions appear identical to the econometrician and that the variance of value draws within an auction is quite small. Then bidders will want to bid close to their valuations in equilibrium because they expect their competitors to have very similar valuations. However, if the auctions differ in an unobservable way that is known to the bidders, their bids will vary considerably across auctions. When I pool data across these ostensibly identical auctions, my model and estimates will imply that the value distribution has a relatively large variance. Thus, in simulations, bidder markdowns and profits would be overestimated.[25]

To avoid these issues, I assume the parameters characterizing auction a are drawn from distributions based on observable characteristics and an auction-specific random effect. Each auction a is characterized by a vector of observable characteristics X_a, a participation cost K_a and a distribution of bidder values $F_a \sim TLN(\mu_a, \sigma_a, 0, \bar{v})$. $TLN(\cdot)$ is a lognormal distribution truncated above at $\bar{v}$.[26]

The random effects are assumed to be uncorrelated with the observable characteristics, which is consistent with the reduced-form discussion above. Specifically, I assume the following distributions for $\theta_a = \{\mu_a, \sigma_a, K_a\}$, conditional on $\Gamma = \{\beta, h, \omega\}$:

$$\mu_a \sim N(\beta^\mu X_a + h^\mu(MonthsRestricted_a), \omega^\mu)$$
$$\sigma_a \sim Weibull\ (exp[\beta^\sigma X_a + h^\sigma(MonthsRestricted_a)], \omega^\sigma)$$
$$K_a \sim Weibull\ (exp[\beta^K X_a + h^K(MonthsRestricted_a)], \omega^K),$$

where X_a are observable auction characteristics. The distributions for σ and K must have non-negative support. I specify these parameters using a Weibull distribution, which is bounded below at zero and can take on a variety of shapes. For simplicity, I only allow the scale parameter to vary with X_a and $h(\cdot)$; the shape parameter is common to all types of auctions.

Given the discussion in the previous subsection and the reduced-form evidence, the restrictions could nonlinearly affect auction outcomes. I will allow the months of seasonal restrictions to enter all three distributions as a restricted cubic spline, mirroring the re-

[25]Krasnokutskaya (2009) details this argument further and presents nonparametric identification results in an environment without selective entry.

[26]In practice, I set $\bar{v} = 1500$, which exceeds any observed bid by 300 percent.

gressions already presented. Despite the parametric assumptions, this flexibility within the distribution should help capture the true effects of regulation on auction outcomes and readily allow the decomposition of the equilibrium effects. In practice, I estimate the structural model with the vector of covariates included in Column (4) of Table 4 and Figure 3. The X vector includes all of these covariates. I again specify $h(\cdot)$ as a restricted cubic spline in months of restrictions.

Formal nonparametric identification of the structural model is established by Gentry and Li (2014). They address the case of imperfectly selective entry (which nests the fully selective entry model estimated here as a polar case) in the presence of unobserved auction heterogeneity in values and participation costs. Their argument related to unobserved auction heterogeneity draws on the approach of Hu, McAdams and Shum (2013). Xu (2013) also demonstrates point identification of the fully selective entry model with slightly weaker data requirements.

Informally, identification of the parameters comes from a combination of the data and the distributional assumptions. The parameters of the value distribution are identified by the covariances of the observed auction characteristics, including seasonal restrictions, with features of the bid data. The level of the observed bids, the distances among the bids (especially between the winning and second-highest bids), and the distance from bids to the reserve price are especially informative for separating out the location and scale parameters of the value distribution. The probability of bidder participation has an intuitive link to the participation costs. The number of potential bidders, which is determined in the long run and assumed exogenous to a given auction, provides additional variation. The unobserved heterogeneity is identified by the distributional assumptions and the variation in bidding patterns among observably similar auctions.

5.4 Empirical Implementation of the Model

Given the parametric assumptions and the equilibrium bid functions, I derive the likelihood of a vector of parameters conditional on the observed auction data and describe the Maximum Simulated Likelihood (MSL) estimator. Simulating the equilibrium is computationally non-trivial, so I use importance sampling to reduce the computational burden.

I observe vectors of bids and participation decisions; however, the goal of the structural estimation is to recover the latent distributions of bidder values and auction participation costs. The equilibrium of the auction implies an inverse-bid function that maps bids and participation decisions into valuations, given a value distribution and participation costs.

Thus, I can calculate the likelihood of observing a given vector of bids and participation decisions conditional on auction-specific variables, $\theta = \{\mu, \sigma, K\}$.[27]

A standard concern in the empirical auction literature is that the regularity conditions for MSL estimation do not hold. In particular, Donald and Paarsch (1993) note that the maximum possible bid in a first-price auction depends on all parameters of the value distribution. Thus, even when the support of the value distribution does not depend on parameters, the support of the equilibrium bid strategy still does. In my case, this concern does not apply: the maximum bid in a given auction is indeed a function of the value distribution parameters θ_a; however, the parameter vector to be estimated via MSL is Γ. Because Γ essentially determines the weights of a mixture distribution over θ, the support is actually independent of Γ in the limit: for any value of Γ, it is possible (if very unlikely) that any bid in $[R, \bar{v})$ could be observed. Essentially, the regularity conditions are satisfied through the richly specified unobserved heterogeneity. In practice, I take 1000 draws per data point, which proves sufficient to ensure a non-zero likelihood for every observation.

To accommodate the unobserved heterogeneity in θ, I simulate the integral representing the likelihood of observing a given vector of bids and bidder participation decisions given a guess of the parameter vector Γ. I maximize the log-likelihood function with respect to Γ:

$$\max_{\Gamma} \; \frac{1}{A} \sum_a L_a$$

$$\text{where } \; L_a = \log\left(\int \ell_a(\theta | b_a) p(\theta | \Gamma, X_a) d\theta \right) \approx \log\left(\frac{1}{S} \sum_{s=1}^{S} \tilde{\ell}_a(\theta_{as} | b_a, X_a) \right)$$

where $\mathbf{b_a}$ is a vector of bids and participation decisions observed in auction a, A is the number of auctions in my sample, and θ_{as} is drawn from the density $p(\theta | X_a, \Gamma)$.[28] It is well-known that MSL is consistent only when the number of draws grows sufficiently quickly relative to the sample size. To minimize this concern, I use 1000 simulation draws per observation.

Although I can express the equilibrium bid function in closed form, traditional Monte Carlo simulation is still computationally burdensome. Each evaluation of the likelihood for the full dataset requires solving for the inverse bid function for hundreds of thousands of auctions for each guess of Γ, which takes a non-substantial amount of computing power

[27] A derivation of the bid density is available in Appendix C.

[28] These likelihoods are conditional on the number of potential bidders (N) and the reserve price (R). However, N and R do not enter the importance sampling process because there are no parameters that explicitly depend on them. Thus, I suppress them for notational clarity.

and/or time. Therefore, I adopt an importance sampling approach: Ackerberg (2009) outlines the technique and a number of potential applications in empirical industrial organization, including structural estimation of auctions. Such an approach has been successfully used to estimate similar auction models by Roberts and Sweeting (forthcoming); Bhattacharya, Roberts and Sweeting (2014); and Gentry and Stroup (2014).

Importance sampling involves a change of variables in the integral above:

$$\int \tilde{\ell}_a(\theta|b_a, X_a)\frac{p(\theta|\Gamma, X_a)}{g(\theta|X_a)}g(\theta|X_a)d\theta \approx \frac{1}{S}\sum_{s=1}^{S}\tilde{\ell}_a(\theta_{as}|b_a, X_a)\frac{p(\theta|\Gamma, X_a)}{g(\theta|X_a)}$$

where θ_{as} are now drawn from an initial importance sampling distribution, $g(\theta|X)$. As the guess of the parameter vector Γ changes, the likelihood that a given simulation would have been drawn changes through $p(\theta|\Gamma, X)$. However, no other terms are affected. Essentially, for a given parameter guess, I re-weight the pool of simulation draws by $\frac{p(\theta|\Gamma,X)}{g(\theta|X)}$ to match the density defined by that guess. For instance, if a given simulation was an unlikely draw from $g(\theta|X)$, but a very likely draw from $p(\theta|\Gamma, X)$, this simulation would receive a large weight. This specification allows me to incorporate substantial auction heterogeneity without needing to solve hundreds of thousands of auctions for every candidate parameter vector. Instead, I simply solve for the appropriate vector of weights, which is an inexpensive operation and has an easily-calculated gradient.

The initial importance sampling densities $g(\theta|X)$ are:

$$\begin{aligned}
\mu &\sim \text{Uniform } (0, 6) \\
\sigma &\sim \text{Uniform } (0.01, 2) \\
K &\sim \text{Uniform } (0, 4)
\end{aligned}$$

The intervals are chosen to include all sets of auction parameters that are reasonable upon inspection of the bid data. I obtain similar results when the initial importance sampling distributions are normal (for μ) and Weibull (for σ and K) distributions based on OLS regressions of the bid data. I simulate a new set of auctions based on these first-stage estimates and re-estimate the model. This two-step procedure can help reduce simulation error, as noted in the literature (Ackerberg, 2009).

6 Results of the Structural Estimation

In this section, I simulate auctions using the estimated structural parameters. The model fits well, and the simulations demonstrate that firms almost fully pass through the compliance costs associated with the restrictions. I also perform decompositions to further assess which mechanisms are most important in explaining changes in winning bids.

6.1 Parameter Estimates and Model Fit

I discuss the implications of the parameter estimates for the relationship between the seasonal restrictions and the distribution of valuations and verify the fit of the model. The structural parameters are presented in Table 7, along with standard errors derived from 100 bootstrap replications. The signs of the elements of β^{μ} are as expected, although the restriction splines are difficult to interpret directly. Thus, Figure 6 shows the effect of restrictions on the mean, standard deviation, and participation cost of a contract with otherwise typical observable characteristics in units of $/MBF. The mean value falls by roughly 12 percent for the most-restricted sales. The standard deviation of the value distribution does not vary appreciably, except for a slight (but statistically insignificant) decline for the most restricted sales. The within-auction standard deviation gives some indication of the degree to which a bidder will be able to shade their bid in equilibrium. If the spread is large, then a heavily shaded bid is unlikely to be undercut because each bidder is quite isolated in the distribution of private values. Finally, the mean participation cost rises from $65 for a typical sale ($0.095/MBF*687 MBF) when unrestricted to $82 and $272 when restricted for 6 and 10 months, respectively. The decomposition in the next subsection will provide a quantitative breakdown of how these effects influence equilibrium bids.

In the case of σ and K, as the observable characteristics of the auction change, the location of the distribution will change through β, but the general shape (determined by ω) will remain the same. There is considerable unobserved auction heterogeneity in terms of μ and σ. The distribution of μ implies that a one standard deviation change in this parameter will change the median of the value distribution by 25 percent. In the case of the participation cost, K, there is also some variation. Most sales have very small participation costs. In the mean auction, the participation cost is approximately 0.5 percent of the median private value; this is consistent with the loggers' general familiarity with and close proximity to the sales. Still, bidding on some contracts is considerably more costly than bidding on others, which could reflect some particularly poorly known or isolated sales.

I draw parameters and calculate bids for 10 auctions per observation in the data, and find that my simulated data fit the real dataset quite well.[29] Table 8 demonstrates that I match the mean observed bid, government revenue, winning bid conditional on at least one bidder, and participation rate fairly closely.

The fit conditional on observable characteristics is also very good. I regress all non-zero winning bids on covariates; the coefficients are extremely similar for simulated and true winning bids. Because the spline coefficients for the restriction variable are difficult to interpret, I plot the spline functions in Figure 7.[30] The simulated spline does a relatively good job of matching the data. In effect, this relationship is the bid effect that I will be decomposing.

6.2 Effect of Restrictions on Agent Payoffs

I simulate a representative set of auctions while varying the extent of seasonal restrictions and directly calculate differences in compliance costs, government revenues, and firm surplus. Going forward, I assume the DNR reservation value for the contract is zero, thus assuming the DNR's payoff is equal to the revenues obtained. In Appendix D, I re-calculate the main results under the assumption that the DNR reservation value for a contract is the reserve price observed in the data. Because the likelihood of a contract receiving no bids is increasing in the restrictions, the paper and the appendix effectively provide bounds on the true effect of restrictions on DNR payoffs. The levels of various auction outcomes are shown in Figure 8. The government revenue from an average (687 MBF) sale falls from \$57,000 to \$50,000 moving from 0 to 10 months of restrictions. Point estimates of firm surplus are nearly constant across months of restrictions.[31]

One measure that summarizes the effect of restrictions on auction performance is the share of surplus captured by the bidders. I calculate the percentage of combined firm surplus and government revenue captured by firms and find that it increases slightly with more restrictions: the firm share is 21.2 percent for unrestricted sales, 22.6 percent at 6.5 months, and 22.3 percent at 10 months. The share is significantly different from the unrestricted share for all but the most restricted sales. This occurs because the firm surplus falls by less than government revenues in percentage terms. Thus, in terms of minimizing information rents, the auction performs slightly worse for more restricted sales. The difference represents a 5.2 to 6.6 percent increase in the firm's relative share of surplus.

[29] Appendix F outlines the details of all simulations based on the estimated parameters.

[30] This specification is analogous to the regression that generates Panel A of Figure 3. See Appendix E for the full vector of estimated coefficients and other measures of model fit.

[31] See Appendix E for a plot of changes in firm surplus and government revenue with accompanying confidence intervals.

I calculate compliance costs as the decrease in the expected valuation of the winning bidder, plus any increase in the total bidder participation costs. These compliance costs, shown in the bottom right panel of Figure 8, are pointwise significantly different from zero at a 95-percent confidence level for the interval between 7 and 9 months of restrictions, inclusive. The imprecision beyond that interval reflects the lack of data in the far right tail, but the pointwise estimate still has a p-value of 0.103 at 10 months. For an average-sized sale, point estimates suggest that compliance costs are $2609 for a sale restricted for 6 months and $8258 for a sale restricted for 10 months, which amount to 5 percent and 15 percent of average unrestricted government revenues, respectively. These compliance costs translate to 17 percent of firm surplus if restricted for 6 months and 54 percent of firm surplus if restricted for 10 months.

Changes in government revenues are roughly equal to the compliance costs; the costs of the policy are borne almost entirely by the state. Firm surplus is estimated to be very similar across the full range of restriction intensity. Overall, this estimate is fairly precise: as shown in Figure 8, for restrictions up to 8 months of the year, the 95 percent confidence interval does not include firm surplus increases or decreases of more than $1500. For sales restricted for 10 months, the point estimate is a decrease in firm surplus of $1010. However, this pointwise estimate is less precise: the 95 percent confidence interval is bounded by an increase in firm surplus of $2500 and a decrease of $4500.

Why is there approximately full pass-through? One reason is that contract supply is modeled as inelastic. State-owned timber is exogenously sold in the medium-run: the timber stands sold for harvest are those most necessary for forest management, subject to a minimum annual volume set out by the legislature. Further, the main margin for adjustment is the reserve price, which is set largely through historical prices, with some objective adjustments based on the DNR cost assessment. This exogeneity is clear in 2006-2007, when reserve prices were still quite high despite the housing market crash. Thus, there is no mechanism by which the state's behavior would directly drive changes in markdowns.

Given these supply conditions, there still could have been greater than or less than full pass-through. As described before, markdowns are determined by the extent to which firms are isolated in the distribution. Because the cost of complying with the restrictions does not substantially affect the within-auction dispersion of private values, the costs are passed through to the state at nearly a one-to-one rate.

6.3 Decomposing the Bid Effect

In this subsection, I delve further into the quantitative importance of different mechanisms in explaining the differences in equilibrium bids. Specifically, I decompose the equilibrium bid effect into the value, competition, and participation threshold effects described in Section 5.2 using several sets of auction simulations. Each set involves simulating 10 auctions corresponding to each data observation over a grid of seasonal restrictions from 0 to 10 months.

I vary three objects: the distribution of value draws, $F_i(v;r)$; the perceived distribution of opponents values, $F_{-i}(v;r)$ as it directly affects the bidder markdowns; and the participation threshold, $v^*(r)$. The participation threshold varies because of changes in $F_{-i}(v;r)$ and $K(r)$. First, I estimate auctions fixing all three objects as though the auctions are unrestricted. Second, I allow the value draws to reflect the level of seasonal restrictions but do not vary the perceived opponent distribution or participation threshold. A heuristic description is that loggers know their valuation including compliance costs, but don't realize that other firms would also face compliance costs. This isolates the value effect. Third, I allow the perceived opponent distribution to reflect seasonal restrictions, but still hold the participation threshold constant. Here, the heuristic is that firms observe their values and compliance costs and make their participation decision. Then, before determining their actual bid, they find out that their competitors are drawing from a distribution that is affected by compliance costs. This change isolates the competition effect. Fourth, I recalculate the participation threshold to reflect the effect of seasonal restrictions on $F_{-i}(v;r)$ and $K(r)$, thus endogenizing the entirety of the bidding decision; this incremental change isolates the participation threshold effect. Table 9 summarizes my approach.

The decomposition in Panel B of Figure 9 indicates that the value effect is the most important mechanism at play. The value effect accounts for 73 to 82 percent of the effect on bids throughout the grid of seasonal restrictions. It does not fully account for compliance costs, however. On net, the adjustment of bidding strategies and participation decisions contribute 18 to 27 percent of the total effect on bids. Breaking this strategic change down further, the competition effect depresses bids by 20 to 38 percent of the net bid effect. Unlike the other two effects, the participation threshold effect increases the bids. When bidders observe that their competitors will be weaker on average, this increases the expected profits of previously marginal participants and reduces the threshold value draw needed to participate. In this case, the participation of these additional bidders pushes bids back up 0 to 15 percent of the net bid effect. Taking an average contract restricted for 10 months as an example, the value effect is -\$5882, the competition effect is -\$2312, and the participation threshold effect is +\$1035.

6.4 Implications for DNR Conservation Policy

These results have immediate implications for DNR conservation policy. First, the costs and revenue effects of seasonal restriction are highly nonlinear in the number of months restricted. Accounting for the observed distribution of restrictions, the mean compliance cost for a contract is \$918 and the revenue effect is \$874. Much of this burden is driven by the most restricted contracts. The marginal compliance cost of the 8th through 10th months of restrictions on a typical sale is roughly \$1600 per month and the loss in revenue is roughly \$1100 per month. In benefit-cost terms, if the marginal conservation and recreational benefit of the tenth month of restrictions is less than \$1600 on a sale of average characteristics, then the state should consider relaxing these restrictions. However, I estimate the costs of the first 4 months of restrictions are quite small and not distinguishable from zero. The state need not consider relaxing these less stringent restrictions, and could even consider the benefits of placing additional restrictions on tracts that would otherwise go unregulated under the existing regime.

Second, firm surplus is somewhat affected. The level of firm surplus is \$1010 lower for the most-restricted contracts, although this point estimate is not statistically significantly from zero. The compliance cost is much larger: it is \$8258 for these contracts. This difference suggests that most of the costs are passed through to the state. Still, in terms of political economy, this loss in surplus lends credence to complaints from loggers regarding the most-restricted sales.

From a standpoint of fostering competition, the auctions for restricted contracts are less effective for the state. The loggers capture a larger *share* of the surplus for contracts restricted 6 or more months per year: they capture a share of auction surplus that is 1 to 1.5 percentage points larger relative to a baseline of 21.2 percent for an unrestricted sale. In percentage terms, the firms are slightly more successful in capturing auction rents for more restricted contracts, even though the level of rents has fallen slightly.

One caveat is in order in terms of policy analysis. The effects discussed in the paper are identified using cross-sectional variation in the seasonal restrictions. While the analysis estimates and decomposes the differences in mean winning bids at various levels of restriction stringency, these estimates are specific to the given context and inform the effects of marginal changes from the DNR's policy. Because the value of an unrestricted contract is partially determined by the nature of the restrictions on other contracts, there are some spillovers.

That is, my estimates cannot reliably predict the impacts of abolishing all restrictions. Summing across the entire sample, the compliance costs amount to $4.8 million. The total negative revenue effects are $4.5 million, or roughly 1.4 percent of the $313 million in timber sale revenues collected over the 10-year sample period. This likely represents an upper bound on the differential impacts of restrictions on winning bids. That being said, the marginal effects estimated here should be valid for a change in the restriction status of a small group of contracts. Further, the decomposition exercise remains valid as a means to understanding the relative importance of differences in competitive pressure and valuations in determining the effects of current DNR policy.

Finally, this welfare analysis does not consider the increase in future revenue to the government (and profits to future bidding firms) from having a healthier forest going forward. Indeed, some of the restrictions can be considered an investment that improves the future quality of the state forest system. As a stand of trees is only harvested every 50-60 years or so, much of this would be deeply discounted in present value terms. However, there are some benefits – such as avoided damages to adjacent areas – that could be realized much sooner.

6.5 Optimal Reserve Prices

The government could attempt to recover revenues lost due to seasonal restrictions by setting optimal reserve prices. I simulate the optimal reserve price separately at each level and compare outcomes.[32] First, I randomly select 500 auctions from the data (i.e., the X vector and number of potential bidders N) with replacement. At each level of restrictions, I draw auction parameters $\{\mu, \sigma, K\}$ and associated valuations for 50 simulations per observed auction. Holding these draws constant, I calculate the outcomes of each simulated auction over a fine grid of possible reserve prices. I then average the outcomes across the 50 simulations to find the optimal reserve price for each observed auction.

Table 10 compares the mean reserve price, government revenue, firm surplus, and number of participating bidders for auctions using the observed reserve prices versus the simulated optimal reserve prices at 0, 3, 6, and 10 months of restrictions. Assuming that the

[32]For the purpose of this section, "optimal reserve price" refers to the reserve price that maximizes the government's expected revenue, as is typical in the optimal auction literature. The true government objective function may balance a number of criteria. Appendix D presents these results under the assumption that v_0 is equal to the observed reserve price, providing an informative bound.

government's valuation is zero, the simulations suggest that the DNR is typically setting reserve prices above the optimum. The average optimal reserve price decreases with the restrictions as the typical bidder's valuation falls. The optimal reserve price entices approximately one additional bidder into the auctions on average.

The simulations show that optimal reserve prices can be used to blunt the revenue impact of seasonal restrictions. Panel A of Figure 10 compares the share of revenues lost due to restrictions when using observed reserve prices versus optimal reserve prices. That is, at each level of restrictions, the revenue is compared to the revenue from an unrestricted sale under same reserve price regime. The revenue gap between restricted and unrestricted sales narrows from 12.2 percent to 9.6 percent when using optimal reserve prices; the state is able to insulate itself somewhat from the costs of the restrictions.

Further, an optimal reserve price policy can increase revenues by a magnitude comparable to the losses incurred due to restrictions. Panel B of Figure 10 presents the average revenues in levels, comparing the different reserve price regimes. Unrestricted sales using observed reserve prices bring in revenues comparable to 8-month-restricted sales using optimal reserve prices. Similarly, 10-month-restricted sales using optimal reserve prices bring in revenues similar to those captured by 5-month-restricted sales using the observed reserve prices.

7 Conclusion

Government contracts are often competitively allocated; however, this process could be undermined by a recent proliferation of environmental objectives. I detail the mechanisms by which compliance with conservation restrictions could either weaken or intensify competitive pressure in auctions for Michigan timber contracts. I find that the restrictions are associated with lower winning bids and fewer bidders. I use a structural model to disentangle the various mechanisms and welfare impacts of the policy. In this context, loggers are able to fully pass the costs of the policy on to the government by modifying their bidding strategies. Importantly, I estimate that compliance costs are nearly zero for all but the most severe restrictions. Restriction-contingent optimal reserve prices partially close the revenue gap between less- and more-restricted contracts, though the gap does persist. These findings highlight the need to consider the presence of strategic firm behavior and nonlinear compliance costs when predicting or evaluating the effects of environmental policy.

These results have broader implications for assessing and predicting the impacts of environmental contracting objectives. Policy evaluations and projections that ignore strategic behavior could misinform the political economy discussion, and inaccurately estimate or project policy costs. A simple *ex post* estimate of the cost of the DNR policy would also require strict assumptions about pass-through: a basic comparison of bids would have underestimated the full costs of the policy by roughly 10 percent. Conversely, using cost estimates to project the impact on bids before implementing the program would also be uninformative. Furthermore, like seasonal restrictions, many policies are characterized by implicit or opportunity costs, which must be estimated using revealed behavior.

There are two reasons that one might expect the impacts of environmental goals to be even larger in other settings. My results suggest that seasonally-differentiated regulations can impose costs on firms by reducing flexibility or forcing production to shift to less-profitable times of year. However, the timber contract restrictions still allow loggers to operate on at least some land year-round; one might expect that a uniformly timed seasonal restriction would be even more costly. This situation applies to other markets in which all production is constrained in the same season, such as oil and gas drilling, commercial fisheries, or any number of industries affected by seasonal ozone regulation.

In addition, the competitive implications of environmental objectives could be more severe if a market is particularly thin. In that case, heterogeneity in firms' compliance costs could play a larger role in distorting rents. Furthermore, different methods of implementation would likely have varying effects on firm competition. For instance, if an agency only considered bids from environmentally-certified firms, as is the case for LEED green building requirements, firms would have to incur a large fixed cost simply to participate. The resulting effect on market structure could negatively impact competition. Future analysis of different types of environmental objectives in a variety of settings could further inform the contracting process.

References

Ackerberg, Daniel, "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*, 2009.

Aral, Karca D., Damian R. Beil, and Luk N. Van Wassenhove, "Total-Cost Procurement Auctions with Sustainability Audits to Inform Bid Markups," *Working Paper*, 2014.

Athey, Susan and Jonathan Levin, "Information and Competition in U.S. Forest Service Timber Auctions," *Journal of Political Economy*, 2001.

—, Dominic Coey, and Jonathan Levin, "Set-Asides and Subsidies in Auctions," *American Economic Journal: Microeconomics*, 2013.

—, Jonathan Levin, and Enrique Seira, "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions," *Quarterly Journal of Economics*, 2011.

Bajari, Pat and Greg Lewis, "Procurement Contracting with Time Incentives: Theory and Evidence," *Quarterly Journal of Economics*, 2011.

Balat, Jorge, "Highway Procurement and the Stimulus Package: Identification and Estimation of Dynamic Auctions with Unobserved Heterogeneity," *Working Paper, Johns Hopkins University*, 2013.

Bhattacharya, Vivek, James W. Roberts, and Andrew Sweeting, "Regulating Bidder Participation in Auctions," *RAND Journal of Economics*, 2014.

Brown, Jennifer, Justine Hastings, Erin Mansur, and Sofia Villas-Boas, "Reformulating Competition? Gasoline Content Regulation and Wholesale Gasoline Prices," *Journal of Environmental Economics and Management*, 2008.

Brown, Ross, Michael Kilgore, Jay Coggins, and Charles Blinn, "The impact of timber-sale tract, policy, and administrative characteristics on state stumpage prices: An econometric analysis," *Forest Policy and Economics*, 2012.

—, —, —, —, and Catherine Pfender, "Assessing State Timber Sale Policies, Programs, and Stumpage Price Drivers," 2010. Department of Forest Resources, University of Minnesota.

Busse, Meghan R. and Nathaniel O. Keohane, "Market effects of environmental regulation: Coal, railroads, and the 1990 Clean Air Act," *RAND Journal of Economics*, 2007.

Donald, Stephen G. and Harry J. Paarsch, "Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions," *International Economic Review*, 1993.

Exec. Order No. 13423, 72 Fed. Reg. 3919, 2007.

Exec. Order No. 13514, 74 Fed. Reg. 52117, 2009.

Exec. Order No. 13693, 80 Fed. Reg. 15871, 2015.

Fabra, Natalia and Mar Reguant, "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, 2014.

Gentry, Matthew and Caleb Stroup, "Entry and Competition in Takeover Auctions," *Mimeo*, 2014.

_ **and Tong Li**, "Identification in Auctions with Selective Entry," *Econometrica*, 2014.

Groeger, Joachim, "A Study of Participation in Dynamic Auctions," *International Economic Review*, 2014.

Haile, Philip A. and Elie Tamer, "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*, 2003.

Holt, Charles A., "Competitive Bidding for Contracts under Alternative Auction Procedures," *Journal of Political Economy*, 1980.

Hu, Yingyao, David McAdams, and Matthew Shum, "Identification of first-price auction with non-separable unobserved heterogeneity," *Journal of Econometrics*, 2013.

Hubbard, Timothy P. and Harry J. Paarsch, "Investigating bid preferences at low-price, sealed-bid auctions with endogenous participation," *International Journal of Industrial Organization*, 2009.

Jofre-Bonet, Mireia and Martin Pesendorfer, "Estimation of a Dynamic Auction Game," *Econometrica*, 2003.

Krasnokutskaya, Elena, "Identification and Estimation in Procurement Auctions Under Unobserved Auction Heterogeneity," *Review of Economic Studies*, 2009.

_ **and Katja Seim**, "Bid Preference Programs and Participation in Procurement," *American Economic Review*, 2011.

Levin, Dan and James L. Smith, "Equilibrium in Auctions with Entry," *American Economic Review*, 1994.

Li, Tong and Xiaoyong Zheng, "Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions," *Journal of Econometrics*, 2012.

Mackes, Kurt, "Sawmill Short Course: Log Inputs – Measurement and Conversion Factors," Technical Report, Colorado State University 2004.

Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand, “Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt,” *American Economic Review*, 2013.

Manuel, Kate M. and L. Elaine Halchin, “Environmental Considerations in Federal Procurement: An Overview of the Legal Authorities and Their Implementation,” Technical Report, Congressional Research Service 2013.

Marion, Justin, “Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions,” *Journal of Public Economics*, 2007.

Paarsch, Harry J., “Deriving an estimate of the optimal reserve price: An application to British Columbian timber sales,” *Journal of Econometrics*, 1997.

Riley, John and William Samuelson, “Optimal Auctions,” *American Economic Review*, 1981.

Roberts, James and Andrew Sweeting, “Bailouts and the Preservation of Competition: The Case of the Federal Timber Contract Payment Modification Act,” *American Economic Journal: Microeconomics*, forthcoming.

Roberts, James W., “Unobserved Heterogeneity and Reserve Prices in Auctions,” *RAND Journal of Economics*, 2013.

Ryan, Stephen, “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 2012.

Samuelson, William F., “Competitive Bidding with Entry Costs,” *Economics Letters*, 1985.

Smith, V. Kerry, Roger H. von Haefen, and Wei Zhu, “Do Environmental Regulations Increase Construction Costs for Federal-Aid Highways? A Statistical Experiment,” *Journal of Transportation and Statistics*, 1999.

Srivastava, Ajit, Dalia Abbas, Fei Pan, and Chris Saffron, “Improving Forest Feedstock Harvesting, Processing, and Hauling Efficiencies,” Technical Report, Forest Feedstock Center of Energy Excellence, Michigan State University 2011.

Sweeting, Andrew and Vivek Bhattacharya, “Selective Entry and Auction Design,” *International Journal of Industrial Organization*, 2015.

United States Department of Defense, "Department of Defense Green Procurement Program Strategy," 2008.

U.S. General Services Administration, "DDS 3 Green Features Fact Sheet," `www.gsa.gov/portal/getMediaData?mediaId=196419` 2015. Accessed: 24 Nov 2015.

Xu, Pai, "Nonparametric Estimation of Entry Cost in First-Price Procurement Auctions," *Journal of Applied Econometrics*, 2013.

Figures and Tables

Figure 1: Sample seasonal restriction

5.2.3.3 - Bark slippage restriction (3/09)
Within Payment Unit(s) 1 and 2, unless changed by written agreement,cutting and skidding are not permitted during the period of April 15 to July 15. This restriction is because of bark slippage.

Figure 2: Distribution of months of restrictions

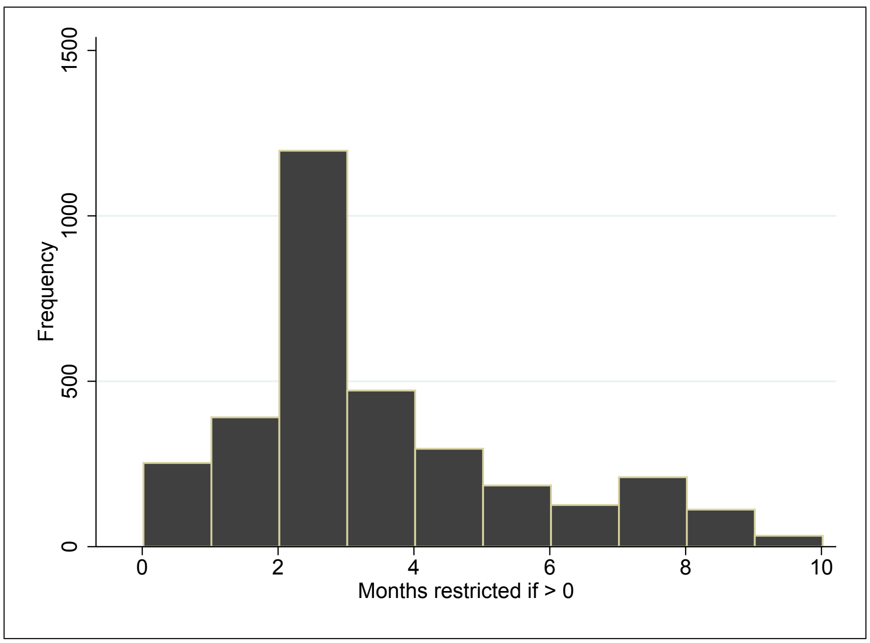

Note: This histogram excludes 1933 auctions with no restrictions.

Figure 3: Reduced-form effect of seasonal restrictions on auction outcomes

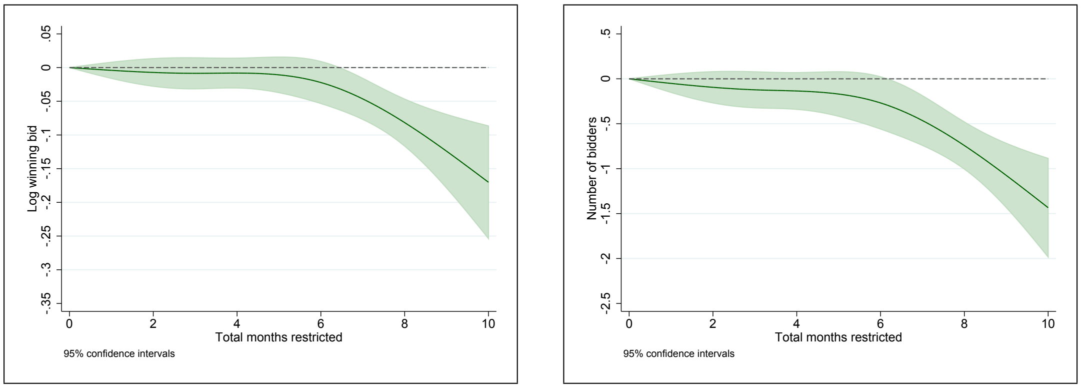

Panel A: ln(winning bid) Panel B: Number of bids

Notes: Figure displays the results of regressing the logarithm of the winning bid (Panel A) or number of bidders (Panel B) on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 4 in terms of control variables and fixed effects included. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions in Panel A and 5207 in Panel B.

Figure 4: Effect of seasonal restrictions, controlling for restriction categories

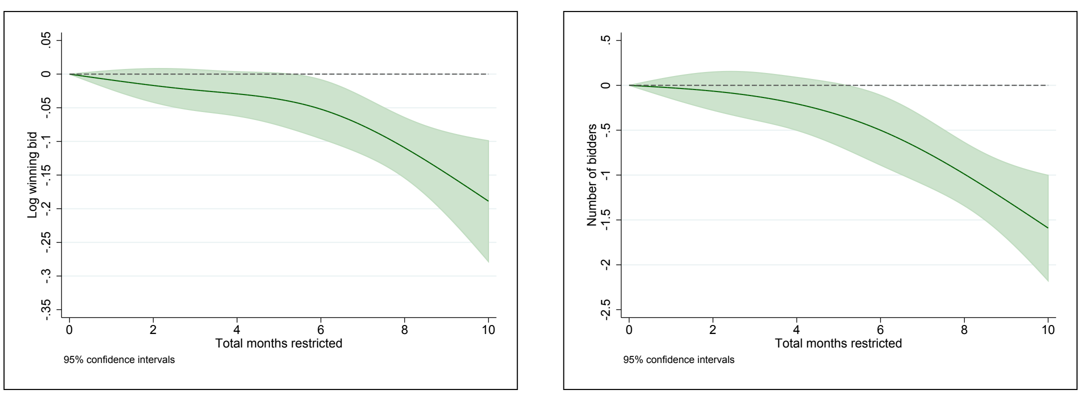

Panel A: ln(winning bid) Panel B: Number of bids

Notes: Figure displays the results of regressing the logarithm of the winning bid (Panel A) or number of bidders (Panel B) on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 6 in terms of control variables and fixed effects included. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions in Panel A and 5207 in Panel B.

Figure 5: Log winning bid, controlling for number of bidders

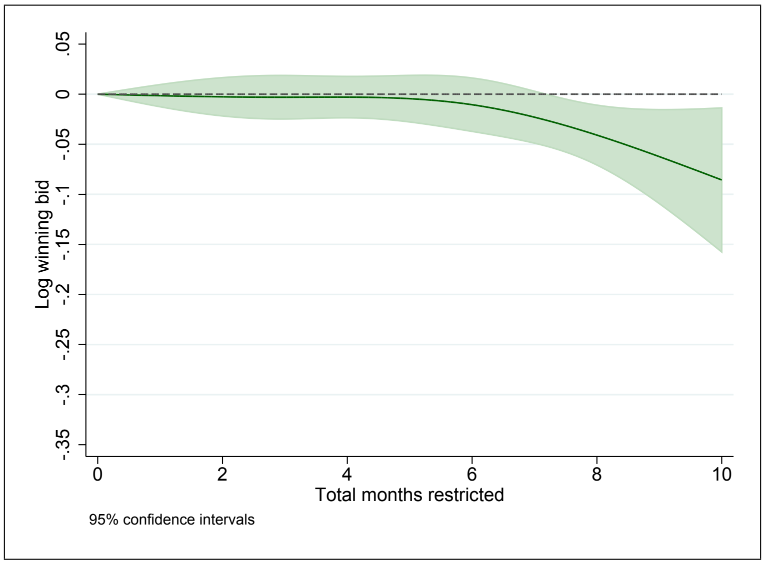

Notes: Figure displays the results of regressing the logarithm of the winning bid on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 6, except that it includes dummies for the number of bids received. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions.

Figure 6: Moments of value distributions at mean X values

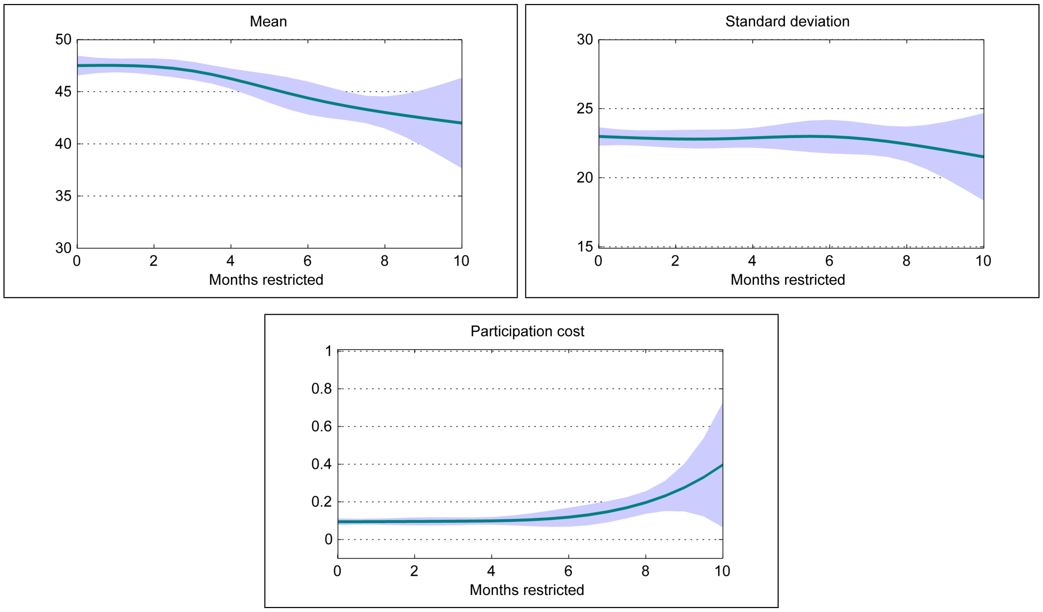

Notes: These are the means and standard deviations of the truncated log-normal value distribution and the participation cost implied by the structural estimates in Table 7, expressed in \$/MBF. They are evaluated for an auction with mean covariate values as the number of months during which the sale is restricted varies. The unobservable components of μ_a, σ_a, K_a are set equal to their mean values. 95% confidence intervals are derived using standard errors from 100 bootstrap replications.

Figure 7: Splines from log winning bid regressions

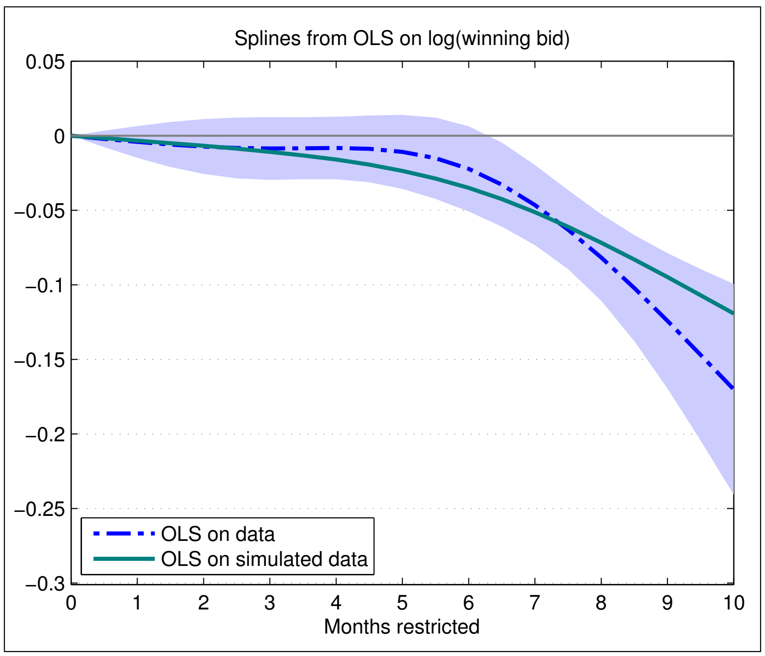

Notes: This figure compares the estimated effect of restrictions in an OLS regression of ln(winning bid) using the real data and the data simulated using the structural estimates in Table 7. 95% confidence interval is for the data (reduced-form) estimate, and are OLS standard errors to be conservative regarding model fit.

Figure 8: Mean auction outcomes, by level of seasonal restrictions

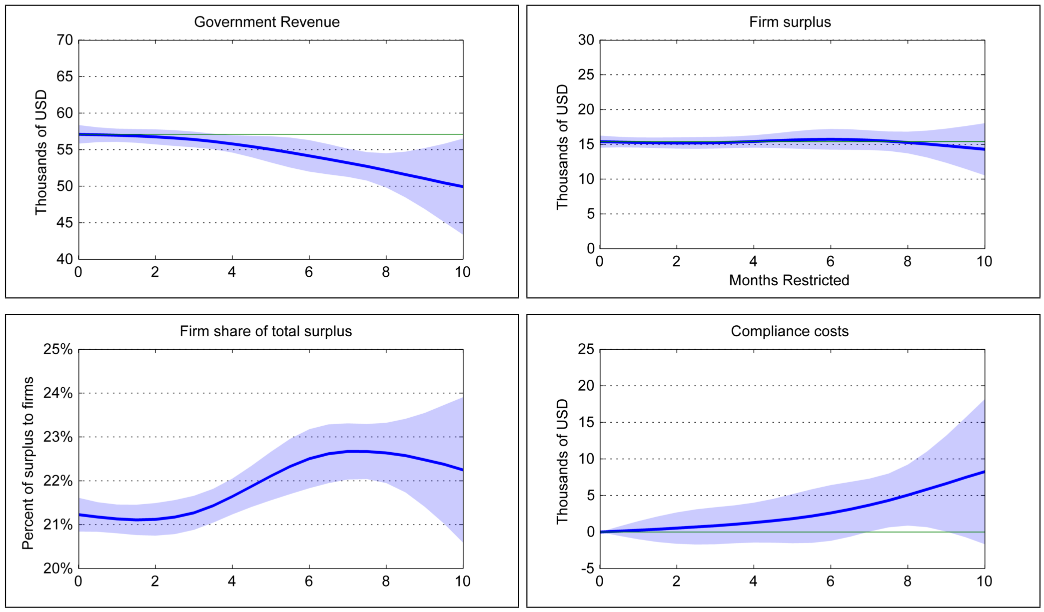

Notes: These are the mean outcomes of auctions simulated according to the distributions implied by Table 7. I simulate 10 auctions per data observation, holding restrictions fixed at zero months. I repeat this at each level of restrictions $\{0.5, 1, \ldots, 9.5, 10\}$. 95% confidence intervals are derived using standard errors from 100 bootstrap replications.

Figure 9: Decomposition of equilibrium bid effect: Mechanisms

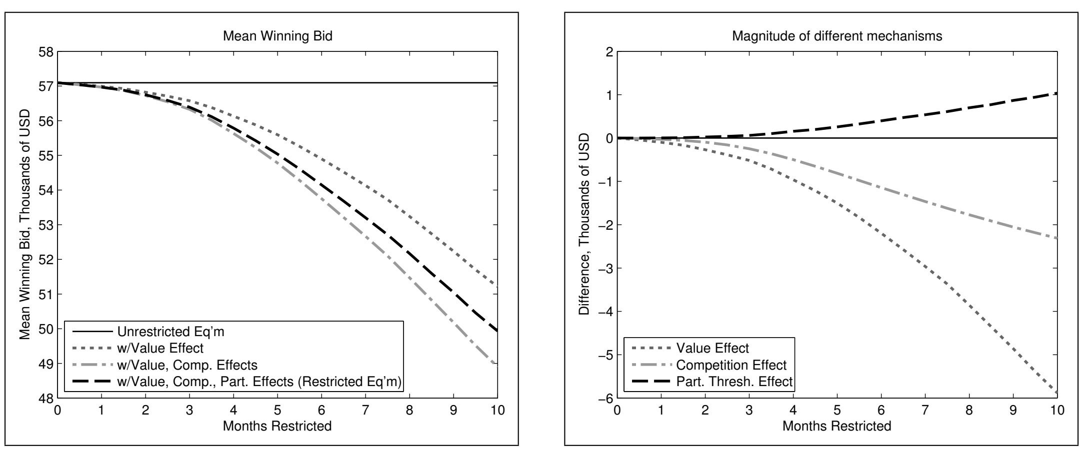

Panel A: Mean winning bid

Panel B: Effect of different mechanisms

Notes: This figure decomposes of the full effect of seasonal restrictions on bidding. Panel A shows the mean winning bid as the number of months restricted vary, as the value, competition, and participation threshold effects are iteratively added in. Panel B shows the changes in the mean winning bid due to each individual effect.

Figure 10: Impact of restrictions under different reserve price regimes, $v_0 = 0$

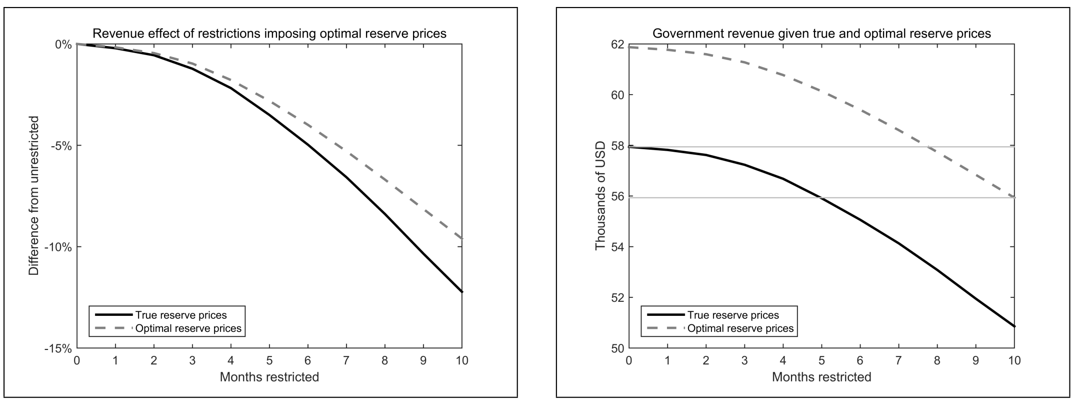

Panel A: Relative impact of restrictions

Panel B: Expected revenues

Notes: I assume the government's value of keeping the contract is equal to zero. These are the relative changes in government revenues as seasonal restrictions vary with optimal reserve prices. I simulate 50 auctions for 500 randomly-drawn data observations, holding restrictions fixed at zero months. I repeat this at each level of restrictions $\{1, \ldots, 10\}$.

Table 1: Major restriction categories

Restriction	Sample auctions affected	Share of sample
Bark slip	1720	33%
Oak wilt concern	847	16%
Soil/wet ground restrictions	642	12%
Winter recreation	295	6%
Forest regeneration	279	5%
Wildlife/endangered species protection	234	4%
Other recreation	87	2%
Nearby private landowner requests	24	$< 1\%$
Misc. others	136	3%
No restrictions	1933	37%

Source: Author's calculations from Michigan DNR timber contracts.

Table 2: Summary statistics

	Obs	Mean	Std. Dev.	P10	P25	P50	P75	P90
Auction Outcomes								
Winning bid ($/MBF)	4750	92.2	50.2	46.8	60.2	79.4	111.3	151.2
Number of bidders	5207	3.9	2.7	1	2	3	6	8
Participation rate	5207	0.2	0.2	0	0.1	0.2	0.3	0.5
Contract Characteristics								
Reserve price ($/MBF)	5207	61.6	32.5	30.3	39.3	53.1	75.4	103.6
Potential bidders	5207	17.8	7	9	13	17	22	27
Months restricted (if >0)	3274	3.7	2.1	1.3	2.4	3	4.5	7.5
Total volume (MBF)	5207	687.4	565.1	189.7	304.9	527	887.9	1371.5
Acres	5207	90.4	65.6	30	44	72	116	176
DNR cost factors	5207	0.7	0.1	0.6	0.6	0.7	0.8	0.9

Notes: Statistics for "Winning bid" exclude 457 auctions that received no bids above the reserve price. Statistics for "Months restricted (if >0)" exclude the 1933 sales with zero months of restrictions. Contract length is missing for 329 observations; the vast majority of these auctions receive no bids.

Table 3: Covariate balance

	Restricted <4 months	Restricted ≥ 4 months	Total
Acres	90.7	89.2	90.4
Upper peninsula	0.36	0.51	0.39
Potential entrants	17.9	17.2	17.8
DNR cost factor	0.72	0.72	0.72
Total volume (MBF)	693.7	663.3	687.4
Vol. share softwood pulp	0.51	0.53	0.52
Vol. share softwood sawlogs	0.093	0.082	0.091
Vol. share hardwood pulp	0.32	0.33	0.32
Vol. share hardwood sawlogs	0.075	0.058	0.071
Species HHI	0.43	0.44	0.43
Percent bid species	0.90	0.90	0.90

Table 4: Linear regressions, log winning bid

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Months restricted	-0.012***	-0.010***	-0.010***	-0.008***	-0.007***	-0.004**
	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
Share Softwood: sawlogs		2.016***	1.721***	1.682***	1.703***	1.808***
		(0.118)	(0.093)	(0.087)	(0.087)	(0.081)
Share Hardwood: sawlogs		1.542***	1.181***	1.207***	1.270***	1.274***
		(0.046)	(0.050)	(0.046)	(0.047)	(0.044)
Share Hardwood: pulpwood		0.403***	0.175***	0.151***	0.174***	0.258***
		(0.033)	(0.032)	(0.029)	(0.026)	(0.025)
Upper peninsula		0.347***	0.296***	0.303***		0.259***
		(0.030)	(0.028)	(0.020)		(0.018)
DNR cost factors		1.122***	1.115***	0.804***	0.908***	0.754***
		(0.082)	(0.068)	(0.061)	(0.055)	(0.054)
Log acres			0.037***	0.040***	0.046***	0.010
			(0.009)	(0.008)	(0.007)	(0.007)
Species-product HHI			0.678***	0.694***	0.683***	0.643***
			(0.042)	(0.038)	(0.036)	(0.034)
Percent bid species			0.812***	0.695***	0.607***	0.596***
			(0.089)	(0.070)	(0.066)	(0.060)
Constant	4.436***	3.062***	1.907***	2.369***	2.315***	2.725***
	(0.018)	(0.063)	(0.096)	(0.086)	(0.090)	(0.076)
Observations	4,751	4,750	4,750	4,750	4,750	4,750
R-squared	0.004	0.414	0.574	0.645	0.663	0.720
Major species dummies	-	-	X	X	X	X
Quarter dummies	-	-	-	X	X	X
Year dummies	-	-	-	X	X	X
Management Unit dummies	-	-	-	-	X	-
Number of bidder dummies	-	-	-	-	-	X

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Standard errors are clustered by county-year.

Table 5: Linear regressions, number of bidders

VARIABLES	(1)	(2)	(3)	(4)	(5)
Months restricted	-0.082***	-0.113***	-0.106***	-0.077***	-0.102***
	(0.023)	(0.022)	(0.021)	(0.017)	(0.016)
Share Softwood: sawlogs		-2.092***	-2.475***	-2.248***	-1.450***
		(0.645)	(0.675)	(0.550)	(0.526)
Share Hardwood: sawlogs		-0.351	-1.380***	-1.404***	-1.154***
		(0.452)	(0.464)	(0.431)	(0.432)
Share Hardwood: pulpwood		-1.507***	-2.094***	-2.105***	-1.781***
		(0.225)	(0.253)	(0.227)	(0.222)
Upper peninsula		1.146***	1.053***	0.964***	
		(0.216)	(0.225)	(0.167)	
DNR cost factors		2.065***	2.056***	0.770*	0.918**
		(0.486)	(0.463)	(0.411)	(0.416)
Log acres			0.593***	0.587***	0.677***
			(0.070)	(0.060)	(0.057)
Species-product HHI			0.757**	0.785***	0.995***
			(0.295)	(0.248)	(0.236)
Percent bid species			2.141***	1.894***	1.517***
			(0.620)	(0.548)	(0.533)
Constant	4.060***	2.891***	-1.285*	0.847	0.797
	(0.126)	(0.364)	(0.727)	(0.724)	(0.721)
Observations	5,209	5,207	5,207	5,207	5,207
R-squared	0.005	0.084	0.133	0.282	0.319
Major species dummies	-	-	X	X	X
Quarter dummies	-	-	-	X	X
Year dummies	-	-	-	X	X
Management Unit dummies	-	-	-	-	X

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Standard errors are clustered by county-year.

Table 6: Linear regressions controlling for restriction categories

VARIABLES	(1) Ln(win bid)	(2) Ln(win bid)	(3) # Bidders	(4) # Bidders
Months restricted	-0.008***	-0.008***	-0.077***	-0.077***
	(0.002)	(0.002)	(0.017)	(0.017)
Share Softwood: sawlogs	1.682***	1.682***	-2.248***	-2.248***
	(0.087)	(0.087)	(0.550)	(0.550)
Share Hardwood: sawlogs	1.207***	1.207***	-1.404***	-1.404***
	(0.046)	(0.046)	(0.431)	(0.431)
Share Hardwood: pulpwood	0.151***	0.151***	-2.105***	-2.105***
	(0.029)	(0.029)	(0.227)	(0.227)
Upper peninsula	0.303***	0.303***	0.964***	0.964***
	(0.020)	(0.020)	(0.167)	(0.167)
DNR cost factors	0.804***	0.804***	0.770*	0.770*
	(0.061)	(0.061)	(0.411)	(0.411)
Log acres	0.040***	0.040***	0.587***	0.587***
	(0.008)	(0.008)	(0.060)	(0.060)
Species-product HHI	0.694***	0.694***	0.785***	0.785***
	(0.038)	(0.038)	(0.248)	(0.248)
Percent bid species	0.695***	0.695***	1.894***	1.894***
	(0.070)	(0.070)	(0.548)	(0.548)
Constant	2.369***	2.369***	0.847	0.847
	(0.086)	(0.086)	(0.724)	(0.724)
Observations	4,750	4,750	5,207	5,207
R-squared	0.645	0.645	0.282	0.282
Restriction Categories	-	X	-	X

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Standard errors are clustered by county-year.

Table 7: Estimated structural parameters

	Covariate	$\mu \sim$ Normal Param	SE	$\sigma \sim$ Weibull Param	SE	$K \sim$ Weibull Param	SE
β							
	Constant	1.136	(0.117)	0.005	(0.156)	0.773	(0.995)
	Restr. Spline 1	0.024	(0.113)	-0.063	(0.139)	0.091	(0.895)
	Restr. Spline 2	-0.765	(0.580)	0.812	(0.696)	0.250	(5.011)
	Restr. Spline 3	1.979	(1.939)	-2.413	(2.298)	6.154	(16.186)
	Upper Peninsula	0.491	(0.025)	-0.258	(0.036)	-0.604	(0.173)
	Ln (Acres)	0.074	(0.010)	-0.019	(0.015)	-0.344	(0.101)
	Softwood Sawlogs	1.867	(0.088)	-0.181	(0.120)	2.914	(0.758)
	Hardwood Pulp	0.198	(0.037)	-0.084	(0.047)	2.596	(0.250)
	Hardwood Sawlogs	1.406	(0.051)	-0.303	(0.076)	1.573	(0.487)
	Pct. Bid Species	0.866	(0.100)	-0.255	(0.143)	-2.249	(0.805)
	Species-Product HHI	0.679	(0.048)	0.058	(0.060)	-0.033	(0.354)
	DNR Cost Factors	0.902	(0.069)	-0.140	(0.087)	-2.399	(0.479)
ω		0.254	(0.004)	5.978	(0.217)	0.938	(0.040)

Notes: Spline variables are the basis functions for a restricted cubic spline on restrictions with knots at 0, 3, 6, and 10 months. Standard errors are calculated from 100 bootstrap replications. μ is normally distributed, while σ and K have Weibull distributions. The β vectors also include parameters for year, quarter of year, and tree species dummies. In the case of the Weibull distributions, the scale parameter is $\exp(\beta X)$ and the shape parameter is ω. Estimation is based on 5,207 auctions, which receive a total of 20,502 non-zero bids.

Table 8: Model fit: sample moments

Moment	Data	Simulation
Mean Observed Bid	79.2	81.4
Mean Auction Revenue	84.1	81.8
Mean Observed Winning Bid	92.2	95.8
Mean Participation Rate	0.24	0.21

Table 9: Decomposition schematic: Mechanisms

	What varies?	Value Realization	Perceived Opponent	Participation Threshold
1	Unrestricted	$F_i(v;0)$	$F_{-i}(v;0)$	$v^*(0)$
2	+ Own values	$F_i(v;r)$	$F_{-i}(v;0)$	$v^*(0)$
3	+ Opponents' value distribution	$F_i(v;r)$	$F_{-i}(v;r)$	$v^*(0)$
4	+ Participation threshold	$F_i(v;r)$	$F_{-i}(v;r)$	$v^*(r)$
	Channel	Derivation		
	Value effect	2-1		
	Comp effect	3-2		
	Participation threshold effect	4-3		
	Net effect	4		

Notes: This table summarizes the decomposition of the equilibrium bid effect. Functions with an argument of zero are evaluated for auctions using the unrestricted parameters. Functions with an argument of r are evaluated over a grid of restrictions from 0 to 10 months. The results of this decomposition are presented in Figure 9.

Table 10: Mean outcomes considering optimal reserve prices, $v_0 = 0$

	None		3 Months		6 Months		10 Months	
Reserve Price Used	Data	Optimal	Data	Optimal	Data	Optimal	Data	Optimal
Reserve Price ($1000s)	42.3	37.5	42.3	37.1	42.3	35.7	42.3	33.5
Gov. Revenue ($1000s)	57.9	61.9	57.2	61.3	55.1	59.4	50.9	55.9
Firm Surplus ($1000s)	15.3	16.6	15.2	16.5	15.6	17.2	14.2	16.0
Number of bidders	3.6	4.5	3.5	4.5	3.4	4.4	3.1	4.4

Notes: Optimal reserve price refers to the reserve price that maximizes expected government revenue. "Revenue" is the winning bid if any bidders participate, and $v_0 = 0$ if not.