Essays on Environmental and Natural Resource Economics

by

Evan Matthew Herrnstadt

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2015

Doctoral Committee:

Associate Professor Ryan Mayer Kellogg, Chair
Associate Professor Jeremy Fox
Assistant Professor Catherine Hausman
Assistant Professor Erich J. Muehlegger, University of California - Davis

©Evan Matthew Herrnstadt

2015

For Karen

ACKNOWLEDGEMENTS

Thanks are owed to my committee members Ryan Kellogg, Jeremy Fox, Catherine Hausman, and Erich Muehlegger (with whom I co-authored the second chapter of this dissertation) for consistent support and feedback; to Dan Ackerberg, Varanya Chaubey, Aaron Flaaen, Joe Golden, Max Kapustin, Jason Kerwin, Justin Ladner, Eric Lewis, Johannes Norling, Karen Stockley, Richard Sweeney, Ophira Vishkin, James Wang, and seminar participants at the University of Michigan, Resources for the Future, the University of Delaware, and the University of Massachusetts-Amherst 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. Much of the analysis in the first and third chapters was performed on the University of Michigan Flux High Performance Computing Cluster. I gratefully acknowledge funding from the Joseph L. Fisher Dissertation Fellowship from Resources for the Future and from the Michigan Institute for Teaching and Research in Economics.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

APPENDIX

ABSTRACT

Essays on Environmental and Natural Resource Economics

by

Evan Matthew Herrnstadt

Chair: Ryan M. Kellogg

This dissertation addresses issues in the economics of the environment and natural resources. The first chapter pertains to the inclusion of environmental objectives into contracts between the government and private firms. In particular, it considers the possibility that conservation restrictions may undermine the goal of fostering competition among private logging firms in timber auctions. Empirically, the policy is costly but is found to be borne primarily by the state without substantial competitive distortions. Importantly, that state can reduce the impact of the policy on revenues by setting reserve prices optimally.

The second chapter is joint work with Professor Erich Muehlegger. We use data on Google search activity related to climate change and shocks to local weather to demonstrate that unusual weather increases the salience of climate change as an issue. Further, we find that recent weather shocks have a significant effect on Congressional votes pertaining to environmental regulation.

The third chapter makes a methodological contribution to the analysis of auction data. Many auctions have a reserve price, below which the seller simply keeps the object of interest. However, it is typically taken for granted that the object will not be re-auctioned later. Empirical researchers

should account for the fact that bidders may respond to this possibility. I present a simple model of repeat auctions, discuss what information can be identified from bidding data, and provide a Monte Carlo simulation of an estimator that addresses this issue. Finally, I examine data on repeat auctions of logging contracts to see whether bidding behavior is consistent with the model. These data could provide a context in which to estimate the model and analyze counterfactual auction policies.

CHAPTER 1.

Conservation Versus Competition? Environmental Objectives in Government Contracting

1.1 Abstract

Government contracts with private firms increasingly incorporate environmental objectives or preferences for sustainable products and producers. At the same time, the government often solicits competitive bids to reduce the rents captured by the firms due to private information. In this paper, I show how environmental objectives can influence equilibrium contract bids through changes to firm costs, strategic bidding behavior, and bidder participation decisions. Using data from Michigan state logging contracts, I find that conservation objectives reduce bidder participation in the contract auctions by up to 35 percent and depress winning bids by up to 17 percent. To disentangle compliance costs from logger margins, I estimate a structural model of the auctions. Simulations based on the estimates imply that the policy imposes economically and statistically significant compliance costs. Loggers are able to completely pass these costs on to the government because compliance costs do not substantially affect the dispersion of private values. However, the use of optimal reserve prices partially mitigates the revenue disparity between more- and less-restricted contracts. Finally, loggers capture a larger share of total auction surplus for restricted contracts, indicating that the policy undermines the state's ability to harness competition to capture surplus.

1.2 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. I identify a large, negative effect on equilibrium bids and logger participation by exploiting the structure of the seasonal restrictions and a rich set of controls.[2] To quantify the relative importance of the cost of complying with the restrictions versus that of weakened competition, I estimate a structural model of the DNR-administered first-price timber auctions. The bidders' value distribution is parameterized so that it depends flexibly on the seasonal restrictions. I find that most of the effect on bids is driven by lower valuations, and that the loggers are able to pass nearly the full cost of compliance

[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, 2013), the authors distinguish between loggers and mills. Conversations with DNR employees suggest that there are few large-scale mills operating in this market.

through to the DNR.

The reduced-form analysis demonstrates that environmental objectives have a negative effect on bids and auction participation and suggests that differential competition could play an important role. 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. Moreover, these estimates are robust to controlling for the underlying reasons for the restrictions, which 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 of 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 to participate in the auction altogether.

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, 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. I also find that setting optimal reserve prices can close some of the revenue gap between more- and less-restricted sales.

The full passthrough 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 compliance costs of their competitors.

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/SIP 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 that promote consideration of environmental factors in federal procurement. In addition, the General Services Administration is considering incorporating bidders' greenhouse gas emissions as a criterion in the federal procurement process. 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 depends on the restrictions, and decompose the reduced-form effect to shine light on the importance of various mechanisms at play. Section 7 concludes the paper.

1.3 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.

1.3.1 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. One indicator of the competitiveness of this process is the share of total surplus that the firm manages to capture in 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, Coey, and Levin, 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, and 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.

1.3.2 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 survey includes infor-

[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.

mation 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. 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.[5]

Seasonal operating restrictions are added to timber contracts to help protect the ecological 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 1.1 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

[5]Further, contracts are transferred between firms in less than 1 percent of contracts.

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 1.1. 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.[6] 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 the Upper Peninsula and the Northern Lower Peninsula suggests that rental activity is limited. In 2009, firms used self-owned equipment for an average of 89 percent of their total operations. An average of

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

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 transact.

1.4 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, which shows that 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.

1.4.1 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).[7] Reserve prices are set using a formula based on recent prices paid for the same species in the same state forest.

Table 1.2 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

[7]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.

transaction.[8] The mean reserve price is \$61.6/MBF, or roughly \$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.[9] 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 is about 83 percent pulpwood by volume. Pulpwood comes from smaller-diameter trees and parts of trees and is sold to make paper products. Sawlogs, which are used to make lumber and utility poles, account for the other 17 percent. There is considerable variation in the sample, both in terms of the proportion of pulpwood versus sawlogs and the proportion of softwood (such as pine) versus hardwood (such as walnut). 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 roughly represent the 1st and 99th percentiles of these variables. I also drop all salvage sales, which specifically market fire-, wind-, or pest-damaged timber. Figure 1.2 presents the number of sample auctions that took place each quarter in the Lower and Upper Peninsulas. Clearly there is some cyclicality in both peninsulas – there are generally more sales held in spring and summer than in winter and fall. Given that the median contract lasts nearly 2.5 years, season of the auction itself should not play a large role in the value of the contract. Still, I include quarter-of-year dummy variables in my primary specification to control for this pattern and find that it does

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

[9] A management unit is usually a 2 to 3 county area: see Figure 1.3 for a map. This general approach is similar to existing work that analyzes entry in timber auctions, such as Roberts and Sweeting (2013) and Athey, Coey, and Levin (2013).

not have a large effect on my results.

1.4.2 Seasonal Restrictions

The DNR does not maintain a database variable that indicates when harvesting operations are allowed. However, the information is 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.[10] This variable captures the first-order driver of lost option value: the number of months during which a sale is inaccessible.[11]

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 1.4 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.

1.5 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 establish suggestive evidence that some of the effect on bids is driven by changes in competition.

[10] 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.

[11] 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. Less 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.

1.5.1 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).[12] 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. Table 1.3 further describes 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.

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 1.5.2, I also leverage the

[12] The full contents of the control vector can be seen in Table 1.4.

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 1.4 and 1.5, respectively).[13] 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 8 percent less revenue 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.[14] Figures 1.5 and 1.6 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.[15] The first four months of restrictions have zero marginal 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.

The effect on the number of bidders suggests that the differences in winning bids might be driven by differences in market thickness. Suppose that seasonal restrictions are more prevalent in areas with fewer active loggers. Then we would expect more restricted sales to attract a different number of bidders and level of bids even without a causal effect of restriction on bidding behavior.

[13] 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.

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

[15] The underlying regressions are analogous to Column (4) of Tables 1.4 and 1.5.

Figure 1.7 shows that a lower share of potential bidders actually participate in auctions for the more-restricted sales. Furthermore, in Figure 1.8 I replicate the regression in Figure 1.5, except that I control for a flexible polynomial of the number of potential bidders. The treatment effect does not change appreciably, indicating that the reduced-form results are not driven by differential market thickness.

1.5.2 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.[16] 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 second-order.

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 combinations of different restriction categories

[16] 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.

in the data, which is satisfied in this context. Table 1.6 presents the number of sales characterized by each pairwise combination of restriction categories. There are 32 unique pairwise combinations of restrictions, and 28 percent of the contracts in the sample are restricted for at least two reasons. 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 1.7: the treatment effect does increase slightly. In Figure 1.9, 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 Figure 1.5. The effect on the number of bidders in Figure 1.10 is also slightly different from the base spline specification in Figure 1.6. 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.

1.5.3 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 1.11 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. This result underlines the importance of directly modeling the reserve price and estimating participation costs.

1.6 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 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.

1.6.1 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 potential bidders that draw independent private values (IPV) v_i from a common distribution $F(v)$. Given this draw, each bidder decides whether to undertake a costly bid-preparation process, which costs K. Participants then submit bids in a first-price auction with public reserve price R, without 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 $b(v; N, R)$.[17] 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

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

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.[18]

The choice of entry model has important consequences for the model's implications, and the validity of the structural estimates.[19] The difference between the models can be understood be 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 fed-

[18] A third option is the affiliated signal model introduced empirically by Roberts and Sweeting (2013). 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.

[19] Roberts and Sweeting (2010) provide a relevant discussion, which I outline here.

eral 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.

1.6.2 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 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 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^* effects 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^*.[20] A different type will now be bidding the reserve price, so the equilibrium bids associated with types above v^* must also change in response.

1.6.3 Parameterizing the Model

I take a parametric approach to estimation similar to Roberts and Sweeting (2013), which allows me to incorporate rich observed and unobserved auction heterogeneity. The observed heterogene-

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

ity 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 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 considerable 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.[21]

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 $\bar{v}$.[22]

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 $\Gamma = \{\boldsymbol{\beta}, \boldsymbol{h}, \boldsymbol{\omega}\}$:

$$\mu_a \sim N(\beta^\mu X_a + h^\mu(MonthsRestricted_a), \omega^\mu)$$

$$\sigma_a \sim Weibull\ (exp[\beta^\sigma X_a + h^\sigma(MonthsRestricted_a)], \omega^\sigma)$$

$$K_a \sim Weibull\ (exp[\beta^K X_a + h^K(MonthsRestricted_a)], \omega^K),$$

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

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

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 regressions 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 1.4 and Figure 1.5. The X vector includes all of these covariates. I again specify $h(\cdot)$ as a restricted cubic spline in months of restrictions.

Informally, identification of the parameters comes from a combination of the data and the distributional assumptions.[23] 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. In particular, the level of the observed bids, the distances among the bids, and the distance from bids to the reserve price are informative. The participation costs are identified by the probability of bidder participation. 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.

1.6.4 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 Simu-

[23] Xu (2013) considers nonparametric identification and estimation of the selective entry model, but does not accommodate unobserved heterogeneity.

lated 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, $\theta = \{\mu, \sigma, K\}$.[24]

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:

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

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

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 $p(\theta|X_a, \Gamma)$.[25] 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-trivial amount of computing power and time. Therefore, I adopt an importance sampling approach: Ackerberg (2009) outlines the technique and demonstrates a num-

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

[25] 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.

ber of 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 (2013); Bhattacharya, Roberts, and Sweeting (2014); and Gentry and Stroup (2014).

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

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

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).

1.7 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.

1.7.1 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 1.8, 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 1.12 shows the effect of restrictions on the mean, standard deviation, and participation cost of a contract with otherwise typical observable characteristics. The mean falls by roughly 12 percent for the most-restricted sales. The standard deviation 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 there is less chance of a more heavily shaded bid being undercut because each bidder is more isolated in the distribution of private values. Because the expected markdown is roughly the expected gap between the first and second-highest valuations, it is difficult to predict the outcome from the first two moments only. 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.

For reference, Figure 1.13 presents the distributions of μ, σ, and K for a representative auction. 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. Table 1.9 demonstrates that I match the mean observed bid, government revenue, winning bid conditional on at least one bidder, and participation rate fairly closely. Figure 1.14 compares the densities of simulated and true winning bids across auctions, conditional on receiving a bid.

The fit conditional on observable characteristics is also very good. In Table 1.10, I show the results of a regression of all non-zero winning bids on covariates; the coefficients are extremely similar for simulated and true winning bids.[26] Because the restriction spline coefficients are difficult to interpret, I plot the spline functions in Figure 1.15. The simulated spline does a relatively good job matching the data. In effect, this relationship is the bid effect that I will be decomposing.

1.7.2 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.[27] The levels of various auction outcomes are shown in Figure 1.16. The government revenue from an

[26] This specification is analogous to the regression that generates Figure 1.5.

[27] In this section of the paper, I assume the DNR's reservation value for the sale is zero. The main results are similar if the DNR values sales at the observed reserve prices. Section 1.7.3 explores this further in the context of optimal reserve prices.

average (687 MBF) sale falls from $57,000 to $50,000 moving from 0 to 10 months of restrictions. Point estimates of firm surplus appear relatively constant across months of restrictions, except for the most-restricted sales.

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 reducing information rents, the auction performs slightly worse for more restricted sales. This represents a 5.2 to 6.6 percent increase in the firm's relative share of surplus.

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 top panel of Figure 1.17, 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, suggesting that 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 1.18, 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 12 percent. For sales restricted for 10 months, the point estimate is a decrease in firm surplus of 7 percent. However, this pointwise estimate is less precise: the 95 percent confidence interval is bounded by an increase in firm surplus of 17 percent and a decrease of 31 percent.

Why is there approximately full pass-through? One reason is that contract supply is modeled as inelastic. Timber is exogenously sold in the medium-run: the timber stands sold for harvest are those most necessary for forest management, subject to the minimum 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 result in 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.

1.7.3 Optimal Reserve Prices

The government could attempt to recover revenues lost due to seasonal restrictions using the auction reserve prices. To allow a consistent comparison across different levels of restriction intensity, I simulate the optimal reserve price separately at each level and compare outcomes.[28] 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

[28] For the purpose of this section, "optimal reserve price" refers to the reserve price that maximizes the government's expected auction surplus given a reservation value v_0, as is typical in the optimal auction literature. In this section, I consider different values for v_0, but consistently refer to the government's expected take (including some chance of receiving v_0) from the auction as "revenue". Of course, the true government objective function may balance a number of competing criteria.

the outcomes of each simulated auction over a fine grid of possible reserve prices.[29] I then average the outcomes across the 50 simulations to find the optimal reserve price for each observed auction.

In most circumstances, it is difficult to infer the auctioneer's true value for the object. Since the DNR has no in-house capacity for timber harvesting, in the preceding analysis, I assumed that the value of not contracting is zero. However, this assumption is less innocuous in the case of an optimal reserve price. Past papers have assumed that the reservation value is the reserve price (Paarsch, 1997; Roberts and Sweeting, 2013), or have considered a range of values between zero and the reserve price (Haile and Tamer, 2003; Roberts, 2013). This assumption impacts the effects of moving to an optimal reserve price. I perform the optimal reserve price analysis assuming that the government's reservation value (v_0) is either zero or the observed reserve price (R_{obs}).

When the government's valuation is zero, the simulations suggest that the DNR is typically setting reserve prices above the optimum. This is perhaps not surprising given that the reserve prices are typically benchmarked using scaled-down *winning bids* from recent comparable auctions. Table 1.11 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. 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. Because the current reserve prices are too high on average, setting an optimal reserve actually improves firm surplus as well as government revenue. In contrast, if $v_0 = R_{obs}$, the observed reserve price is too low, as shown in Table 1.12. The state gains from setting a higher reserve price, but this comes at the expense of firm surplus.

Optimal reserve prices can be used to blunt the revenue impact of seasonal restrictions. Figures 1.19 and 1.20 compare the share of revenues lost due to restrictions when using observed reserve prices versus optimal reserve prices if $v_0 = 0$ and $v_0 = R_{obs}$, respectively. That is, at each level of restrictions, the revenue is compared to the revenue from an unrestricted sale under same

[29] In practice, this grid ranges from 0 to 600 percent of the observed reserve price at 0.1 percent increments.

reserve price regime. When $v_0 = 0$, the gap between restricted and unrestricted sales narrows from 12.2 percent to 9.6 percent. When $v_0 = R_{obs}$, the effect is much smaller, but still positive. This suggests that for reasonable values of v_0, flexible reserve prices allow the state to run the auction in a way that best accommodates 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. Figures 1.21 and 1.22 present 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.

1.7.4 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 1.6.2 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 with 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 1.13 summarizes my approach.

The decomposition in Figure 1.23 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.

1.7.5 Policy Implications

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. The mean compliance costs for a contract amount to \$918. Taken across the entire sample, the compliance costs amount to \$4.8 million. The mean revenue lost due to restrictions is \$874 per sale. Overall, the seasonal restrictions program reduced Michigan DNR timber revenue by \$4.5 million, or roughly 1.4 percent of the \$313 million in timber sale revenues collected over the 10-year period. However, much of the burden is due to 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, 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.

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. However, simulated optimal reserve prices suggest that the DNR could increase its returns from all contracts, and that the improvements are even more pronounced for the more-restricted sales. If the DNR's reservation values for the contracts are sufficiently high, this optimal reserve policy would reduce firm cost passthrough.

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. This suggests that the competitive performance of the auction is undermined somewhat by the restrictions.

1.8 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. 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 not inform the political economy discussion. They also could not assess the extent to which the auction leverages the benefits of competition for the state. 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.

1.9 Figures

Figure 1.1: 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 1.2: Sales per quarter, by peninsula

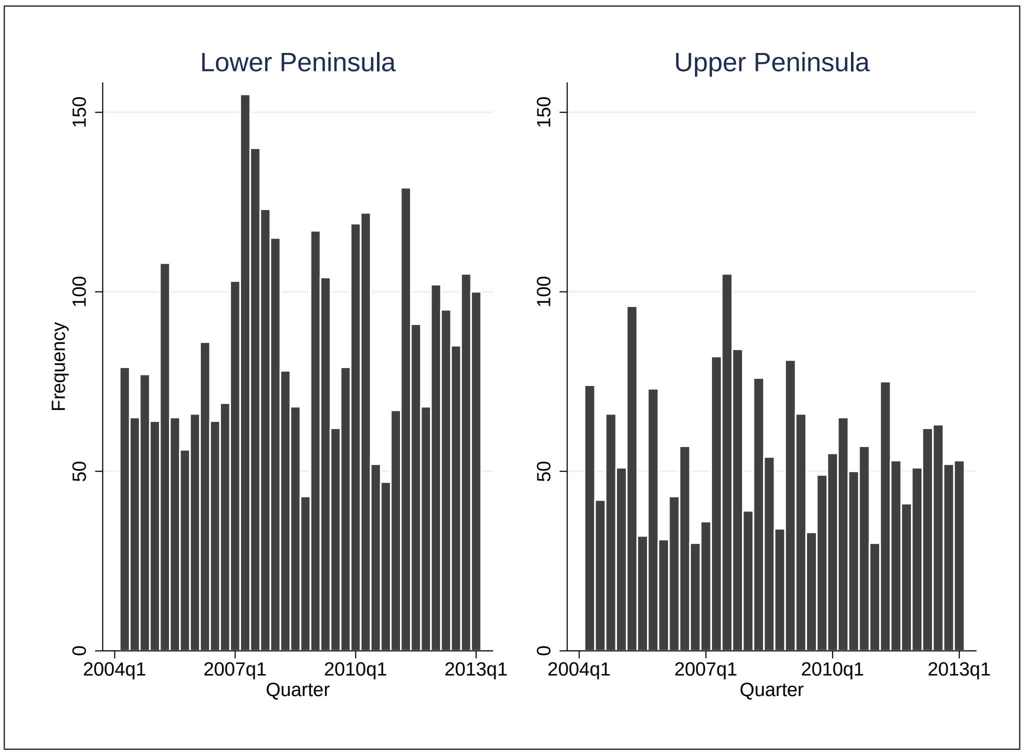

This displays the number of sample contracts auctioned in each quarter from 2004Q2-2013Q1.

Figure 1.3: Michigan DNR Forest Management Units

Source: Michigan DNR. I define a logger to be a potential bidder if they bid on a contract in the same quarter-management unit combination as the contract of interest.

Figure 1.4: Distribution of months of restrictions

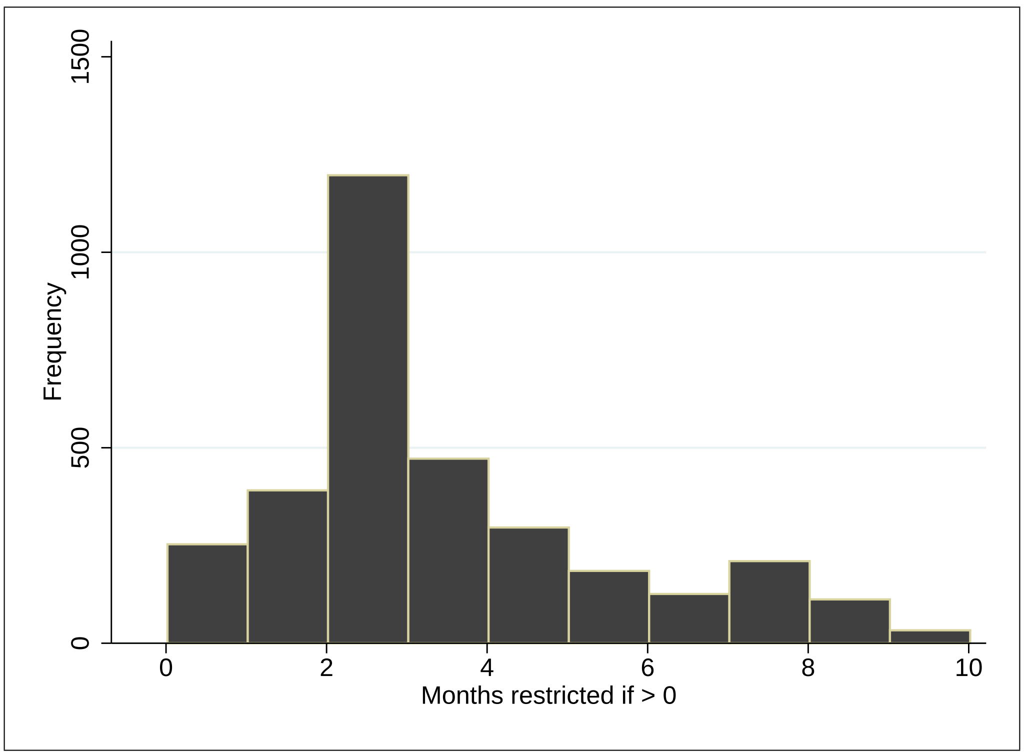

This histogram excludes 1933 auctions with no restrictions.

Figure 1.5: Log winning bid

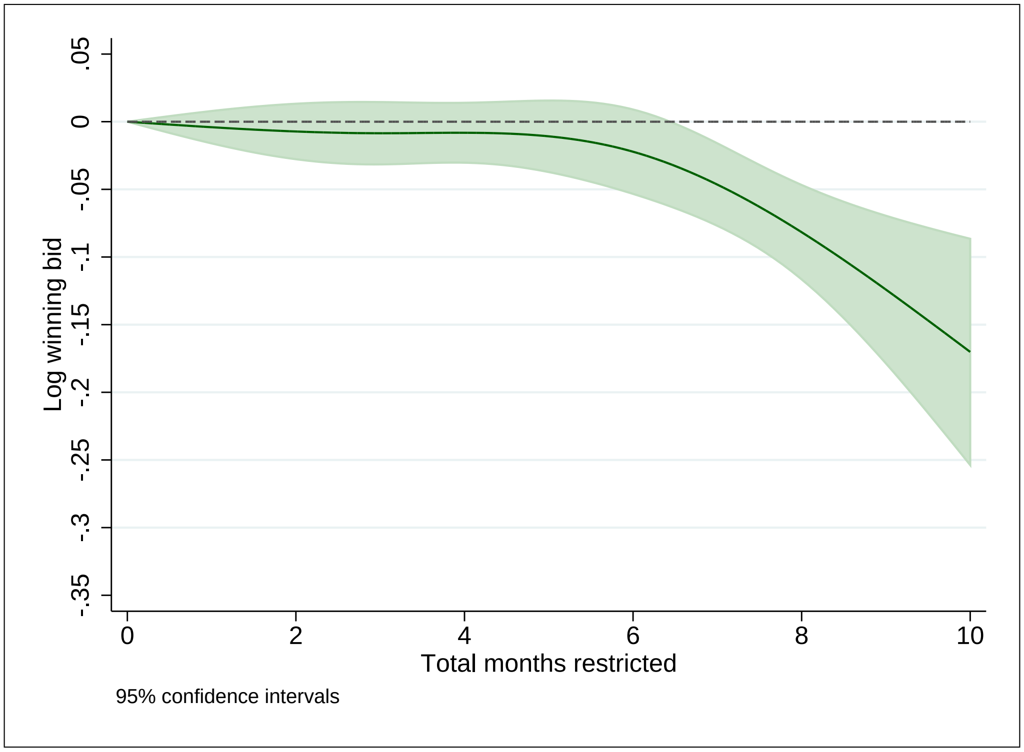

Note: Figure displays the results of regressing the logarithm of the winning bid on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 1.4. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions.

Figure 1.6: Number of bidders

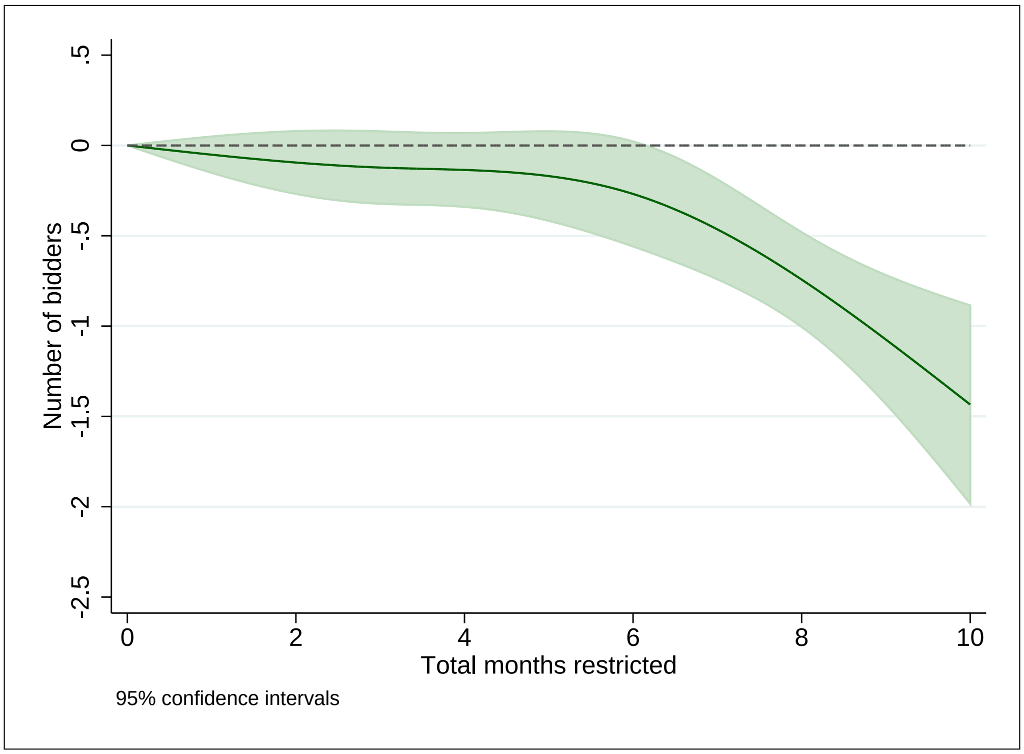

Note: Figure displays the results of regressing the number of bidders (up to 10) on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 1.5. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 5207 auctions.

Figure 1.7: Participation rate

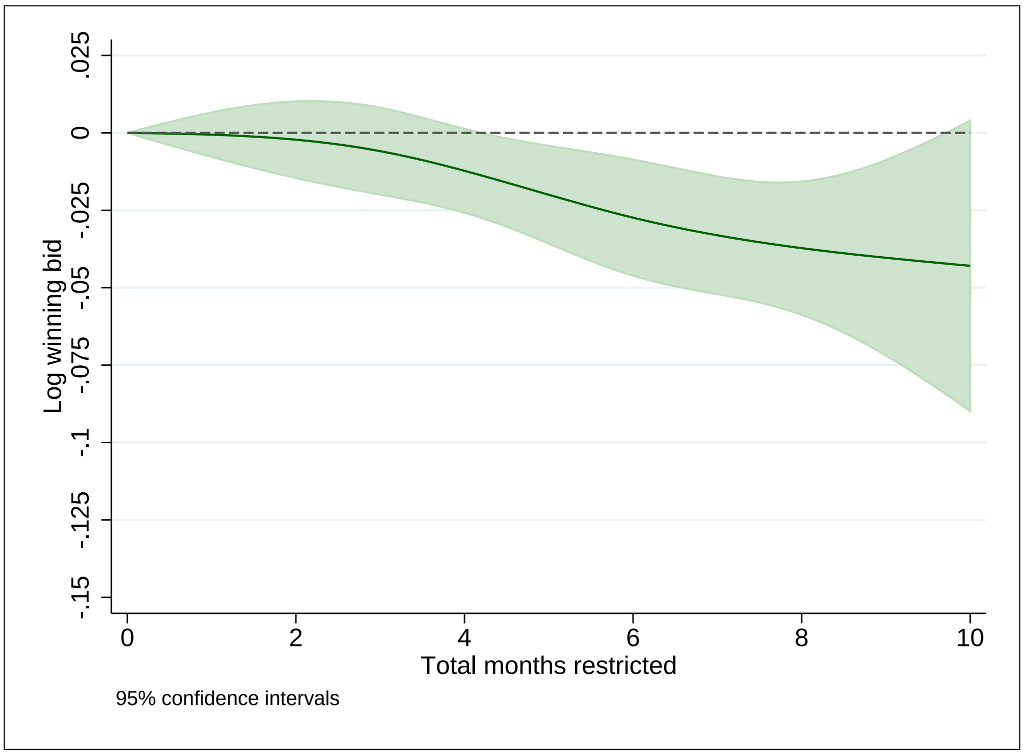

Note: This figure regresses the bidder participation rate (number of bids received/number of potential bidders) on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. A potential bidder is defined as any logger than bids in an auction in the same calendar quarter and management unit as the auction of interest. The specification is analogous to column (4) of Table 1.5. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 5207 auctions.

Figure 1.8: Log winning bid, controlling for potential bidders

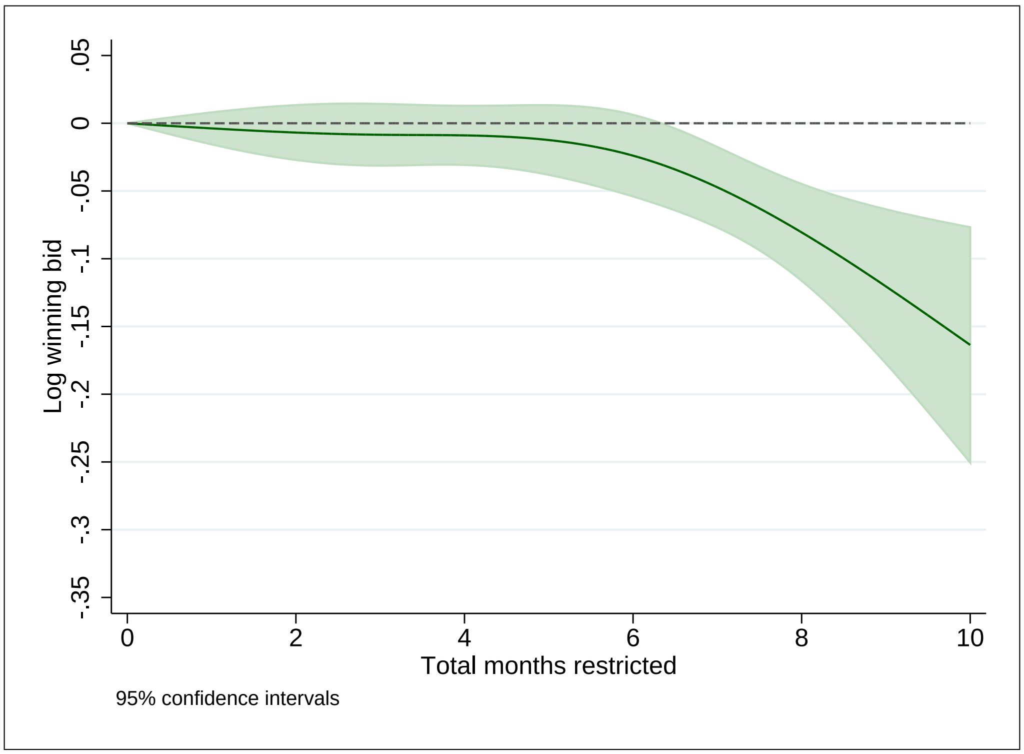

Note: Figure displays the results of regressing the logarithm of the winning bid on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 1.4, except that it includes a fifth-degree Chebyshev polynomial in potential bidders. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions.

Figure 1.9: Log winning bid, controlling for restriction categories

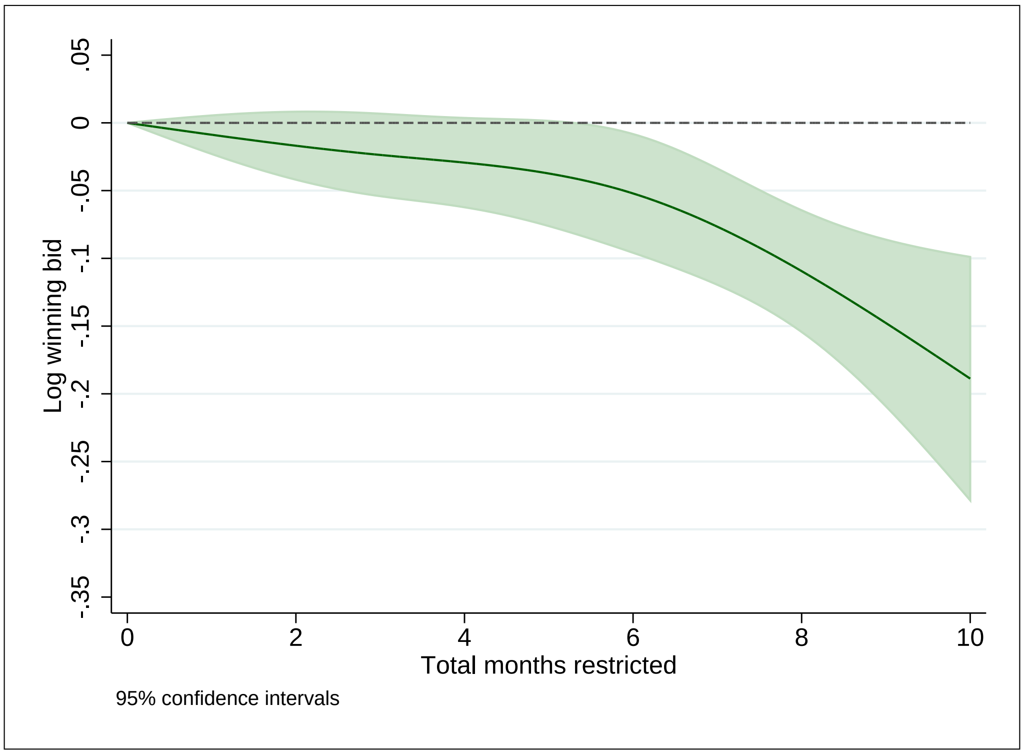

Note: Figure displays the results of regressing the logarithm of the winning bid on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 1.7, and includes dummies for restriction categories. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions.

Figure 1.10: Number of bidders, controlling for restriction categories

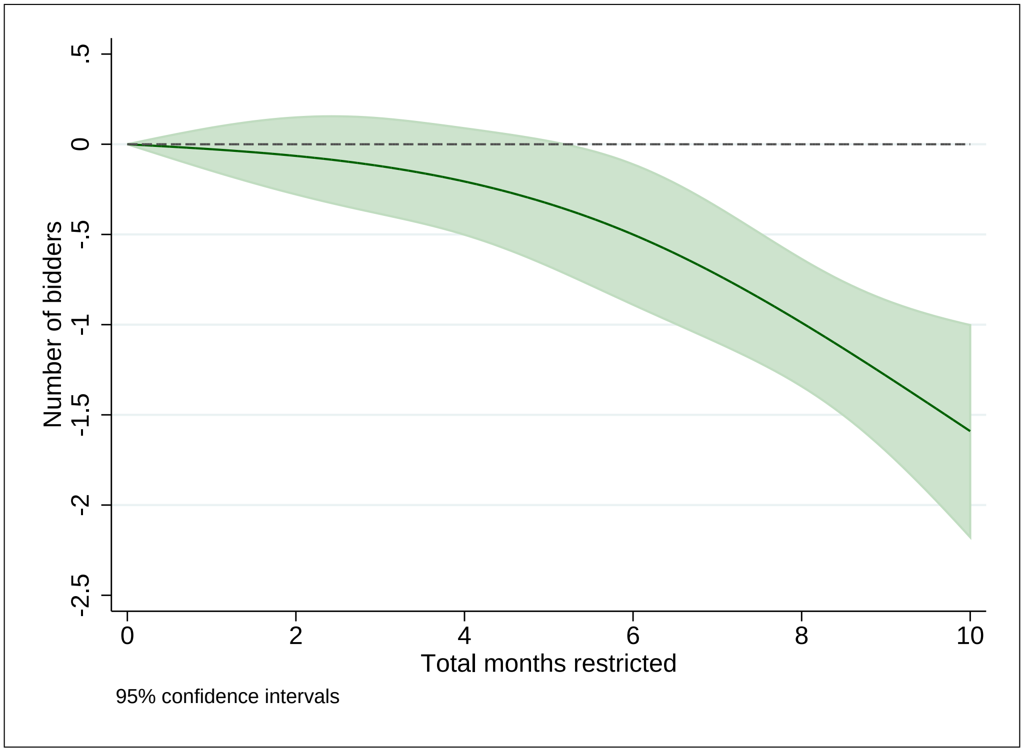

Note: Figure displays the results of regressing the number of bidders (up to 10) on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 1.7, and includes dummies for restriction categories. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 5207 auctions.

Figure 1.11: Log winning bid, controlling for number of bidders

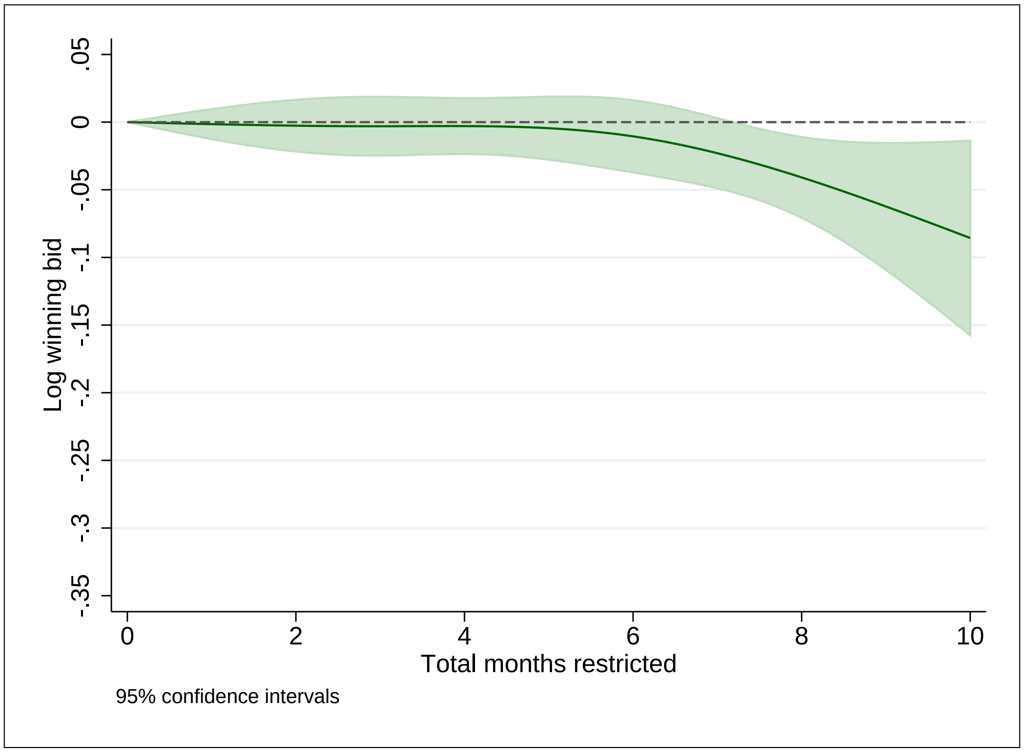

Note: Figure displays the results of regressing the logarithm of the winning bid on a cubic spline in months restricted with knots at 0, 3, 6, and 10 months. The specification is analogous to column (4) of Table 1.7, except that it includes dummies for the number of bids received. The shaded area is the 95% confidence interval implied by standard errors that are clustered by county-year. Sample size is 4750 auctions.

Figure 1.12: Moments of value distributions at mean X values

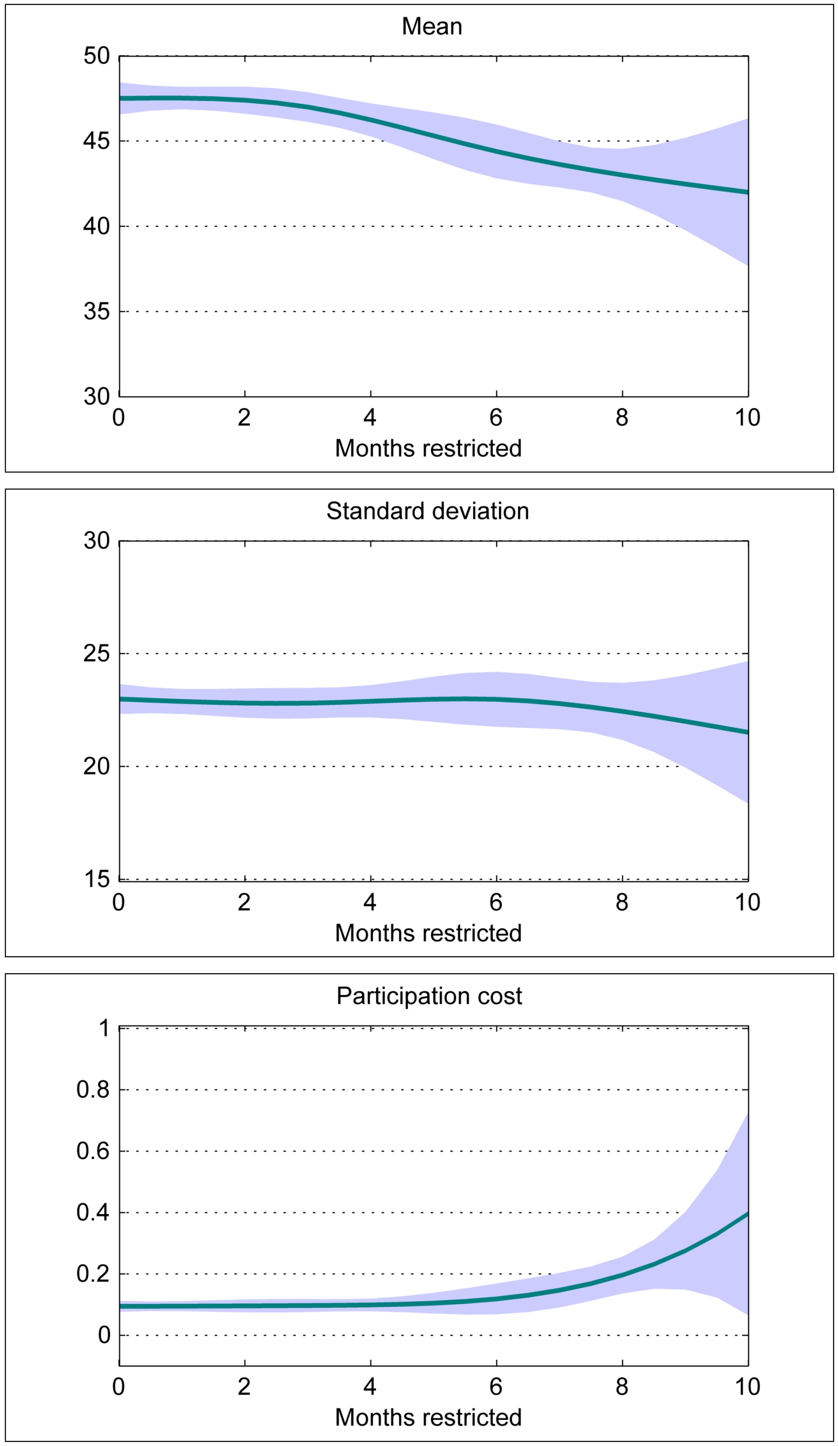

Note: These are the means and standard deviations of the truncated log-normal value distribution and the participation cost implied by the structural estimates in Table 1.8. They are evaluated for an auction with mean covariate values as the number of months during which the sale is restricted varies. 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 1.13: Distribution of auction-specific valuation parameters

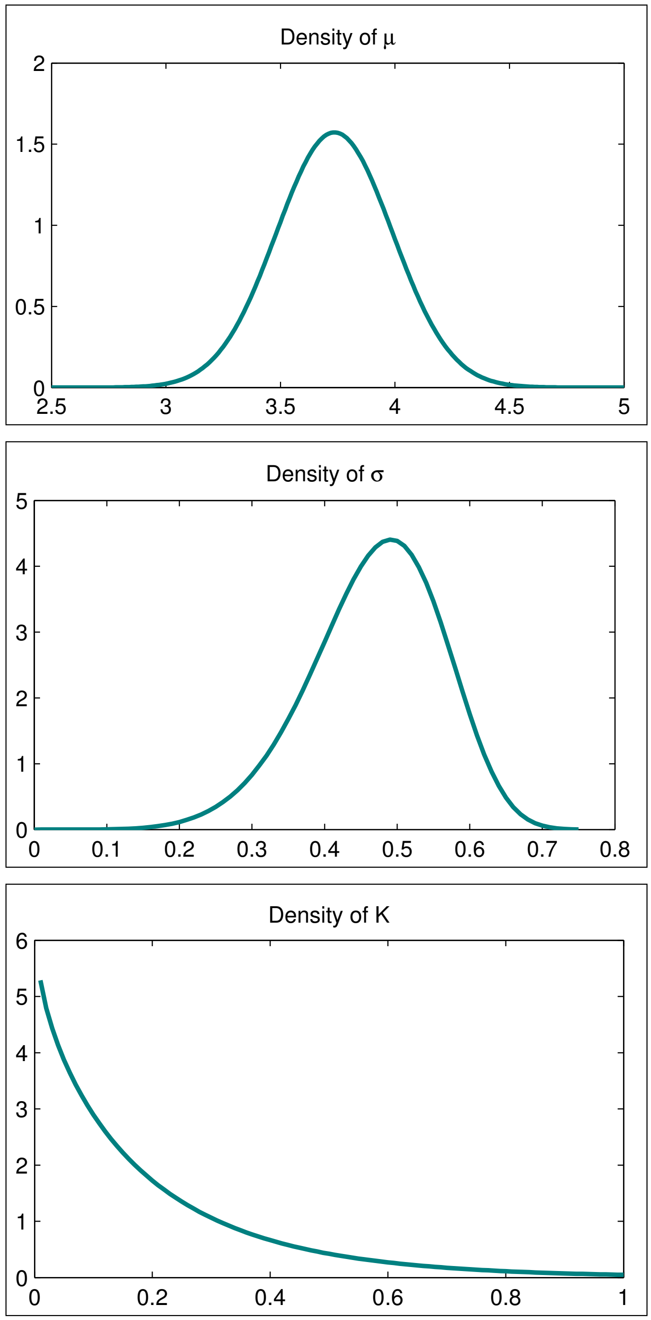

Note: These are the distributions of μ_a, σ_a, and K_a implied by the structural estimates in Table 1.8 for an auction with mean values of the auction covariates. Recall that μ_a has a normal distribution, while σ_a and K_a have Weibull distributions.

Figure 1.14: Comparing winning bids in all auctions receiving bids

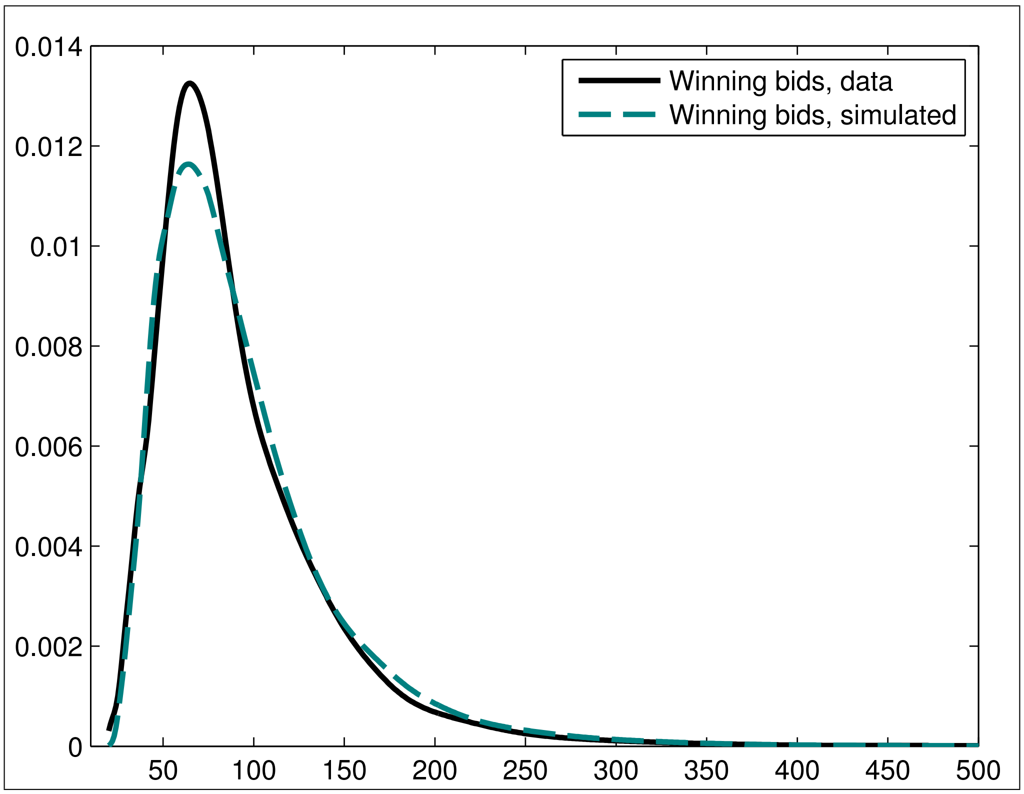

Note: These figures compare the outcomes of auctions simulated according to the distributions implied by Table 1.8 to those in the data. I simulate 10 auctions per data observation.

Figure 1.15: Splines from log winning bid regressions

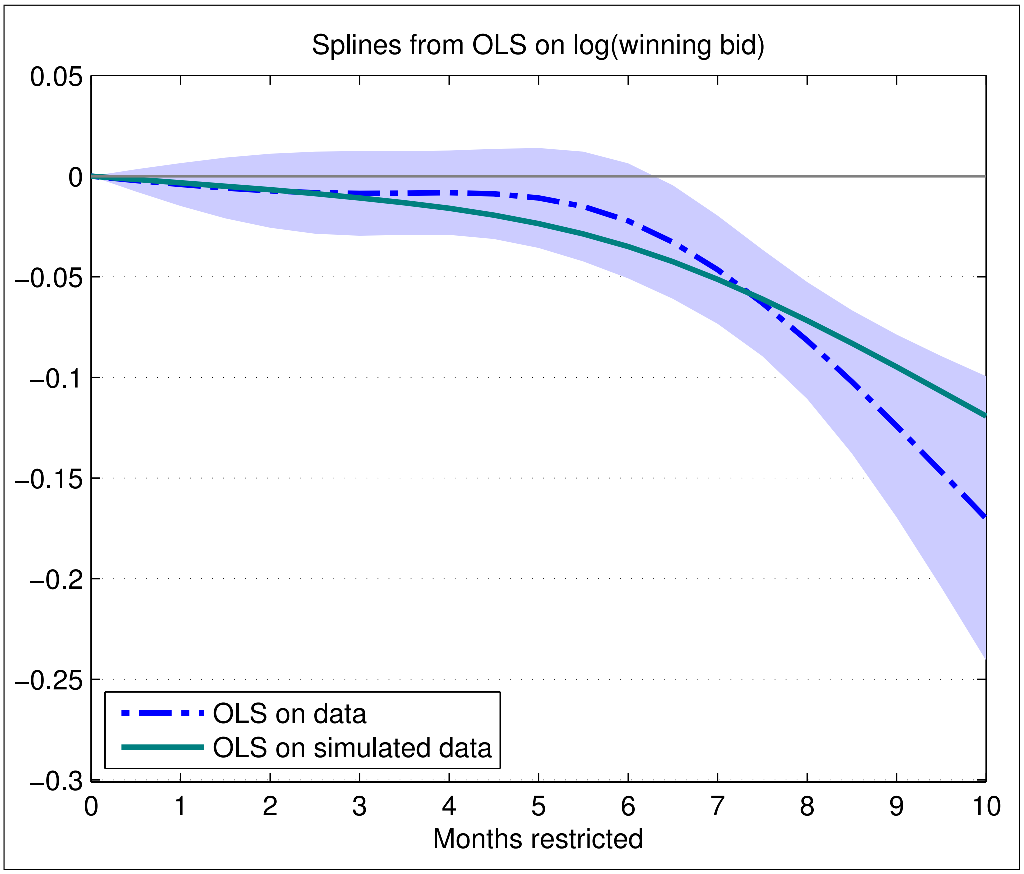

Note: This figure compares the estimated effect of restrictions in an OLS regression of ln(winning bid) using the real data and the data simulated using the structural estimates in Table 1.8. 95% confidence interval is for the data (reduced-form) estimate, and are OLS standard errors to be conservative regarding model fit.

Figure 1.16: Mean auction outcomes, by level of seasonal restrictions

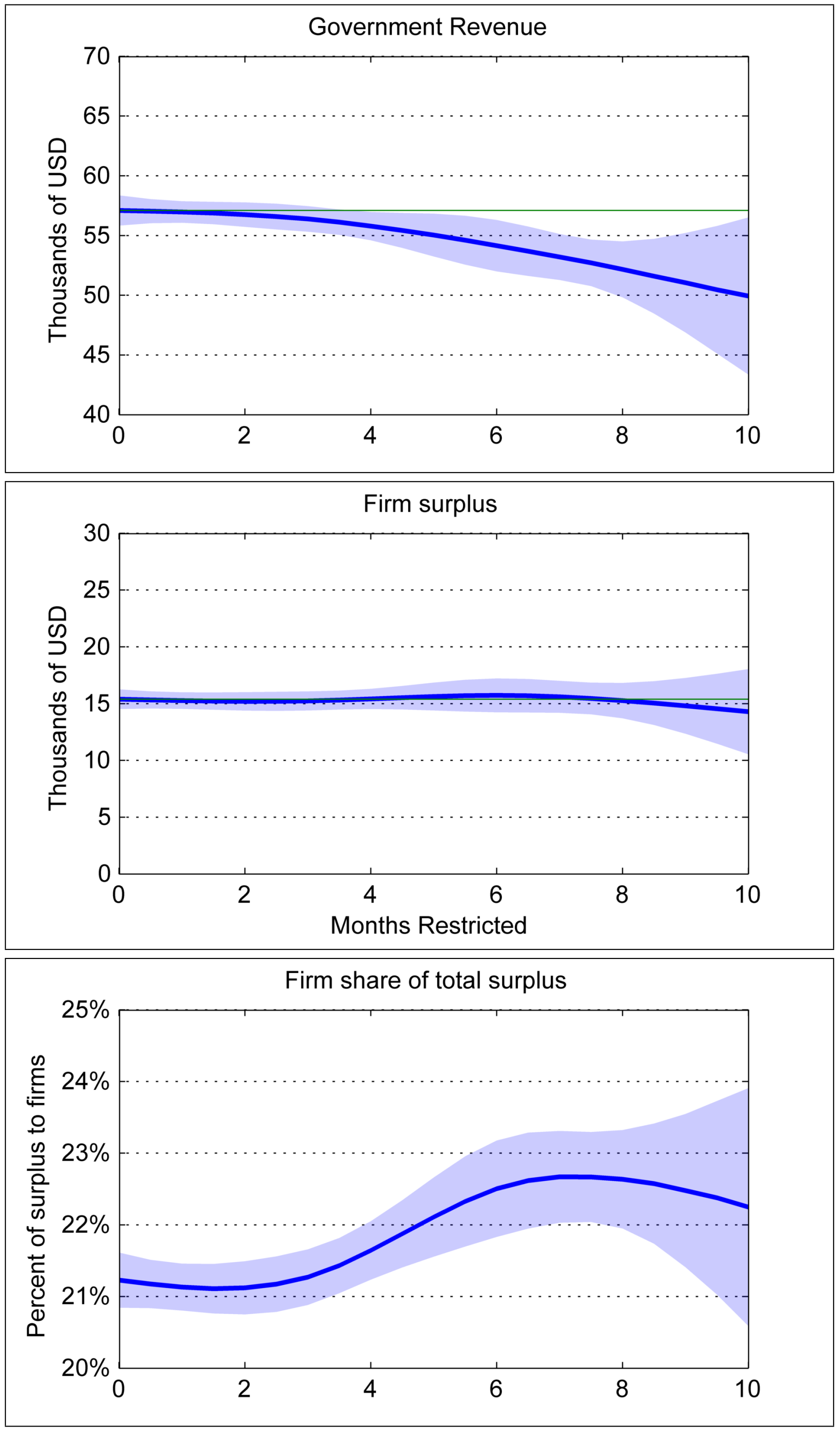

Note: These are the mean outcomes of auctions simulated according to the distributions implied by Table 1.8. I simulate 10 auctions per data observation, holding restrictions fixed at zero months. I repeat this at each level of restrictions $\{0.5, 1, \ldots, 9.5, 10\}$. 95% confidence intervals are derived using standard errors from 100 bootstrap replications.

Figure 1.17: Changes in mean auction outcomes, by level of seasonal restrictions

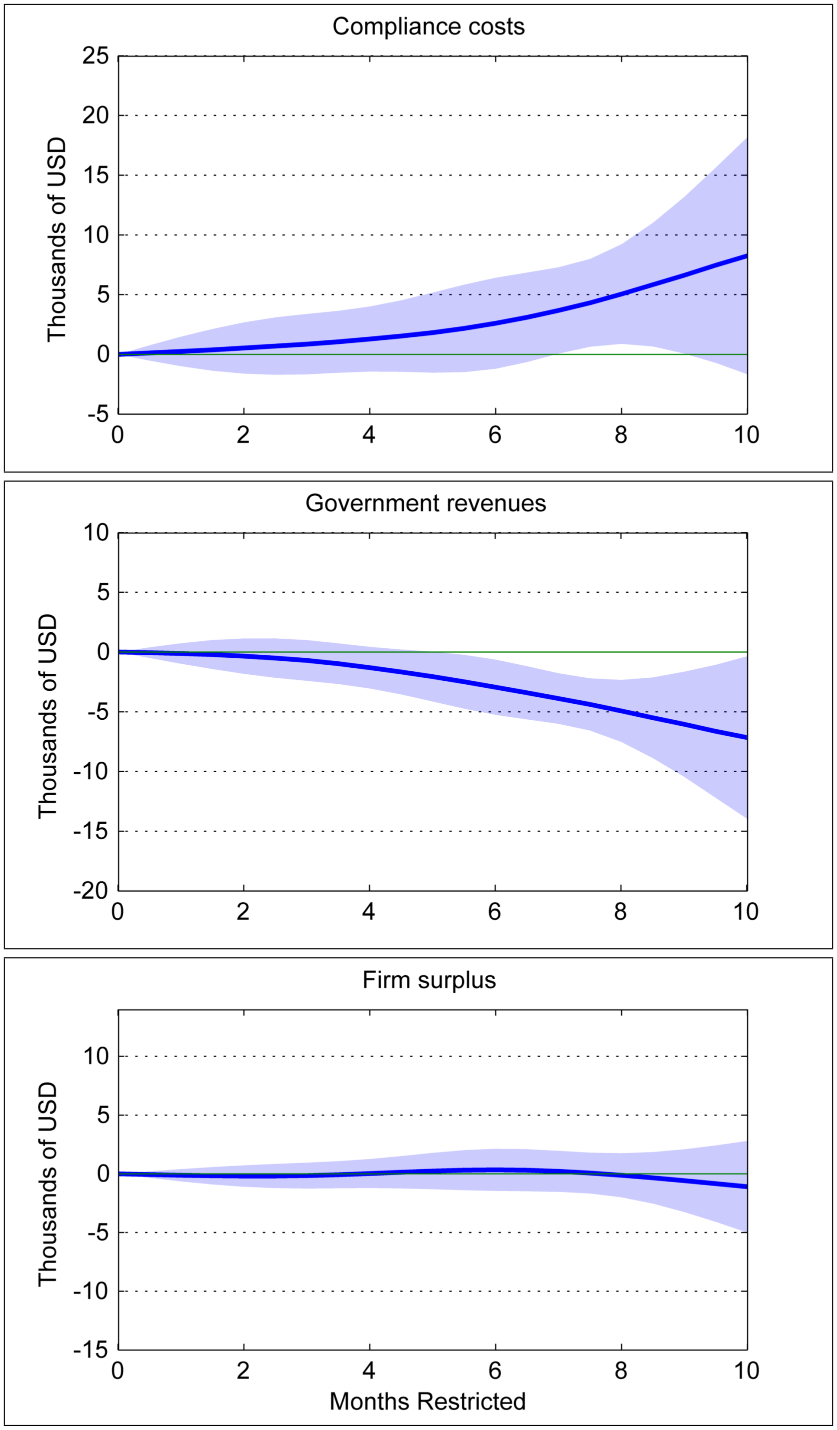

Note: These are the level changes in auction outcomes as seasonal restrictions vary. Compliance costs are the (negative) change in the valuation of the winning bidder, plus any change in the total participation costs incurred by bidders. I simulate 10 auctions per data observation, holding restrictions fixed at zero months: these simulated outcomes are normalized to 0 for the purposes of the figure. 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 1.18: Relative changes in surplus

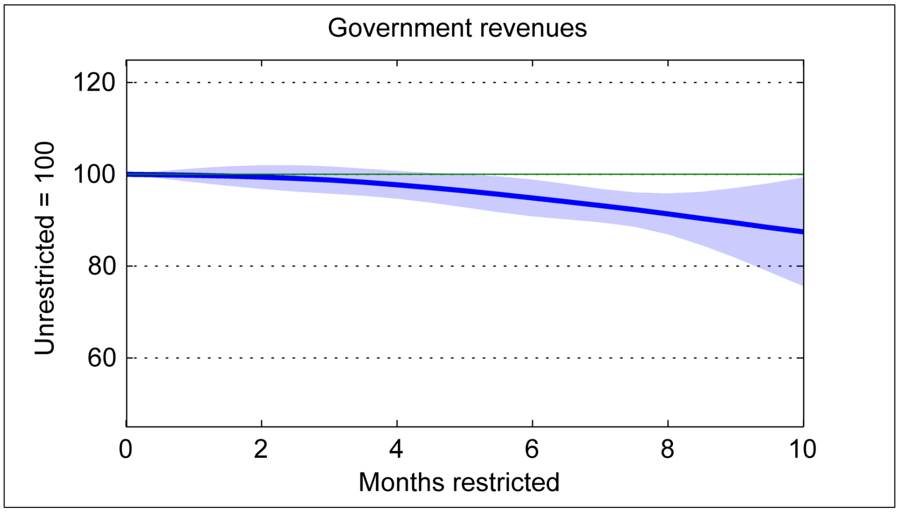

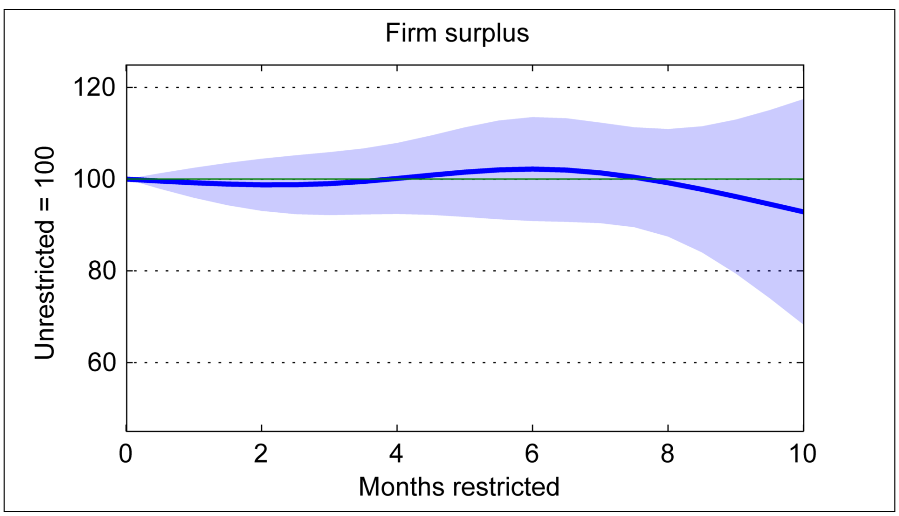

Note: These are the relative changes in surplus measures as seasonal restrictions vary. I simulate 10 auctions per data observation, holding restrictions fixed at zero months: these simulated outcomes are normalized to 100 for the purposes of the figure. I repeat this at each level of restrictions $\{0.5, 1, \ldots, 9.5, 10\}$. 95% confidence intervals are derived using standard errors from 100 bootstrap replications.

Figure 1.19: Relative impact of restrictions under different reserve price regimes, $v_0 = 0$

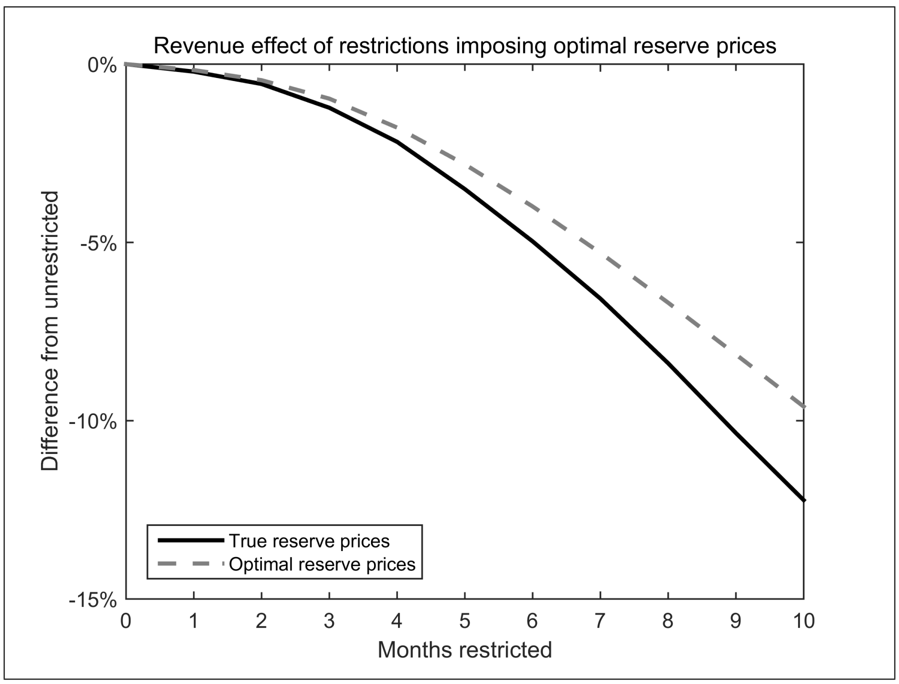

Note: I assume the government's value of keeping the contract is equal to zero. These are the relative changes in government revenues as seasonal restrictions vary with optimal reserve prices. I simulate 50 auctions for 500 randomly-drawn data observations, holding restrictions fixed at zero months.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

Figure 1.20: Relative impact of restrictions under different reserve price regimes, $v_0 = R_{obs}$

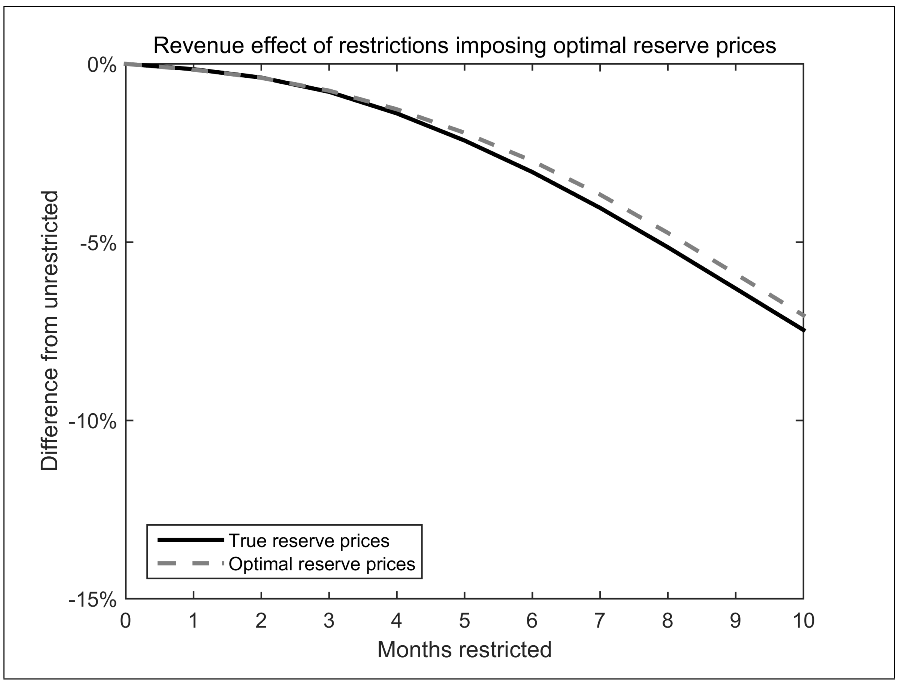

Note: I assume the government's value of keeping the contract is equal to the observed reserve price. These are the relative changes in government revenues as seasonal restrictions vary with optimal reserve prices. I simulate 50 auctions for 500 randomly-drawn data observations, holding restrictions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

Figure 1.21: Revenues across reserve price regimes, $v_0 = 0$

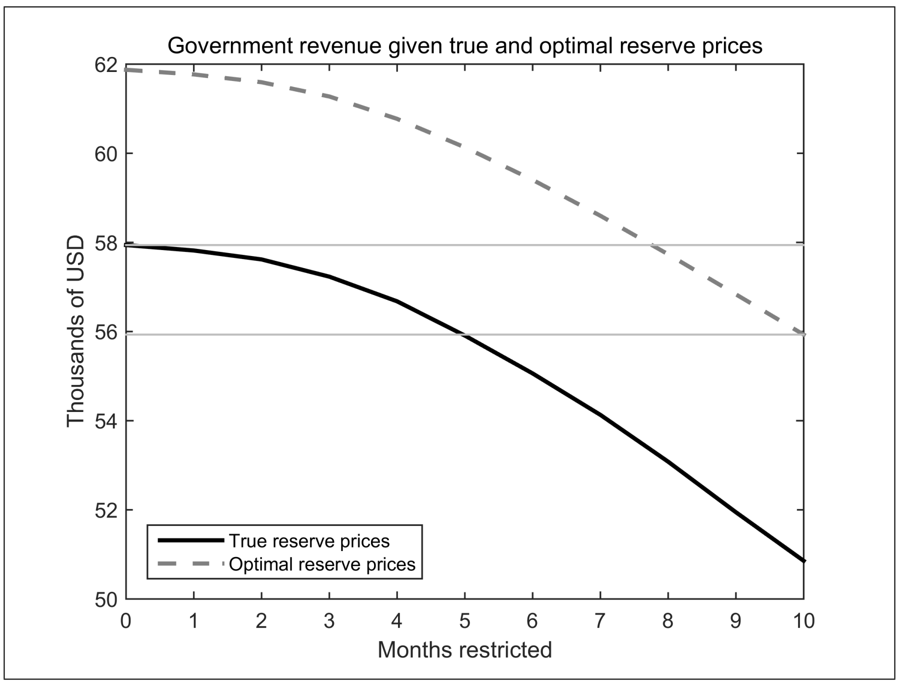

Note: I assume the government's value of keeping the contract is equal to zero. These are the average government revenues as seasonal restrictions vary with and without optimal reserve prices. I simulate 50 auctions for 500 randomly-drawn data observations, holding restrictions fixed at zero months. I repeat this at each level of restrictions $\{1, \ldots, 10\}$.

Figure 1.22: Revenues across reserve price regimes, $v_0 = R_{obs}$

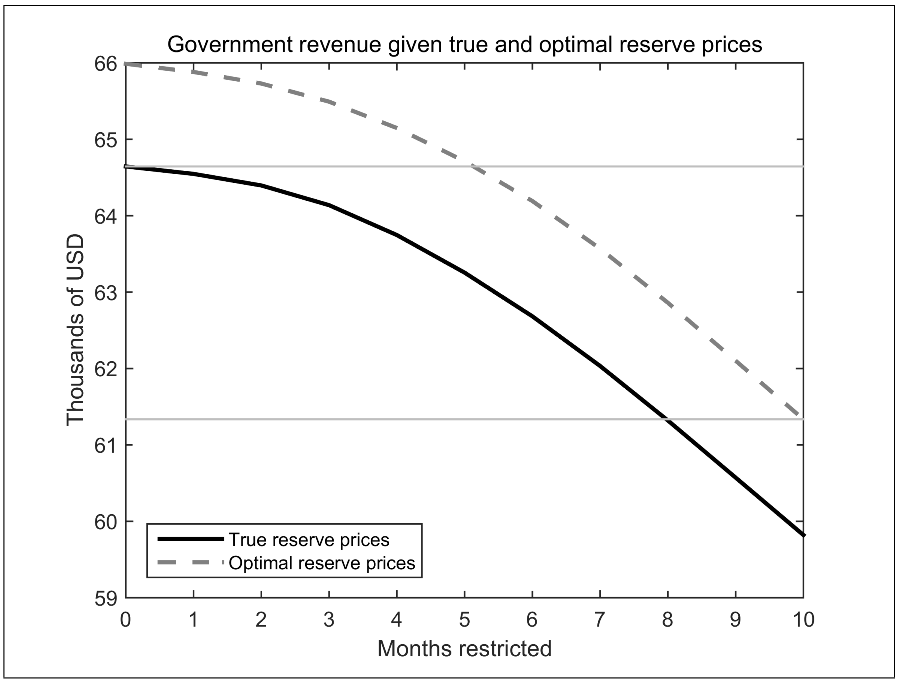

Note: I assume the government's value of keeping the contract is equal to the observed reserve price. These are the average government revenues as seasonal restrictions vary with and without optimal reserve prices. I simulate 50 auctions for 500 randomly-drawn data observations, holding restrictions fixed at zero months. I repeat this at each level of restrictions $\{1, \ldots, 10\}$.

Figure 1.23: Decomposition of equilibrium bid effect: Mechanisms

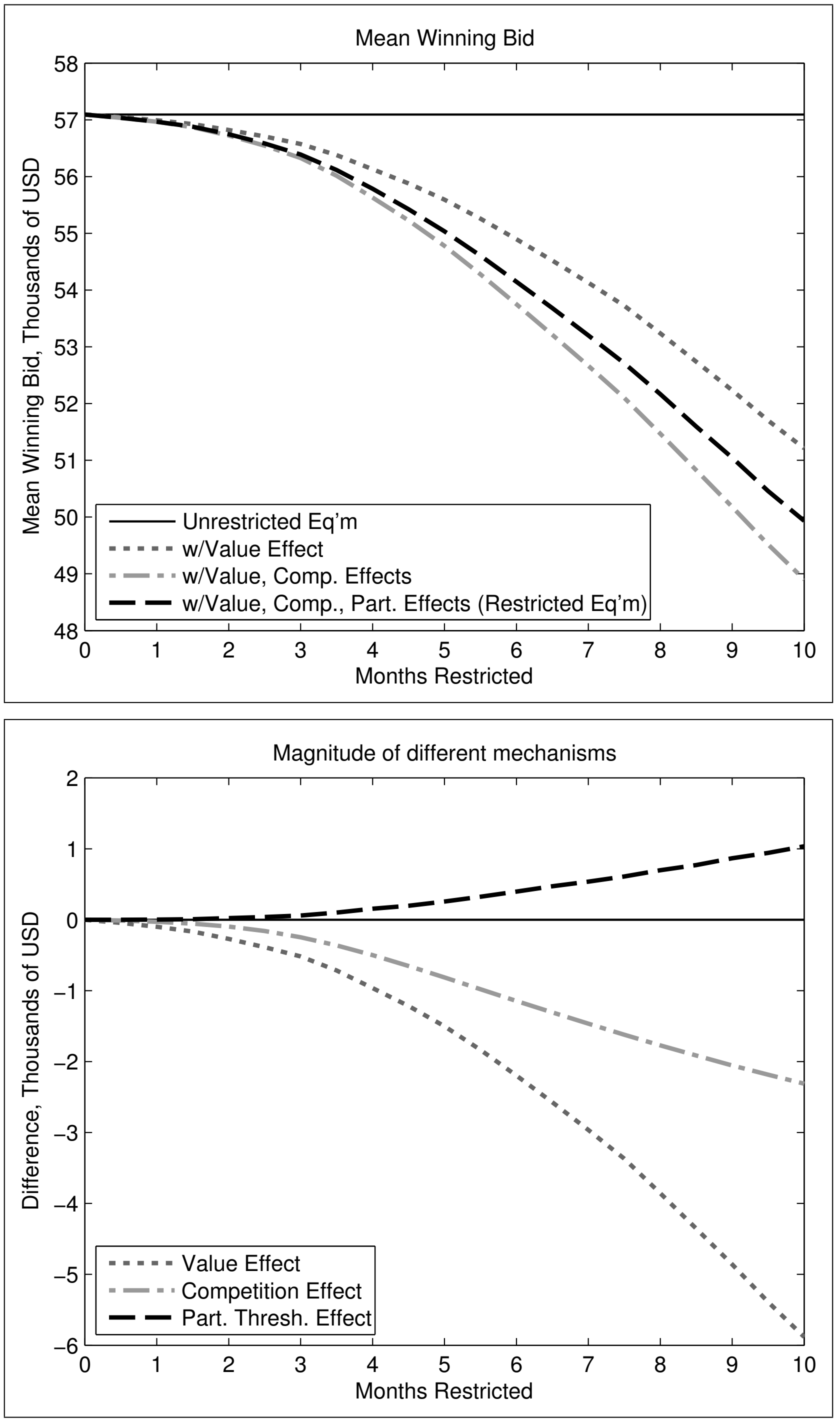

Note: This figure decomposes of the full effect of seasonal restrictions on bidding. The top panel shows the mean winning bid as the number of months restricted vary, as the value, competition, and participation threshold effects are iteratively added in. The bottom panel shows the changes in the mean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

1.10 Tables

Table 1.1: 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 1.2: Summary statistics

	Obs	Mean	Std. Dev.	P10	P25	P50	P75	P90
<u>Auction Outcomes</u>								
Winning bid ($/MBF)	4750	92.2	50.2	46.8	60.2	79.4	111.3	151.2
Number of bidders	5207	3.9	2.7	1	2	3	6	8
Participation rate	5207	0.2	0.2	0	0.1	0.2	0.3	0.5
<u>Contract Characteristics</u>								
Reserve price ($/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
Share soft pulpwood	5207	0.5	0.3	0	0.2	0.6	0.8	0.9
Share soft sawlogs	5207	0.1	0.1	0	0	0	0.1	0.3
Share hard pulpwood	5207	0.3	0.3	0	0	0.2	0.6	0.9
Share hard sawlogs	5207	0.1	0.1	0	0	0	0.1	0.2
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
Contract length (years)	4878	2.5	0.6	1.9	2.1	2.4	3	3.3

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 1.3: DNR Cost Factor Criteria

Factor	**Determinants**
Felling & Bucking (cutting trees into logs)	Logs per tree, density of underbrush and uncut trees
Skidding (moving logs) Problems	Slope and rockiness of terrain, wetness of area
Skidding Distance	
Road Maintenance and Minor Construction	Filling in damaged roads or considerable snow plowing
Distance to High Quality Road	Distance, terrain, forest road quality
Distance to Mill or Processing Plant	
Quantity	MBF/acre
Quality	Tree quality grade

Note: adapted from DNR internal Information Circular #4207

Table 1.4: Linear regressions, log winning bid

VARIABLES	(1)	(2)	(3)	(4)	(5)
Months restricted	-0.010***	-0.010***	-0.010***	-0.008***	-0.007***
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
Share Softwood: sawlogs	2.016***	1.938***	1.721***	1.682***	1.703***
	(0.118)	(0.100)	(0.093)	(0.087)	(0.087)
Share Hardwood: sawlogs	1.542***	1.407***	1.181***	1.207***	1.270***
	(0.046)	(0.047)	(0.050)	(0.046)	(0.047)
Share Hardwood: pulpwood	0.403***	0.196***	0.175***	0.151***	0.174***
	(0.033)	(0.031)	(0.032)	(0.029)	(0.026)
Upper peninsula	0.347***	0.297***	0.296***	0.303***	
	(0.030)	(0.027)	(0.028)	(0.020)	
DNR cost factors	1.122***	1.106***	1.115***	0.804***	0.908***
	(0.082)	(0.078)	(0.068)	(0.061)	(0.055)
Log acres		0.052***	0.037***	0.040***	0.046***
		(0.010)	(0.009)	(0.008)	(0.007)
Species-product HHI		0.635***	0.678***	0.694***	0.683***
		(0.042)	(0.042)	(0.038)	(0.036)
Percent bid species		0.733***	0.812***	0.695***	0.607***
		(0.095)	(0.089)	(0.070)	(0.066)
Constant	3.062***	2.023***	1.907***	2.369***	2.315***
	(0.063)	(0.099)	(0.096)	(0.086)	(0.090)
Observations	4,750	4,750	4,750	4,750	4,750
R-squared	0.414	0.521	0.574	0.645	0.663
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 1.5: Linear regressions, number of bidders

VARIABLES	(1)	(2)	(3)	(4)	(5)
Months restricted	-0.113***	-0.111***	-0.106***	-0.077***	-0.102***
	(0.022)	(0.022)	(0.021)	(0.017)	(0.016)
Share Softwood: sawlogs	-2.092***	-2.687***	-2.475***	-2.248***	-1.450***
	(0.645)	(0.643)	(0.675)	(0.550)	(0.526)
Share Hardwood: sawlogs	-0.351	-0.760*	-1.380***	-1.404***	-1.154***
	(0.452)	(0.445)	(0.464)	(0.431)	(0.432)
Share Hardwood: pulpwood	-1.507***	-1.990***	-2.094***	-2.105***	-1.781***
	(0.225)	(0.241)	(0.253)	(0.227)	(0.222)
Upper peninsula	1.146***	0.944***	1.053***	0.964***	
	(0.216)	(0.214)	(0.225)	(0.167)	
DNR cost factors	2.065***	2.076***	2.056***	0.770*	0.918**
	(0.486)	(0.475)	(0.463)	(0.411)	(0.416)
Log acres		0.516***	0.593***	0.587***	0.677***
		(0.073)	(0.070)	(0.060)	(0.057)
Species-product HHI		1.109***	0.757**	0.785***	0.995***
		(0.281)	(0.295)	(0.248)	(0.236)
Percent bid species		2.681***	2.141***	1.894***	1.517***
		(0.658)	(0.620)	(0.548)	(0.533)
Constant	2.891***	-1.894***	-1.285*	0.847	0.797
	(0.364)	(0.717)	(0.727)	(0.724)	(0.721)
Observations	5,207	5,207	5,207	5,207	5,207
R-squared	0.084	0.114	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 1.6: Pairwise combinations of restrictions

	Oak Wilt	Bark Slip	Priv. Prop.	Winter Rec.	Other Rec.	Regrowth	Soil/ Wetness	Misc. Other	Sole Restr.
Oak Wilt	.	.	.	.	.	.	.	.	494
Bark Slip	207	.	.	.	.	.	.	.	703
Priv. Prop.	0	8	.	.	.	.	.	.	4
Winter Rec.	29	116	1	.	.	.	.	.	58
Other Rec.	12	21	2	31	.	.	.	.	13
Regrowth	22	57	1	39	10	.	.	.	64
Soil/Wetness	65	268	15	34	12	78	.	.	234
Misc. Other	4	2	0	1	0	2	0	.	129
Wildlife	17	69	3	14	9	17	26	2	91

Note: 1471 out of 5207 sales (28%) in the sample have more than one major type of restriction.

Table 1.7: Linear regressions controlling for restriction categories

VARIABLES	(1) Ln(win bid)	(2) Ln(win bid)	(3) # Bidders	(4) # Bidders
Months restricted	-0.008***	-0.013***	-0.077***	-0.123***
	(0.002)	(0.003)	(0.017)	(0.023)
Share Softwood: sawlogs	1.682***	1.691***	-2.248***	-2.126***
	(0.087)	(0.087)	(0.550)	(0.551)
Share Hardwood: sawlogs	1.207***	1.190***	-1.404***	-1.446***
	(0.046)	(0.046)	(0.431)	(0.438)
Share Hardwood: pulpwood	0.151***	0.147***	-2.105***	-2.119***
	(0.029)	(0.029)	(0.227)	(0.225)
Upper peninsula	0.303***	0.297***	0.964***	0.936***
	(0.020)	(0.020)	(0.167)	(0.165)
DNR cost factors	0.804***	0.798***	0.770*	0.730*
	(0.061)	(0.061)	(0.411)	(0.410)
Log acres	0.040***	0.038***	0.587***	0.575***
	(0.008)	(0.008)	(0.060)	(0.060)
Species-product HHI	0.694***	0.695***	0.785***	0.790***
	(0.038)	(0.038)	(0.248)	(0.249)
Percent bid species	0.695***	0.690***	1.894***	1.902***
	(0.070)	(0.069)	(0.548)	(0.550)
Constant	2.369***	2.390***	0.847	0.938
	(0.086)	(0.086)	(0.724)	(0.733)
Observations	4,750	4,750	5,207	5,207
R-squared	0.645	0.648	0.282	0.287
Restriction Categories	-	X	-	X

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

Table 1.8: 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(βX) and the shape parameter is ω. Estimation is based on 5,207 auctions, which receive a total of 20,502 non-zero bids.

Table 1.9: 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 1.10: Model fit: Winning bid OLS

Covariate	Data	Simulation
Constant	2.37	2.37
Restr. Spline 1	-0.05	-0.04
Restr. Spline 2	0.18	-0.06
Restr. Spline 3	-1.34	-0.24
Upper Peninsula	0.30	0.31
Ln (Acres)	0.04	0.04
Softwood Sawlogs	1.66	1.77
Hardwood Pulp	0.15	0.17
Hardwood Sawlogs	1.21	1.21
Pct. Bid Species	0.69	0.66
Species-Product HHI	0.69	0.71
DNR Cost Factors	0.80	0.87

Notes: Dependent variable is the log of the winning bid. Includes only auctions with at least one bidder.

Table 1.11: 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.

Table 1.12: Mean outcomes considering optimal reserve prices, $v_0 = R_{obs}$

	None		3 Months		6 Months		10 Months	
Reserve Price Used	Data	Optimal	Data	Optimal	Data	Optimal	Data	Optimal
Reserve Price ($1000s)	42.3	67.0	42.3	66.8	42.3	68.3	42.3	67.4
Gov. Revenue ($1000s)	64.6	66.0	64.1	65.5	62.7	64.2	59.8	61.3
Firm Surplus ($1000s)	15.3	12.6	15.2	12.4	15.6	12.5	14.2	11.1
Number of bidders	3.6	2.2	3.5	2.2	3.4	2.0	3.1	1.8

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

Table 1.13: 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 1.23.

CHAPTER 2.

Weather, Salience of Climate Change, and Congressional Voting

2.1 Abstract

Climate change is a complex long-run phenomenon. The speed and severity with which it is occurring is difficult to observe, complicating the formation of beliefs for individuals. We use Google search intensity data as a proxy for the salience of climate change and examine how search patterns vary with unusual local weather. We find that searches for "climate change" and "global warming" increase with extreme temperatures and unusual lack of snow. Furthermore, we demonstrate that effects of abnormal weather extend beyond search behavior to observable action on environmental issues. We examine the voting records of members of the U.S. Congress from 2004 to 2011 and find that members are more likely to take a pro-environment stance on votes when their home state experiences unusual weather.

2.2 Introduction

Anthropogenic climate change is one of the most difficult policy problems that humanity faces today. The costs and benefits of mitigating carbon emissions are highly uncertain. The relevant pollutants are globally mixing, which creates an enormous collective action problem. Finally, the process of climate change unfolds over several decades. Because the impacts of climate change manifest themselves as gradual changes in the distribution of weather outcomes, it can be difficult for individuals to observe whether climate change is occurring. In addition, climate change is a one-time event, and individuals cannot possibly draw on prior experience to guide their perceptions. However, public support and understanding are vital to the successful creation and implementation of climate change mitigation and adaptation policies.

Given these complications, people may seek a proxy by which to update their opinion. Unusual weather could be used (rightly or wrongly) as an observable, short-term analog to climate change. Indeed, Hansen, Sato, and Ruedy (2012) describe the effect of climate change as changing the weights on a pair of dice that determine short-run realizations of weather. In this paper, we estimate the effect of unusual weather conditions on salience of climate change. We proxy for salience using a search intensity index created by Google for the terms "climate change" and "global warming". Controlling for a wide variety of fixed effects to account for spurious geographic and seasonal relationships and broad temporal trends, our results are remarkably robust and suggest that short-run weather phenomena do in fact affect the extent to which people think about climate change.

Furthermore, we demonstrate that the effects of weather extend beyond search behavior to the voting records of U.S. Congressional members. Examining *within-member* variation in support for 207 environmental votes tracked by the League of Conservation Voters (LCV) between 2004 and 2011, we find evidence that voting on environmental issues is correlated with recent unusual weather in a representative's home state. Reassuringly, the correlation between weather and voting does not extend to votes unrelated to the environment. Although the effect is modest in size, our results suggest that that search intensity may provide a useful proxy for voter and legislator concerns and demonstrates an important link between unusual weather and political action on environmental policy.

Our work relates to several other papers. A series of papers estimate the extent to which individuals respond to short-run weather in forming their beliefs about climate change. Deryugina (2013) uses an annual Gallup poll to determine whether individuals respond to weather fluctuations by Bayesian updating their expectations about climate change. She finds that while short-term weather fluctuations do not affect individuals' beliefs, longer spells of unusually warm weather do have an impact. She also examines heterogeneity by political affiliation and finds that the effect is confined largely to conservative respondents. Hamilton and Stampone (2013) analyze a series of polls of New Hampshire residents. Interestingly, they find that political independents are the only subgroup that respond to recent weather cues in forming their opinions regarding climate change. Owen, Conover, Videras, and Wu (2012) find that respondents to a pair of surveys in August 2009 and October 2007 are more likely to support environmentally-protective policy if their state experienced a heat wave or drought during the most recent summer. They also find that people who regularly access more sources of news information are less responsive to weather cues. Egan and Mullin (2012) also find evidence of a response.

A separate literature demonstrates the value of internet search data in modeling economic behavior. Choi and Varian (2009) demonstrate that Google Insights data can be used to predict demand for automobiles, retail sales, home sales, and travel behavior. After several papers demonstrated the efficacy of using Google searches to predict flu outbreaks, Google itself established the Google Flu Trends tool.[1] Most relevant to our analysis is Kahn and Kotchen (2011). They find that when a state's unemployment rate increases, Google search activity for "global warming" decreases and search activity for "unemployment" increases. That is, concerns about economic conditions "crowd out" attention to the issue of climate change. These results focus on unrelated trends that compete with climate change for attention. In contrast, we examine a factor (weather) that *directly attracts* attention to climate change because it is a series of realizations of the broader climate process.

Our paper makes two contributions. Previous studies of climate beliefs and weather use survey waves that are either infrequent or limited to a specific geographic location. In contrast, search intensity is reported weekly for each state – higher frequency reporting provides us much more identifying variation with which to estimate the relationship between weather and search intensity flexibly and to better control for unobserved heterogeneity that might be correlated with both

[1] http://www.google.org/flutrends

weather and search activity.

Our empirical results suggest this flexibility is important along several dimensions. First, the variation in the data allows us to simultaneously estimate the effects of temperature, precipitation and snowfall. For instance, given a response to an unusually warm winter, we can estimate the relative contributions of warmer-than-average temperatures separately from the effect of a lack of snow. These various channels may have completely different implications. For example, if the response is entirely due to lack of snow and not higher temperatures, this could limit the relevant geographic (and climatic) range to that in which snowfall now regularly occurs. Further, snowfalls are easily-observed individual events; in contrast, an increased frequency of extreme temperatures might not be as discernible. Second, we find evidence that the search intensity responds asymmetrically to unusually high and low temperatures and snowfall. For instance, in the winter, unusually cold and warm weather are both correlated with increased search; this would be obscured by a fully-linear specification. Third, the effects of weather on search intensity vary by season.

Our second contribution to is to provide an important link between weather and search behavior to observable actions related to the environment – specifically, the voting behavior of members of the U.S. Congress on environmental bills. Previous work has focused on individual attitudes as the explanatory variable of interest, but has not established a link between weather and tangible changes in behavior. Our work helps to fill an important gap. Controlling for member fixed effects, we find that U.S. congressional members are more likely to cast a pro-environment vote when their home state experiences unusual weather and search intensity in their home state is high. Reassuringly, the effects are specific to environmental legislation, and in particular, environmental regulation most closely related to climate change – we do not find similar effects of weather or search intensity on non-environmental legislation, nor do we find strong effects for environmental legislation unrelated to climate policy or industrial emissions. Although the effects we estimate are modest in size (as would be expected) and may not affect the ultimate outcome of the vote, our results suggest that extreme local weather (or the issue salience it generates) is a factor legislators may consider when voting on environmental issues. Furthermore, our results suggest that internet search intensity may provide a useful proxy for the salience of issues to the broader public.

Our paper proceeds as follows. Section 2.3 describes the data and econometric approach. Section 2.4 presents the empirical results related to weather and search intensity. Finally, Section 2.5 examines the relationship between extreme weather, individual search behavior and voting of members of Congress on environmental issues. In the online appendix, we demonstrate the robustness of our results to a number of different specifications.

2.3 Methodology

2.3.1 Data

Search intensity data Our proxy for climate change salience uses the Google Insights (now part of Google Trends) search index. This tool is outlined in Stephens-Davidowitz (2013). Essentially, Google Insights tracks the relative frequency with which a given search term is submitted. In most of our specifications, we use the index for searches of ("global warming"+"climate change") at the state-week level. The index is constructed to facilitate accurate comparisons across periods and locations; that is, a given search term is scaled by the overall level of search activity in each state. The advantage of this approach is that a populous state, such as California, will not have a mechanically higher search index than a less populous state, such as Iowa. Thus, our measure of the search term corresponds to search intensity, conditional on overall search activity. Google censors search terms that do not surpass a certain threshold in terms of *absolute* search volume. This affects approximately 20% of our sample from 2004-2011, but is most relevant in 2004-2006 for sparsely populated states in the Great Plains and Rocky Mountain regions.[2]

We can use several data sources to get a sense of nationwide search magnitudes during our study period. Google Adwords, a service for potential advertisers, reports that U.S. users googled "climate change" or "global warming" approximately 185,000 time per month in 2013. To estimate

[2] In the online appendix, we re-run our regressions using only state-years for which complete data is available and find that our results do not change substantively.

total search volumes for our study period, we adjust total searches in 2013 for changes in search intensity (tracked by Google Trends) and changes in nationwide Google search volumes from comScore, a market research company that tracks media and internet trends. Although total Google searches rose from 134 billion searches in 2011 to 154 billion searches in 2013, search intensity for "climate change" and "global warming" fell by approximately 20 percent during the period. The two changes roughly offset each other - our best guess is that relevant searches averaged roughly two hundred thousand per month in 2011. Using a similar methodology, we estimate that at the peak, searches for the two terms averaged approximately half-a-million searches in January and February of 2007.

Weather data Our weather data come from the National Climatic Data Center (NCDC). The NCDC collects daily weather station data for over 10,000 U.S. weather stations. The typical station records minimum and maximum daily temperature, precipitation and in some cases, snowfall, snow depth and other meteorologic variables. For purposes of this paper, we limit our analysis to 6,624 stations with data on minimum and maximum temperatures from 2004 to 2011. The stations are located throughout the 50 states – Rhode Island has the fewest stations (8) and California, the most (370). For each daily station record, we calculate the deviation of maximum daily temperature, precipitation, snowfall and snow depth from a 10-year baseline from 1994 to 2003 and matched by day of the year. To match the search intensity data, we aggregate up to the state-week level.

Summary Statistics To illustrate one dimension of our weather variation, we plot monthly average temperature deviations from the 1994-2003 baseline in Figure 2.1 going back to 1974. The solid line is the lagged 12-month moving average deviation. The dotted line is a linear trend and illustrates that temperatures have been increasing on average since 1974. This trend is less pronounced if we focus solely on the last two decades. Although average temperatures have risen since 1974, the warmest 12-month period in U.S. history prior to 2012 stretched from late-1999 to late-2000, during our 10-year baseline period.[3]

We present summary statistics of the weather and search variables for our regression sample in Table 2.1. The weather variables are presented as deviations from the 10-year baseline covering 1994-2003, matched by state-calendar week. Relative to baseline, the period from 2004 through 2011 was similar in terms of temperature and slightly snowier, on average. As one would expect, there is substantial week-to-week variation around the baseline.

The relationship between our sample and the baseline differs by season. Relative to the 10-year baseline, winter has been slightly colder than normal, while spring, summer, and fall have been slightly warmer. The standard deviation of the temperature variable is of the same order of magnitude for all seasons, and suggests that there is considerable variation around the mean. As one would expect, snowfall and snow depth are most variable in the winter, somewhat less variable in the spring and fall, and quite tightly distributed in the summer.

2.3.2 Empirical Approach

In essence, we want to identify the effect of unusual short-run weather on the relevance of climate change in the eye of the general public, using the Google search intensity index outlined above as a proxy for salience. We take a largely agnostic stance on the mechanisms underlying a possible relationship. Weather could affect search intensity through channels such as personal experience, exposure to news coverage of extreme weather, or interactions with friends and family.

We simultaneously estimate effects for the maximum temperature, precipitation, snowfall, and snow depth. Table 2.2 presents basic correlations among the explanatory variables. As one would expect, deviations in temperature, snowfall, and snow depth are correlated with one another. However, the frequency of our panel provides sufficient independent variation to estimate the coefficients on each precisely.

[3]Source: http://www1.ncdc.noaa.gov/pub/data/cmb/images/us/2012/jul/warmest_12months.png

The base specification for state s, week w, month m, year y can be expressed as:

$$INDEX_{s,wmy} = \sum_j \beta^j DEV^j_{s,wmy} + \alpha_{my} + \gamma_{sm} + \varepsilon_{s,wmy} \tag{2.1}$$

where j indexes the four weather variables, $DEV^j_{s,wmy}$ is the deviation from the historical mean for measure j, β^j is the effect of measure j on the climate change search intensity index, and α_{ym} and γ_{sm} are fixed effects. In our main specification, we relax the linearity of the relationship of the index on the deviation variables by allowing for asymmetric effects depending on the sign of the deviation:

$$INDEX_{s,wmy} = \sum_j \beta^{nj} NEGDEV^j_{s,wmy} + \sum_j \beta^{pj} POSDEV^j_{s,wmy} + \alpha_{my} + \gamma_{sm} + \varepsilon_{s,wmy} \tag{2.2}$$

where $NEGDEV^j = I(DEV^j < 0) * |DEV^j|$ and $POSDEV^j = I(DEV^j > 0) * |DEV^j|$. Thus, the coefficients $\{\beta^{nj}, \beta^{pj}\}$ are the effect of the *magnitude* of negative / positive deviation from the 10-year weather baseline on search intensity.

We graphically illustrate the basic idea behind our empirical strategy. Figure 2.2 plots kernel-smoothed time trends of the residuals of search index and average snowfall for Colorado from October 2006 through April 2007 after conditioning on year-month and state-month of year fixed effects. Through early December, snowfall tracks close to the 10-year baseline. In late December, relative search activity is halved during a series of weeks with unusually high snowfall. However, as snowfall becomes more scarce in late January and February, search activity increases again.

A first potential concern with our analysis that Google searchers may not be representative of the general public. Past analyses such as Choi and Varian (2009) and Kahn and Kotchen (2011) suggest that Google search is sufficiently in the mainstream to be useful for this sort of analysis. In addition, we are not making claims as to whether local weather will help support for climate change reach some crucial electoral threshold. Rather, we examine whether very short-run weather events have the capability to affect the salience and prominence of climate change. Compared with 2010 Census data, the distribution of Google searchers skews away from those over 65 years of age, and toward those 18-25. The shares in the 25-44 and 45-65 age groups are roughly the same as in the population.[4]

In addition, one might be concerned that there may be underlying seasonal or geographic correlations that are purely coincidental. For instance, as displayed in Table 2.1, recent summers have been hot compared with baseline means while recent winters have not. During our sample period, the Conference of the Parties to the United Nations Framework Convention on Climate Change convened during November and December in each year. If this highly climate-relevant event results in a spike in news coverage and search activity, we would incorrectly estimate a negative relationship between maximum temperature and climate search intensity. Similarly, if states with more urban areas have had systematically different weather deviations than more rural states, we might misattribute a correlation between weather differences and differences in political ideology as reflected in interest in climate change.

To address these concerns, we employ a variety of fixed effects to control for such possible sources of bias. In our preferred specification, we include year-month fixed effects and state-month of year fixed effects. The variation identifying our primary estimates controls for broad national trends during a given month, and monthly seasonality at the state level. For a given January week in Iowa, we consider the covariance in how unusual search and weather are among all January weeks in Iowa, controlling for nationwide means in that specific month. The year-month effects capture changes in nationwide attitudes toward climate change, average internet penetration, and changes in the makeup of internet users over time. The state-month of year effects control for state-specific seasonality in weather deviations and climate change search intensity.

Finally, search activity by climate skeptics could affect the implications of our results. In Figure 4, we compare the national time-series of our primary search with one that nets out several potential skeptical searches. As is clear from the figure, these explicitly skeptical searches comprise a small fraction of the total searches. The window indicated in the figure does display one week of particularly high skeptical search activity: it corresponds to the "Climategate" incident. Our

[4]Google search demographics from comScore, via http://blog.pmdigital.com/2010/08/who-uses-google-yahoo-and-bing. Census demographics from http://www.census.gov/prod/cen2010/briefs/c2010br-03.pdf

results are robust to omitting this period. Of course, we cannot hope to identify all such searches; there exists a strong current of skepticism among parts of the U.S. population. Our results do encompass the causal effect of weather shocks on the search habits of such skeptics. However, our interpretation of changes in search intensity as a proxy for issue salience does not change.

2.4 Weather and Search Intensity Results

The results from the base specification are presented in Panel A of Table 2.3. The first column is a simple specification in which climate-related search intensity is modeled as a linear function of deviations from historical weather patterns. Perhaps surprisingly, in the aggregate, higher temperatures (relative to the baseline) are associated with lower search intensity. The coefficient on snowfall is as expected, in that unusually low snowfall is related to more climate change searches.

We relax the initial specification in two ways. First, we run our analysis separately for each season of the year in columns (2) through (5) of Panel A. This allows the effects of unusual weather on search intensity to have different magnitudes and signs across seasons. For example, unusually warm weather in the winter might be far more noticeable in the winter than in the spring. We find that the effects vary considerably by season. While lower temperatures are still negatively related to search in the winter, the opposite is true in the summer. The effect of unusually low snow depth is now statistically significant in the winter and fall, but not in the spring (or summer).[5]

Second, we allow the effect of weather to vary asymmetrically with respect to positive and negative deviations from the 10-year baseline. Although results from Panel A of Table 2.3 provide evidence that short-run weather shocks are correlated with search intensity, if search intensity responds differently to positive and negative deviations from the baseline, these specifications may mask the true effect. The bias would be particularly pronounced if search intensity is a function of the absolute deviation of weather from the long run average. To this end, Panel B of Table 2.3

[5]For completeness, we also present coefficients for each month of the year in the online appendix. Providing further flexibility in estimate the coefficients by month does not provide any additional insights beyond the estimation by season.

presents results that allow positive and negative weather deviations to have asymmetric linear effects on search intensity. To be clear, our specification regresses search intensity on the absolute value of positive and negative deviations. If the coefficients for the positive and the negative deviation in snowfall are both positive, then the relationship between snowfall and search intensity is "V"-shaped.

Both positive and negative deviations from the baseline average temperature are positively associated with search intensity. The negative temperature deviation coefficient from column (1) of Panel A is driven by the fact that the search-inducing effect of a negative deviation dominates the effect of a positive deviation. Search intensity seems to respond weakly to unusually dry weather. The coefficients on snowfall and snow depth especially illustrate the importance of allowing for asymmetric effects. The negative snowfall and snow depth coefficients from column (1) of Panel A would suggest that there is more search activity in normal weeks than in especially snowy weeks. However, when we allow asymmetric effects, we find that weeks of abundant snowfall and snow depth do not seem to differ from a normal week in terms of search intensity. Instead, the flexible specification demonstrates that the effect in Panel A is driven by weeks with a notable lack of snow. The coefficients on negative deviations in snowfall and snow depth are roughly four times larger than their counterparts in Panel A, and the coefficient on snow depth is now statistically significant.

We again run separate regressions for each season and present the results in columns (2) through (5) of Panel B. We interpret the magnitude of the coefficients in the following manner. The search index is simply the number of searches involving climate change or global warming as a share of total search activity, scaled by some unknown coefficient. We assume that climate-related searches are a small proportion of total search activity. Thus, a 10% increase in the search index corresponds to a 10% increase in climate-related searches. We will consider the effect of weather shocks on the mean week in percentage terms. For instance, in the winter, the mean search index is 43.02. An 4.302-unit increase in the search index during the winter would correspond to a 10% increase in climate-related search over the mean week.

As before, we find substantial variation in the effect of abnormal weather across seasons. In the winter, search intensity responds positively to both unusually cold and warm weather. In particular, a winter week that is 4°C colder than normal (1 standard deviation of our temperature variable) would result in an increase in the search index of 6.54, or a 15.2% increase in climate-related search activity relative to the mean week. Similarly, a week that is 4°C warmer than normal would result in an increase in the search index of 2.20, or about 5.1%. Much of the effect of warm winter weather operates through a lack of snowfall. Indeed, a winter week that has less snowfall than average by only 10mm (roughly 1 standard deviation) is also associated with an increase of roughly 2.56 (6.0%) in the search index; a week in which the average snow depth is lower than usual by 1 standard deviation (roughly 70mm each day) is associated with an increase of 5.32 (12.4%) in the search intensity. These magnitudes suggest that weather shocks are actually responsible for fairly large movements in climate-related search activity relative to the mean week.

Responses during other seasons demonstrate different patterns. In the spring, weather does not actually seem to have much of an impact: none of the coefficients are statistically significant. This confirms a main result of Deryugina (2013), who finds that beliefs elicited in a March survey are not affected by very short-run weather deviations. In the summer, search responds strongly to extremely hot temperatures, but not to cool temperatures. Negative deviations in summer precipitation are associated with less search. Finally, in the fall search increases with unusually low snowfall and snow depth. This is consistent with search responding to steadily warm fall weather that delays the first snowfall or a heat wave that results in unexpectedly extreme temperatures.

In the online appendix, we provide a number of robustness checks. We run separate regressions for each month of the year. We also repeat our analysis including several different of combinations of fixed effects. Finally, we perform our analysis at the city level for the 25 largest cities in the U.S. Our results prove to be quite robust to all of these alternative specifications.

It is important to note that these results are consistent with several alternative models of economic behavior and belief updating. Despite the high temporal frequency, the aggregate nature of the search data does not allow us to make strong conclusions about the particular method by which people adjust their beliefs. As an example, our results may reflect rational re-evaluation of beliefs of climate change by individuals who were previously skeptical. Unusual weather may cause them to update their beliefs and search online for more information about climate change.[6] Equally plausible, though, are alternative explanations for relationship between unusual weather and search activity. Evidence in favor of some of these explanations already exists in the literature. Kahn and Kotchen (2010), for example, propose that concerns about environmental concerns fight for individuals' limited attention - they find evidence that concerns about climate change may be crowded out by economic concerns. More generally, we think it unlikely that a single explanation would fully explain the relationship between abnormal weather and search behavior. Thus, we refrain from advancing a particular story or explanation for the results, although we believe that this is an interesting avenue for future research.

2.5 Weather, Search Intensity and Voting Behavior

We now pivot from examining the relationship between abnormal weather and internet search activity to examining observable action on environmental issues, specifically the voting behavior of members of the U.S. Congress. In this section, we extend our approach from the previous section demonstrate that atypical weather is correlated with the voting behavior of members of the U.S. Congress on environmental issues.

Our analysis directly relates to two literatures. A long literature in political science suggests "issue salience" plays an important role in voter engagement (Brians and Wattenberg, 1996), attitudes towards elected officials (George Edwards and Welch, 1995) and policymaking (Burstein, 2003). Specifically, issues which voters perceive as particularly relevant are correlated with election turnout, approval ratings and political action on issues. Second, our results relate to the literature on classic political economy originating with Stigler (1971) and Peltzman (1976) that postu-

[6] If we believe, though, that this is the only driver of search activity and that there is an initial stock of "climate skeptics," we might expect that stock to deplete over time and more unusual weather occurs and consequently, the effect of unusual weather may diminish. In our data, we do not find strong statistical evidence that the effects of unusual weather on search behavior diminish over time, although we acknowledge that this does not provide definitive evidence against this explanation.

late that voting behavior is driven both by individual ideology and the need to represent constituent interests.

Our primary source of voting data comes from the League of Conservation Voter ("LCV") scorecards. For each member of Congress and each vote on bills, resolutions, motions and amendments related to the environment, the LCV records a member's vote and identifies whether the vote represents a pro- or anti-environment position. LCV scorecards (and voting scorecards more generally) have been used extensively in the literature (see Kahn, 2002; Levitt, 1996; Kalt and Zupan, 1984) to identify members of Congress who tend to take pro- or anti-environmental stances. For our analysis, we use constructed a panel of all the members of the U.S. House of Representative or the U.S. Senate. For each congressperson, we track his or her vote on 207 environmental votes scored by the LCV between 2004 and 2011.[7] Democrats tend to receive high LCV ratings and Republicans tend to receive low LCV ratings – the mean ratings for Democrats and Republicans are 89.7 and 14.1 on a scale of 0 (uniform voting against environmental positions) to 100 (uniform voting in favor of environmental positions), but LCV scores vary within political party substantially. Of congressional members in office for more than a single year in the 2004-2011 period, Dan Boren (House, OK) was the lowest rated Democrat at 32.7 and Christopher Shays (House, CT) was the highest rated Republican at 88.1.

We consider a linear probability model and regress pro-environment voting as a function of weather in a member's home state.[8] All specifications include congressional member fixed effects. Consequently, identification comes from within-member variation – we test whether member i's vote on environmental vote v is correlated with anomalous weather conditions in their home state s at a similar point in time t. We also include varying sets of time fixed effects to flexibly control for state-invariant shifts in the propensity to vote in favor of environmental regulation.

We use two approaches to test for the relationship between anomalous weather and congressional voting. First, we directly regress voting on the weather variables from the previous section.

[7] We exclude eight votes that are tracked by the LCV, but not directly related to environmental issues, such as the reauthorization of the Childrens' Health Insurance Program or the nomination of federal judges. Our results are robust to the inclusion of these six votes.

[8] We obtain qualitatively similar results using a probit model.

As before, we allow for an asymmetric relationship between the dependent variable and positive and negative weather deviations. Formally, our we consider the specification

$$Pro - Env.\, Vote_{i,v} = \alpha_i + \sum_j \beta^{nj} NEGDEV^j_{s,t} + \sum_j \beta^{pj} POSDEV^j_{s,t} + \varepsilon_{i,v} \tag{2.3}$$

where j denotes each weather variable and $NEGDEV^j_{s,t}$ and $POSDEV^j_{s,t}$ represent positive and negative deviations from the 10-year baseline.

Table 2.4 presents the main results relating voting on environmental issues to weather and search intensity. Panel A presents the results of the linear probability model of pro-environment voting on weather, member fixed effects and successive sets of time fixed effects. Unusually low temperatures in a member's home state are correlated with a greater likelihood of voting against environmental legislation or motions. Unusually low snowfall in a member's home state is correlated with an increased likelihood of voting in favor. The magnitudes are modest but significant and persist with the inclusion of year-month fixed effects that subsume the effect of national weather or news spuriously correlated with weather that occurs in the month of the environmental vote. Snowfall one standard deviation below the mean during winter months in associated with an 1.5 percentage point increase in the likelihood of voting in favor of environmental legislation. The eight weather variables are highly significant, collectively, in the specifications in columns (1) and (2). In the specification in column (3), the p-value on the F-test of the weather variables is 0.147, slightly above conventional levels for significance.

As a second approach, we construct an "index" of the abnormality of recent weather in a state. For our index, we project search intensity onto four lags of the local climate deviations for temperature, precipitation, snowfall and snow depth.[9]

In essence, the projection consolidates unusual rainfall, temperatures and snowfall into a single summary statistic. This procedure creates a more parsimonious measure of abnormal weather; we use this measure to clarify the relationship with voting behavior and allow heterogeneity to enter

[9]The F-statistic for the joint test of the coefficients on the weather variables in equation (2.4) is 26.53.

in a concise way. As a result, we interpret the coefficient on the projected weather variables as the reduced-form effect of any combination of collectively abnormal weather variables that would induce a one-point change in search intensity.

It is important to note that the interpretation of the coefficient in this context differs from that of an instrumental variable regression. A true IV regression would estimate the causal effect of one particular channel (in our case, search intensity) on voting. Rather, our approach measures the collective effect of weather through a number of different channels. The projection allows us to treat unusual realizations of temperature, precipitation and snowfall comparably.

Our approach is similar to a number of recent papers that project one or more covariates onto a single variable to analyze a reduced-form effect. Madestam, Shoag, Veuger, and Yanagizawa-Drott (2013) examine political protests and representative voting. They project the size of Tea Party tax day protests on rainfall but cautiously interpret the coefficient on protest size, noting that rainfall may affect both the size of the protest and "quality" of the protest. Chodorow-Reich (2014) compares post-financial crisis employment at firms as a function of the exposure of a firms' banking partners to the financial crisis. As a proxy for a bank's exposure, the paper projects the change in annualized loans between 2005 and 2009 onto a set of pre-crisis covariates plausibly related to a bank's financial strength. Again, the author notes that the "second stage" does not identify a particular causal pathway. Rather, the projected change in annualized loans is interpreted as a summary statistic for a number of factors related to the financial strength of the bank.

A second advantage of the projection of search intensity onto the weather variables is to mitigate concerns of reverse causality. If internet searches related to climate change are partially driven by actions taken by Congress or by the voting of particular members, a positive correlation between search intensity and Congressional voting may simply reflect constituents' interest in the position taken by their representative. In contrast, the projection only relies on variation in search intensity correlated with lagged weather variables.

Formally, we estimate the following:

$$SI_{s,t} = \gamma_s + \sum_{k=1}^{4}\sum_{j} \lambda^{njk} NEGDEV^{j}_{s,t-k} + \sum_{k=1}^{4}\sum_{j} \lambda^{pjk} POSDEV^{j}_{s,t-k} + \nu_{s,t} \quad (2.4)$$

$$Pro - Env.\, Vote_{i,v} = \alpha_i + \beta\widehat{SI}_{s,t} + \varepsilon_{i,v} \quad (2.5)$$

In Panel B, we present the coefficients estimated by (2.5), using member fixed effects and a similar series of time fixed effects to those in Panel A. [10] Again, we find that the weather-correlated component of search intensity is correlated with voting in favor of environmental legislation – a one standard deviation increase in scaled search intensity (.28) is associated with a 8.4 percentage point increase in the likelihood of voting in favor of environmental legislation. As with the weather variables, the magnitude of the coefficient declines with the inclusion of finer time fixed effects. Month by year fixed effects subsume the effect of national weather events, such as the 2012 U.S. summer heatwave. Identifying the coefficient off of this within-month variation only, a one standard deviation increase in search intensity is associated with a 2.3 percentage point increase in the likelihood of voting in favor of environmental legislation. As a point of reference, Hussain and Laband (2005) examine 33 LCV votes whose costs are confined to a small set of states. Senators who represent one of those states are 15% less likely to cast a pro-environment vote. Given the extreme political circumstances involved in those votes, our effect (one-eighth as large for a 1 standard deviation increase in search, one-quarter as large for a 2 standard deviation increase in search) appears non-negligible.

While we find evidence of a strong positive correlation between weather-driven search intensity and likelihood of pro-environmental voting, as we note able we are not interpreting the regression results above as an IV estimate of the causal effect of search intensity on voting behavior. Rather, there are several possible pathways that could be driving this correlation. First, the weather might directly affect the Congressperson's ideology/beliefs in the same way that it affects his/her constituents' beliefs. Second, the weather might directly affect voting through Congressional beliefs,

[10] For ease of presentation, we rescale the Google Search Intensity a scores by a factor of 100 - values of 0 and 100 in the original index correspond to values of 0 and 1 in the rescaled index.

thus leading to increased search through increased constituent awareness or local media coverage. Third, the increased search activity could indicate increased constituent pressure on the Congressperson to vote in an environmentally-favorable way. Our data do not allow us to distinguish among these three effects, but they are all relevant in that they reflect an impact (whether direct or indirect) of weather shocks on legislative behavior.

One concern with these results is that the timing of votes may be endogenous. All of our previous results condition on an environmental vote being held – for endogeneity to spuriously drive our results, Congress would have to schedule favored environmental votes in weeks following extreme weather and unfavored environmental votes in other weeks. Because we only include year-month and member fixed effects, our identification strategy would be vulnerable to such a phenomenon. Although we cannot observe whether a particular environmental vote is preferred for other reasons, we can examine whether the timing of environmental votes seems to follow extreme weather overall. We regress contemporaneous and one-week lagged weather deviations on an indicator for whether an LCV vote occurred, controlling for year-month and state fixed effects. The idea is to compare weeks within a calendar month, and see if LCV votes happen following weeks with more extreme weather. We do not find evidence that this is the case. Given this finding, reverse causality would only be problematic if those environmental votes that are inherently more favored overall also tend to be scheduled after especially extreme weeks of weather.

A second concern is other factors that might drive a spurious relationship between the timing of votes and unusual weather. It is still possible that local weather is spuriously correlated with changing political preferences at the state-level and hence, within-member voting on environmental legislation. As a placebo test, we examine voting data from the American Conservative Union (ACU). Similar to the LCV, the ACU tracks "a wide range of issues before Congress to determine which issues and votes serve as a dividing line to help separate those members of the U.S. House and Senate who protect liberty as conservatives and those who are truly liberal."[11] For the placebo test, we use the 350 non-environmental votes tracked by the ACU from 2004 to 2011.[12]

[11] http://conservative.org/legislative-ratings/

[12] The ACU tracks votes related to immigration, the minimum wage, family planning, religious freedom and other

If general political preferences are shifting at the same time as unusual weather, we should expect that the weather-correlated variation in search intensity would be correlated with voting on non-environmental votes tracked by the ACU. Table 2.5 presents the results of an identical specification to Panel B of Table 2.4 using Congressional member voting on non-environmental issues tracked by the ACU rather than environmental votes tracked by the LCV. Columns (1) through (3) use all of the non-environmental votes tracked by the ACU; columns (4) through (6) use only the non-environmental votes tracked by the ACU that occur in the *same week* as the environmental votes tracked by the LCV. We do not find the weather-correlated component of search intensity to be strongly correlated with taking liberal or conservative positions on votes unrelated to the environment, even when restricting the set of votes to those occurring in the same week as the LCV votes. Thus, we do not find strong evidence that suggests that our results are driven by changes in *general* voter preferences that are spuriously correlated with unusual weather.

Finally, we consider two possible sources of heterogeneity in the response of voting to unusual weather, drawing on the political economy literature originating with Stigler (1971) and Peltzman (1976) that postulate that voting behavior is driven both by individual ideology and the need to represent constituent interests. Similar to more recent empirical articles on voting behavior, such as Kalt and Zupan (1984) and Levitt (1996), we posit that the weight a representative places on individual ideology and constituent interests vary with respect to the position of the representative and the nature of the issue on which the vote is taken. For example, incumbents facing re-election may weight constituent interests highly as might a representative facing a vote that demonstrates dedication to his or her district.

In our context, we examine two sources of heterogeneity. First, we allow the response to extreme weather to vary by congressional member characteristics. If left-leaning constituents care more about environmental issues, we posit that representatives from these districts may face greater pressure in response to abnormal weather. Moreover, we might expect that representatives facing

issues unrelated to the environment. In addition, the ACU tracks 44 votes related to the environment issues that are also tracked by the LCV (e.g. HR 2643: Allowing the Dept. of the Interior to issue new leases for offshore natural gas development); we omit these 44 votes from the placebo test.

2-year re-election cycles in small, geographically contained districts may face greater re-election pressure from constituents than senators and consequently, may be more responsive to short-lived weather anomalies.

The specifications in Table 2.6 test whether the strength of the correlation between search intensity and voting behavior differs by the characteristics of the Congressional member. We interact the weather-correlated component of search intensity with whether the member of Congress is a Democrat, a member of the Senate, and with the member's LCV score over the 2004-2011 period. In one of the three specifications, we find that the correlation between anomalous weather in a member's home state and voting on environmental legislation is significantly stronger in the House than the Senate. This is consistent with the hypothesis that six-year terms in the Senate that may make Senators less responsive to short-lived changes in constituent interests. We also find strong evidence that the response to unusual weather also differs by political affiliation. The correlation between voting and home-state search intensity is significantly stronger for Democrats than Republicans. As a refinement, we allow for the response to unusual weather to differ for each ten-percentage point bins of LCV ratings. These coefficients are plotted with 95% confidence intervals in Figure 2.5.[13] As before, positive values indicate that a member is more likely to take a pro-environment stance when home-state search intensity is high and less likely to take a pro-environment stance when home-state search intensity is low. Although we find little evidence of correlation between voting and home-state search intensity for members with LCV ratings below 50 percent, we find a positive and strongly significant relationship for members that take a pro-environment stance slightly more than half the time. Unsurprisingly, the correlation diminishes for members with very high LCV ratings – these members almost always vote in favor of environment legislation.

A second source of heterogeneity examines the characteristics of the votes themselves. The LCV tracks a wide variety of votes related to the environment, only a subset of which relate to climate change or air pollution more generally. If a congressional member's vote acts as a verifiable

[13]The specification generating the coefficients estimates in the figure is a refinement of specification (3), and includes member fixed effects.

signal to constituents, we might expect the effect of unusual weather to be greater for policies that directly relate to climate change or pollution. To test this hypothesis, we hand-classify the 207 votes into three categories: (1) 18 votes directly related to climate change or carbon policy, (2) 84 votes related to industrial pollution or regulation, and (3) 105 votes related to the environment, but unrelated to industrial policy or carbon emissions, such as wetland protection.[14] A second hypothesis relates to the votes that are particularly close to passage. As opposed to votes on issues with more clear bipartisan support or resistance, party leadership may "coordinate" caucus voting behavior on issues that are very close to passage or defeat. Thus, we would expect that Congressional members may have increased latitude when voting on bills or motions that are expected to handily pass or fail. Although no clear guidelines exist for what constitutes a "close" vote, we define votes that passed or failed by less than five percent of the vote to be "close." Figure 2.4 plots the histogram of pro-environment vote share for all 207 of the issues tracked by the LCV between 2004 and 2011. Graphically, the votes falling between the dotted lines represent the issues close to passage. Using this criterion, 73 of the 207 the votes are classified as close votes. Votes that are close are roughly equally distributed across all three categories of environmental votes.

Table 2.7 presents the results allowing for the effects of anomalous weather to vary based on vote characteristics. As before, the three columns correspond to specifications without fixed effects, with year fixed effects, and with month-year fixed effects. Focusing on our preferred specification in column (3), we find that anomalous weather is uncorrelated with voting for the least-relevant group of environmental issues. In contrast, we find a significant, positive correlation between anomalous weather and voting on bills and motions that are more closely related to carbon emissions or industrial pollution. We estimate that a one standard deviation increase in the search intensity is correlated with a 8.4 percentage point increase in the likelihood of a representative taking a pro-environment stance on a vote related to industrial pollution and an 11 percentage point increase in the likelihood related to carbon emissions policy. Although we cannot distinguish whether anomalous weather affects voting through constituent preferences or a representative's

[14] A list of all the votes and classifications are available from the authors by request.

own beliefs, we find the strongest correlation between voting and extreme weather exactly where political economy would suggest. In addition, we find suggestive evidence of diminished influence of unusual weather on bills and motions very close to passage.

It is important to qualify the results above in two respects. First, the correlation between voting and search intensity reflects the voting of individual members, conditional on the actual legislation brought to a vote. While we find that members of Congress (and in particular, Democrats) are more likely to vote in favor of environmental regulation when home-state relative search intensity for global warming or climate change tends to be high, we cannot assess whether this implies discrete changes in the passage of legislation or the changes in the content of legislation brought to a vote. Most of the votes tracked by the LCV were passed or defeated with substantial support; in these cases, the vote of a single member is unlikely to be marginal ex-ante and members of Congress may have more latitude to take a position contrary to the position of their party. Only 15 percent of the votes tracked by the LCV were passed (or defeated) by less than a five percentage point margin. Members (and caucuses) may behave differently for votes close to passage or defeat. We nonetheless feel that the observed relationship to marginal voting behavior is meaningful. The relationship illustrates that abnormal weather or high search intensity is related to important, observable behavior on environmental issues. Although the political economy of the legislative process makes it unlikely that the marginal effect of an individual Congressional member would translate into discrete changes in policy, our results suggest that search activity may be a useful proxy for constituent concern and the salience of particular policy issues.

Second, while we identify an effect of abnormal weather on pro-environment voting, it is beyond the reach of our existing data to map a clear causal chain from weather to legislative action. As we note above, we are not arguing that search activity itself is solely responsible for the changes in voting behavior we identify, but rather that search activity (once instrumented) represents a possible proxy for the abnormality of weather. As the previous literature (Kahn, 2002; Levitt, 1996; Kalt and Zupan, 1984) notes, many factors drive the voting of legislators, from ideological preferences and interactions with concerned constituents to longer-run concerns about re-election and the ability to generate campaign contributions. Whether the link to voting behavior arises because constituents express greater concern for the environment or legislators themselves change their personal views is a topic for future research. That said, the short run nature of our identifying variation does suggest that the effect is not entirely driven by a long-run shift in ideological preferences or a desire to demonstrate a *consistent* pro-environment stance to voters.

2.6 Conclusion

Anthropogenic climate change remains a societal threat and major policy challenge. Public opinion on the existence and severity of climate change has fluctuated considerably over recent decades. Forming accurate beliefs about a long-term one-time event such as climate change places an enormous informational burden on the actor. Unusual weather is an observable, short-term analog that could be used to update one's opinion regarding climate change.

This paper tests the extent to which the salience of climate change is affected by such short-run weather deviations. We use Google Insights search data to proxy for salience, which allows us to perform our analysis at the state-week level. We find that search intensity does indeed respond to weather deviations. Further, the high temporal resolution of our data allows us to provide a number of novel insights. The effect of weather on search intensity varies substantially across the seasons. Unusually cold temperatures have a large effect only in the fall and winter; unusually warm weeks are associated with increased search only in the winter and summer. There does not appear to be much of a relationship between spring weather and search.

We demonstrate that similar patterns exist in the environmental voting record of members of the U.S. Congress. We find that members, and in particular Democrats, are more likely to vote in favor of environmental legislation when their home state experiences anomalous weather or high search activity related to global warming and climate change. The effect of unusual weather is stronger for environmental regulation closely related to climate change or industrial emissions than environmental regulation unrelated to industrial or carbon policy and absent for votes unrelated to environmental policy. In addition, the effects are less strong for "close" votes for which political concerns and vote coordination by party leadership seem to outweigh the effects of unusual weather. While modest in size, the results provide an important, policy-relevant link between anomalous weather and observable action on environmental issues. In addition, the results suggest that search activity may be a useful proxy for the salience of particular policy issues, an important political consideration that is typically difficult to assess.

2.7 Figures

Figure 2.1: Average temperature deviations, 1974-2011

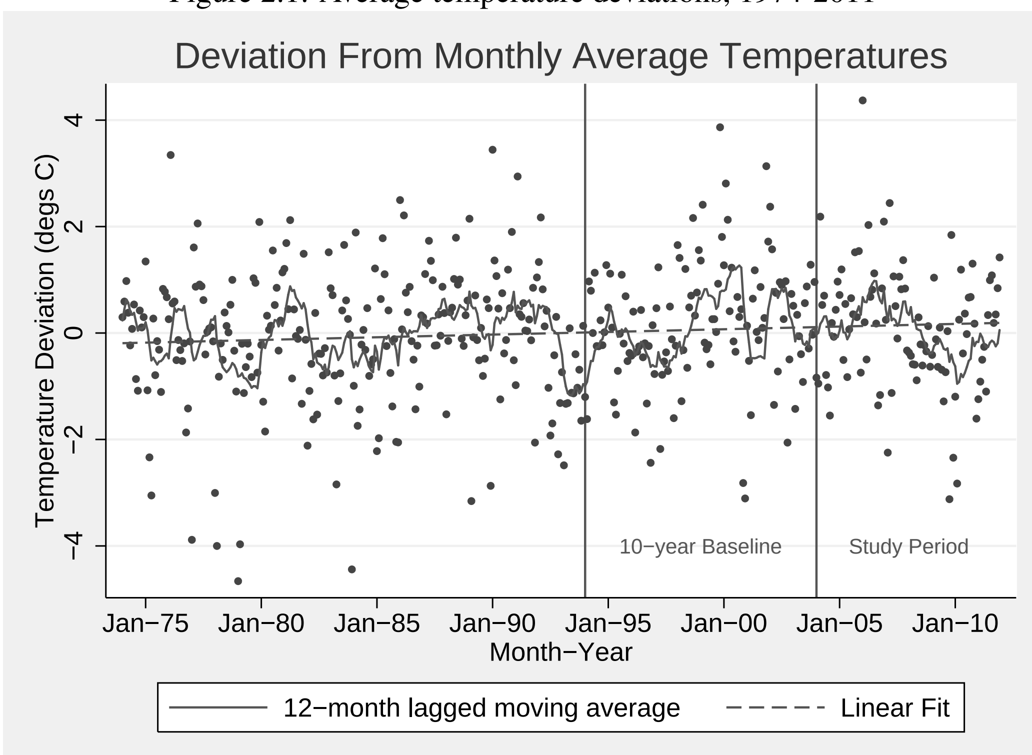

Figure 2.2: Plot of residuals: Colorado, Oct. 2006-Apr. 2007

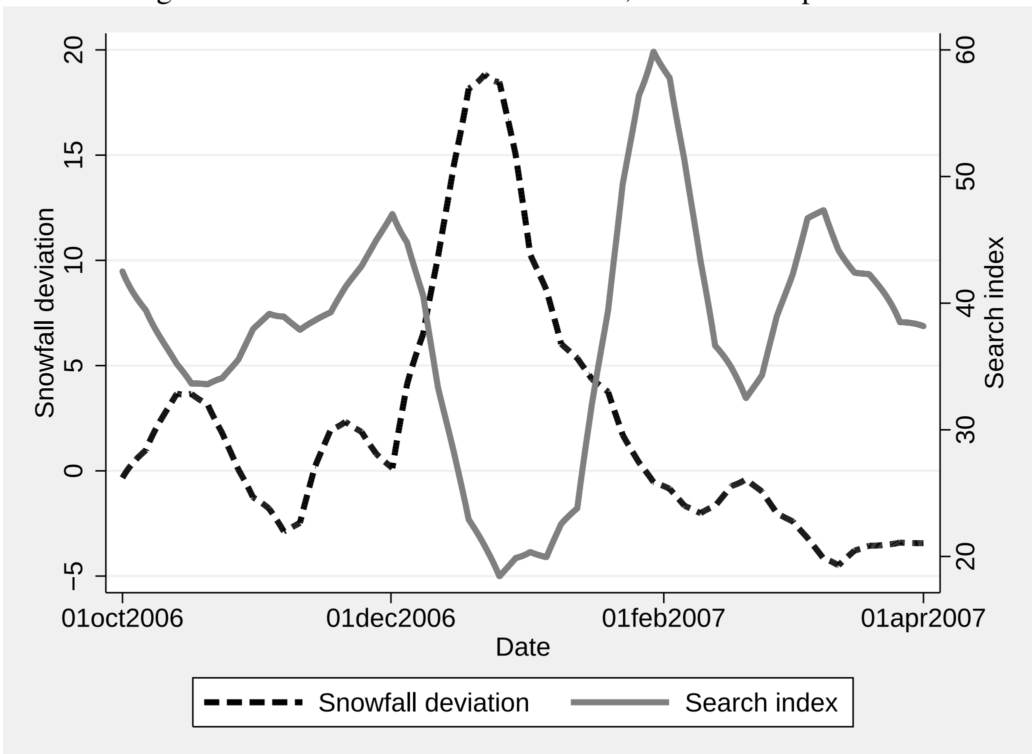

Figure 2.3: All climate-related searches compared to skeptical searches

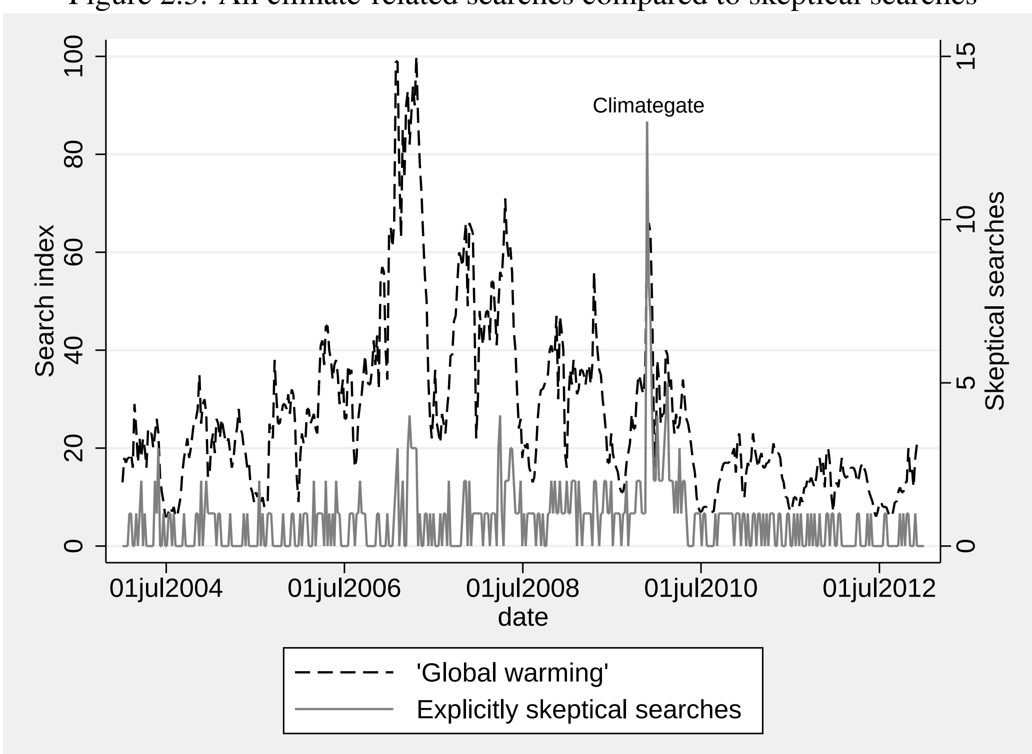

Figure 2.4: Environmental vote share for LCV-tracked votes

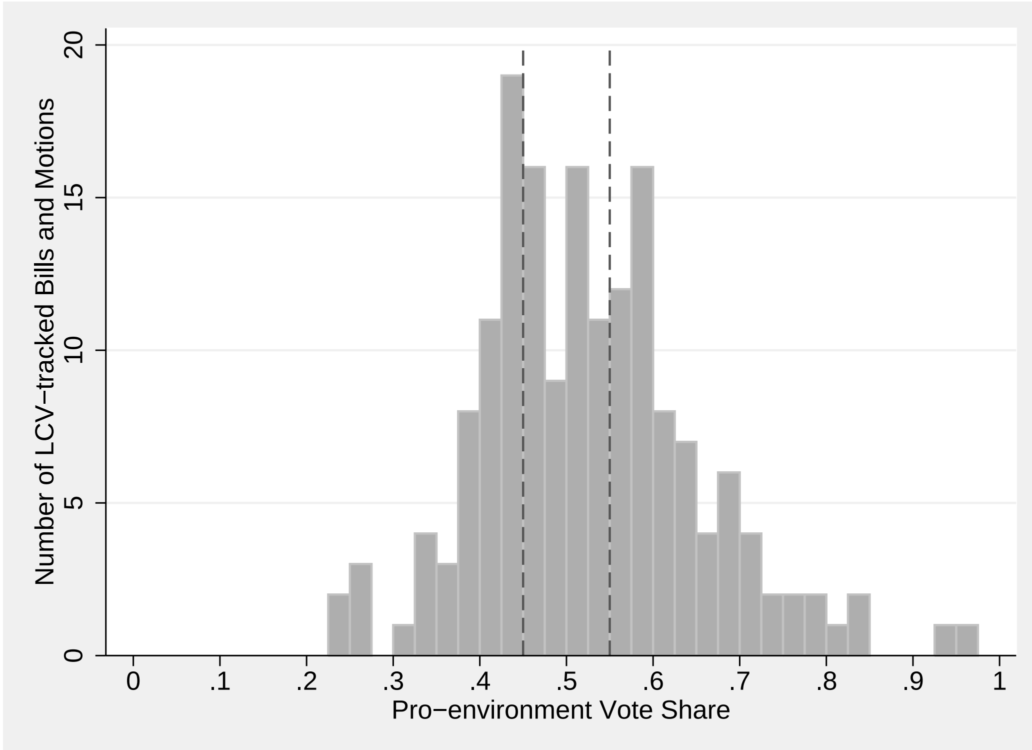

Figure 2.5: Estimated effect of search on voting by member's overall LCV score

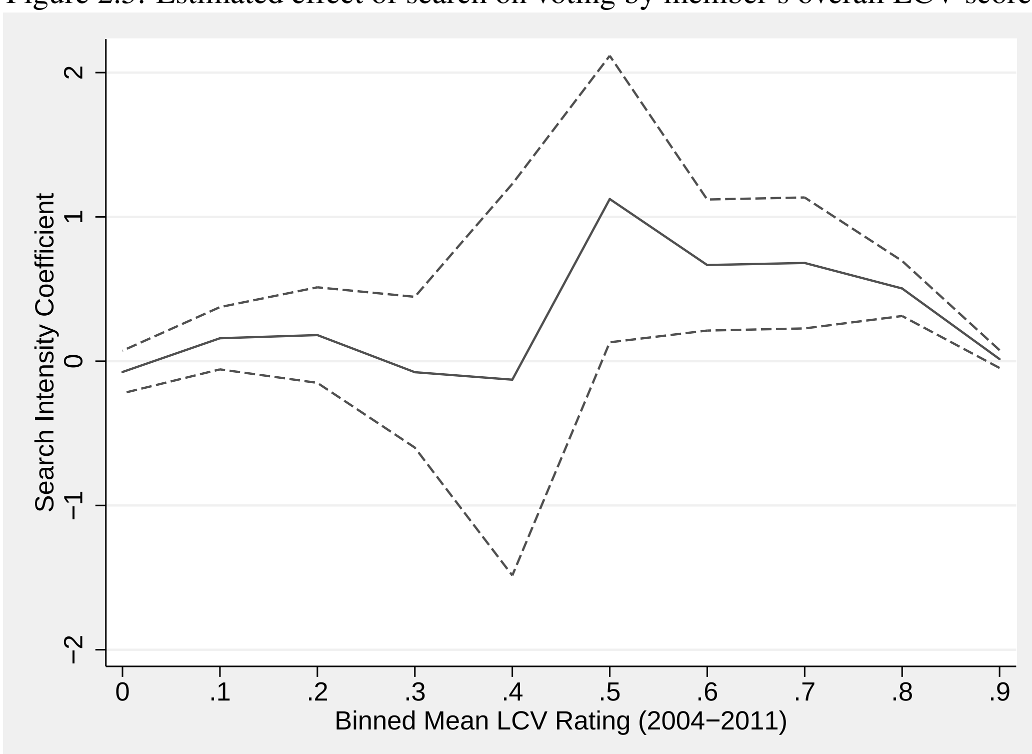

2.8 Tables

Table 2.1: Descriptive Statistics, Full Sample

		Max. Temp. (°C)	Precip. (mm)	Snowfall (mm)	Snow Depth (mm)	Google Search Index
Full sample:	Mean	-0.0331	0.1044	0.4892	4.3219	39.651
(N = 16546)	SD	3.2633	2.7115	6.0449	42.619	28.284
Winter:	Mean	-0.7083	-0.0260	1.7908	14.411	43.021
(N = 4269)	SD	4.0552	2.4107	10.141	73.259	29.640
Spring:	Mean	0.2705	0.1031	0.0685	2.8974	47.101
(N = 4320)	SD	3.2284	2.7223	4.8793	37.368	30.495
Summer:	Mean	0.1994	0.1362	0.0006	0.2317	23.318
(N = 3485)	SD	2.1956	2.6152	0.0525	3.7803	16.440
Fall:	Mean	0.1371	0.2055	0.0337	-0.7467	41.964
(N = 4472)	SD	3.0479	3.0242	3.4447	10.026	27.114

Notes: All weather variables are deviations from the 10-year baseline covering 1994-2003. Sample period is from 2004-2011.

Table 2.2: Weather correlations

	Max temp	Precip	Snowfall
Max temp	·	·	·
Precip	-0.1077	·	·
Snowfall	-0.3204	0.1248	·
Snow depth	-0.2704	0.0312	0.4760

Notes: All weather variables are deviations from the 10-year baseline covering 1994-2003. Sample period is from 2004-2011.

Table 2.3: Effect of weather deviations on search intensity

Panel A: Linear Specification

	(1) All Seasons	(2) Winter	(3) Spring	(4) Summer	(5) Fall
Max Temp, deg. C	-0.240***	-0.654***	-0.074	0.232*	-0.046
	(0.064)	(0.112)	(0.089)	(0.120)	(0.088)
Precip., mm	-0.007	0.177*	0.020	0.106*	-0.026
	(0.047)	(0.097)	(0.086)	(0.058)	(0.091)
Snowfall, mm	-0.042*	-0.094***	0.025	5.690*	0.065
	(0.021)	(0.023)	(0.052)	(3.384)	(0.090)
Snow Depth, mm	-0.018*	-0.024**	-0.001	-0.009	-0.118***
	(0.009)	(0.009)	(0.022)	(0.020)	(0.040)
Constant	23.781***	22.758***	62.722***	93.443***	58.605***
	(1.541)	(1.572)	(1.256)	(1.974)	(1.607)
Observations	16,546	4,269	4,320	3,485	4,472
R-squared	0.761	0.684	0.780	0.781	0.737

Panel B: Asymmetric Specification

	(1) All Seasons	(2) Winter	(3) Spring	(4) Summer	(5) Fall
Pos dev, Max Temp, deg. C	0.292***	0.547***	-0.110	0.707***	0.082
	(0.078)	(0.155)	(0.142)	(0.163)	(0.178)
Neg dev, Max Temp, deg. C	0.806***	1.634***	-0.026	0.322*	0.246*
	(0.101)	(0.172)	(0.139)	(0.192)	(0.133)
Pos dev, Precip., mm	0.057	0.844***	-0.096	0.018	-0.080
	(0.073)	(0.160)	(0.133)	(0.074)	(0.110)
Neg dev, Precip., mm	0.122	1.048***	-0.311	-0.285**	-0.159
	(0.124)	(0.244)	(0.203)	(0.139)	(0.161)
Pos dev, Snowfall, mm	0.003	-0.047	0.024	7.961*	0.132*
	(0.029)	(0.029)	(0.066)	(4.055)	(0.070)
Neg dev, Snowfall, mm	0.285***	0.256**	-0.039	11.321	0.774**
	(0.088)	(0.101)	(0.117)	(29.207)	(0.317)
Pos dev, Snow Depth, mm	-0.007	-0.003	-0.032*	0.027	0.019
	(0.009)	(0.009)	(0.019)	(0.037)	(0.055)
Neg dev, Snow Depth, mm	0.045**	0.076***	-0.067	0.540	0.300**
	(0.020)	(0.023)	(0.060)	(0.527)	(0.122)
Constant	21.408***	16.192***	27.028***	90.708***	60.040***
	(1.662)	(1.874)	(1.218)	(2.086)	(1.416)
Observations	16,546	4,269	4,320	3,485	4,472
R-squared	0.763	0.696	0.781	0.783	0.741

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Dependent variable is the Google search index. All regressions also include year * month FE and state * month of year FE. Standard errors are clustered at the state level.

Table 2.4: Environmental Votes, Local Weather and Search Intensity

	(1)	(2)	(3)
Panel A: Weather variables			
Pos Dev, Max Temp	0.00004	0.00032	-0.00092
	(0.00208)	(0.00195)	(0.00166)
Neg Dev, Max Temp	-0.00467***	-0.00588***	-0.00257
	(0.00148)	(0.00150)	(0.00159)
Pos Dev, Snowfall	-0.00048	-0.00024	-0.00040
	(0.00038)	(0.00036)	(0.00032)
Neg Dev, Snowfall	0.00385**	0.00265	0.00190
	(0.00189)	(0.00181)	(0.00186)
Pos Dev, Precipitation	0.00194	0.00075	0.00053
	(0.00139)	(0.00128)	(0.00124)
Neg Dev, Precipitation	-0.00116	0.00024	-0.00113
	(0.00234)	(0.00219)	(0.00260)
Pos Dev, Snow Depth	0.00002	0.00003	0.00004
	(0.00006)	(0.00007)	(0.00005)
Neg Dev, Snow Depth	0.00012	0.00013	-0.00001
	(0.00021)	(0.00020)	(0.00021)
F-test p-value	$<$0.001	$<$0.001	0.147
Observations	61173	61173	61173
R-Squared	0.654	0.657	0.672
Panel B: Weather-correlated with Search Intensity			
Climate Change Search Intensity/100	0.313***	0.254***	0.111**
	(0.0601)	(0.0557)	(0.0475)
Observations	61148	61148	61148
R-Squared	0.655	0.657	0.672

Notes: *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is a binary variable indicating whether a representative voted for the LCV-endorsed position. All specifications include representative fixed effects. In addition, column (2) includes year * month fixed effects and column (3) includes year * week fixed effects. Standard errors are clustered at the state level.

Table 2.5: ACU Votes, Local Weather and Search Intensity

	All ACU-tracked votes			Same-week ACU-tracked votes		
	(1)	(2)	(3)	(4)	(5)	(6)
Climate Change Search Intensity/100	-0.0253	-0.00977	0.0500*	0.0249	-0.0165	-0.00344
	(0.0258)	(0.0265)	(0.0280)	(0.0349)	(0.0375)	(0.0494)
Observations	90143	90143	90143	41509	41509	41509
R-Squared	0.551	0.554	0.569	0.542	0.549	0.573

Notes: *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is a binary variable indicating whether a representative voted for the ACU-endorsed position. All specifications include representative fixed effects. In addition, columns (2) and (5) include year * month fixed effects and columns (3) and (6) include year * week fixed effects. Standard errors are clustered at the state level.

Table 2.6: Environmental Votes and Search Intensity, by Representative Characteristics

	(1)	(2)	(3)
Climate Change Search Intensity/100	0.167**	0.104	-0.0584
	(0.0770)	(0.0742)	(0.0684)
Senate * Search Intensity/100	-0.111	-0.131*	-0.0461
	(0.0709)	(0.0712)	(0.0687)
Democrat * Search Intensity/100	0.276***	0.289***	0.298***
	(0.0959)	(0.0950)	(0.0888)
Observations	61148	61148	61148
R-Squared	0.655	0.658	0.672

Notes: *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is a binary variable indicating whether a representative voted for the LCV-endorsed position. All specifications include representative fixed effects. In addition, column (2) includes year * month fixed effects and column (3) includes year * week fixed effects. Standard errors are clustered at the state level.

Table 2.7: Environmental Votes and Search Intensity, by Vote Characteristics

	(1)	(2)	(3)
Other Vote * Search Intensity/100	0.295***	0.213***	-0.0191
	(0.0710)	(0.0660)	(0.0667)
Industrial Regulation * Search Intensity/100	0.370***	0.316***	0.310***
	(0.0783)	(0.0779)	(0.0853)
Climate Change * Search Intensity/100	0.290***	0.283**	0.395**
	(0.102)	(0.112)	(0.172)
Close vote * Search Intensity/100	-0.0969***	-0.0654***	-0.101***
	(0.0172)	(0.0159)	(0.0211)
Observations	61148	61148	61148
R-Squared	0.655	0.658	0.672

Notes: *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is a binary variable indicating whether a representative voted for the LCV-endorsed position. All specifications include representative fixed effects. In addition, column (2) includes year * month fixed effects and column (3) includes year * week fixed effects. Standard errors are clustered at the state level.

CHAPTER 3.

Seller Commitment and the Empirical Analysis of First-Price Auctions

3.1 Abstract

The empirical auction literature focuses on recovering firm valuations, which are used to calculate surplus division, optimal reserve prices, and the impacts of an increased number of competing bidders. In the case of a binding reserve price, the standard assumption is that the auctioneer keeps the object if the reserve price is met. In this paper, I consider a model in which the seller may re-auction object once if the initial reserve price is not met. I argue that the value distributions are nonparametrically identified without additional modeling assumptions using standard arguments in the literature. However, the distribution of shocks to values between the first- and second-round auctions is not identified using only bid data. I impose a parametric assumption and propose a semiparametric estimation procedure. A Monte Carlo simulation shows that the estimator performs well, and demonstrates that ignoring the seller commitment problem will lead to biased estimates of bidder valuations. Finally, I discuss how previous methods using reserve prices can be adapted to identify the first-round distribution of unobserved heterogeneity and allow it to persist into the second round. Future work will relax some of the model's assumptions and apply the estimator to data from timber auctions.

3.2 Introduction

Sellers will often use an auction in order to organize and increase competition for their assets. In this context, features such as reserve prices and entry fees can be used to improve the expected payoff to the seller. Much of the theoretical literature and all of the empirical literature has assumed that the potential bidders believe the seller will permanently withhold the object in the event that no bidder meets the reserve price.

Unfortunately, this assumption fails in many important settings, especially in the context of auctions for natural resource extraction rights. Porter (1995) notes that 47% of federal offshore oil and gas leases that do not receive a sufficient bid are eventually re-auctioned. According to McAfee and Vincent (1997), the US Forest Service responds to a no-bid auction by re-auctioning the tract and reducing the reserve price by 10 percent. The Michigan Department of Natural Resources (DNR) typically reduces a timber sale's reserve price by roughly 20 to 40 percent in response to a no-bid auction (Heym, 2013). Beyond natural resources, another example of re-auctioning arises in FDIC sales of distressed real estate assets (McAfee, Quan, and Vincent, 2002) that do not receive bids at the initial reserve price. The theoretical literature that *has* addressed this issue has shown that the optimal choice of auction parameters can be considerably different when this assumption is relaxed.

In this paper, I consider the importance of this assumption in the empirical analysis of first-price auctions. I specify and solve a two-round model in which potential bidders decide whether to participate in a first-round auction based on the probability that the item may be re-auctioned in the future. Using tools from the structural empirical auction literature, I find that bidder valuations are nonparametrically identified identified in my model, as is the cutoff valuation that makes a bidder indifferent between bidding and waiting. However, I find that if valuations receive shocks between the first- and second-round auctions, the distribution of those shocks is not nonparametrically identified. Instead, I explore an alternative semiparametric approach and test its performance using a Monte Carlo simulation. I also describe a potential application to Michigan state timber auctions and discuss how first-round reserve prices can be used to identify the distribution of persistent unobserved heterogeneity. In future work, I will extend the model, apply the estimation procedure to the Michigan timber data, and examine the implications of the model for optimal auction design.

When bidders know the object may be re-auctioned, some may decide that it is better not to bid, in hopes that they are able to bid in a second auction in which they have a (possibly) higher valuation, against weaker revealed competition, and a lower reserve price. The possibility of waiting for a future auction will bias naive estimates of the value distribution. Intuitively, without re-auctioning, a bidder with a value at the reserve price is indifferent between bidding and not bidding. When re-auctioning is possible, the bidder who is indifferent between bidding and waiting will typically have a value strictly *above* the reserve price. Thus, the model will make different predictions about how bids near the reserve price are mapped back into their corresponding values.

Further, bidding behavior in the second-round auction will be conditioned on the knowledge that competitors' values are adversely selected from the original value distribution. While non-parametric techniques that invert bids into values using necessary conditions (Guerre, Perrigne, and Vuong, 2000) will still recover the correct valuations for both rounds of the auction, these techniques will not provide sufficient information for the evaluation of most counterfactual policies in this context. For example, these dynamic considerations change the calculus of important outcomes such as the optimal reserve price and the total surplus in the market; the distribution of between-period shocks is a key element to these analyses. Any changes to the mechanism that affects the types of firms willing to eschew bidding in the first round (for instance, changes in number of potential bidders, entry fees, or reserve prices) requires estimates of how equilibrium bidding will change, and what the new distribution of second-round values will be.

As noted by Athey, Cramton, and Ingraham (2002) in their guide to setting reserve (upset) prices in British Columbia Timber auctions,

> "When a tract does not sell at auction, the Ministry acquires certain information. If bidders' values for a particular tract are forever constant, and if the bidders are bidding under the equilibrium described above, then the Ministry learns that the value of the tract is less than the upset it set...overall, the expected future resale value of a tract that fails to sell today depends on the correlation between today's values and tomorrow's values by firms...if bidder values are correlated over time, little is gained by setting a

high upset price and continually reoffering a tract in the hopes it will sell at that high price."

This excerpt emphasizes the importance of recovering the relationship between bidder valuations for a tract today and valuations for a tract at a future re-auctioning date. The correlation between first- and second- price values is necessary for any sort of counterfactual policy simulation.

If the government cannot commit to keeping the object off the market after the first auction, it needs to set an optimal *pair* of reserve prices. The first-round reserve price obviously affects revenue in the first round; however, it also changes the distribution of firms associated with an auction that receives no bid, and thus the maximum revenue that can be extracted in the case that a second-round auction is held. Similarly, the second-round reserve price affects both first- and second-round bidding behavior.

I discuss applying the estimation procedure to Michigan DNR timber auctions, and future work will carry out the full analysis. The data are quite rich: all firm identities, bids, and a comprehensive set of observable characteristics are available. Of the 4616 first-round auctions, 480 receive no bids. This is a modest share, but it provides enough data to estimate a distribution of second-auction bids. Further, the possibility of multiple rounds could also substantially affect bidding behavior in the 742 auctions that only receive one bid and the 846 auctions that only receive two. The DNR's reserve price process is relatively well-defined. The general procedure following a no-bid auction is to post the sale at the original reserve price for 60 days, then re-auction (roughly 6 months after the original auction) with a reserve price that is 20 to 40 percent lower than before (Heym, 2013). Given that the bidders know that a no-bid auction is very likely to be sold again, this setting is an example in which we would expect bidding to be influenced by the seller commitment problem.

This paper relates to several strands of the existing auction literature. First, there is a theoretical literature pertaining to the design of optimal auctions when the auctioneer cannot commit to the design of the auctions. An early paper in this area is by McAfee and Vincent (1997), who characterize the equilibrium sequence of reserve prices and show that revenues are eroded by this inability to commit in a sense similar to the Coase conjecture. Burguet and Sákovics (1996) show that when value discovery is costly, an auctioneer can increase welfare and revenues by announcing that a second auction without a reserve price will be held if no bids are submitted. More recently, Skreta (2013) extended McAfee and Vincent's result when the seller can choose any mechanism. Vartiainen (2011) considers an even more extreme version of non-commitment: the seller cannot even commit to honor the *current* mechanism; that is, a valid bid above the reserve price might not result in a transaction.

My approach complements and extends these studies in two ways. First, I approach the problem from an empirical angle. I do not assume that sellers set optimal reserve prices; however, my structural model assumes bidding behavior that is optimal when bidders account for the re-auctioning of an unsold contract. Second, my model allows for shocks to bidder valuations over time. The aforementioned papers assume that bidders have constant valuations. In an empirical context, this is not realistic: firm and market conditions change over time, and bidding behavior in a follow-up auction will reflect this. Future work will use structural estimates to calculate the optimal seller strategies characterized in the existing theoretical work, and empirically quantify the importance of the commitment problem.

A second strand of related literature is the literature on the structural empirical analysis of first-price auction data. The approach pioneered by Guerre, Perrigne, and Vuong (2000) uses necessary conditions to map observed bids into valuations without directly calculating the equilibrium of the game. This approach has been generalized to account for dynamics (Jofre-Bonet and Pesendorfer, 2003; Balat, 2013), endogenous participation (Xu, 2013), and unobserved auction heterogeneity (Krasnokutskaya, 2009; Roberts, 2013). I will contribute to this literature by showing how to extend this approach to the common case where a seller cannot credibly commit to keep an object off the market.

Finally, this issue is of particular interest in the context of natural resource extraction. As noted above, sales of timber stands and oil/natural gas leases often go unsold and are re-auctioned. These are the settings of many economic policies of consequence. Analyses based on bid data in this case could be limited if there is a strategic incentive to withhold one's bid. I consider this in the context of timber, but the approach could be applied in other contexts, such as oil lease auctions, perhaps exploiting recent identification approaches for auctions with interdependent costs (Somaini, 2014).

3.3 Model Setup and Equilibrium

To empirically address these issues, I specify a simple model of a first-price auction with (potentially) two rounds of bidding. Bidders are *ex ante* homogeneous, and draw independent private values. If no one is willing to meet the reserve price in the first round, the object is re-auctioned in the next period. If no one is willing to meet the reserve price in *this* auction, then the object remains with the auctioneer, as is assumed in a typical one-period model.[1]

3.3.1 Model

The model proceeds under the following assumptions:

1. **Bidder valuations:** Bidder values in the first round are independent private values, $v_{1i} \sim F_1$. Bidders receive i.i.d. shocks, $\omega \sim F_\omega$, to their valuations between rounds such that $v_2(v_{1i}, \omega_i) \sim F_2$, with $E[v_{2i}] = v_{1i}$.
2. **Reserve Prices:** $R_2 \leq R_1$ (the auctioneer sets a lower reserve in the second round).
3. **Discounting:** Bidders and the auctioneer discount the second period at rate δ.
4. **Timing:**
   (a) N potential bidders observe reserve price R_1 their own valuations v_{1i}, and the distribution of other bidders' valuations F_1.
   (b) Bidders submit a bid $b_1(v_{1i})$ in a sealed-bid first-price auction.
   (c) If no bids above R_1 are received, the auctioneer sets a new reserve price R_2.
   (d) The same N potential bidders receive shocks ω_i, learn their new values $v_2(v_{1i}, \omega_i)$, and observe the distribution of shocks F_ω.
   (e) Bidders submit bids $b_2(v_{2i})$ in a second-round sealed-bid first-price auction.

[1] In section 3.8, I discuss relaxing several of these assumptions.

I restrict the analysis to increasing symmetric bidding equilibria within each auction. The full equilibrium will be a first-round cutoff value, v^*; and two bidding functions, b_1 and b_2. Typically, v^* will exceed the reserve price R_1. This is driven partially by the option value associated with the shocks to valuations; that is, a bidder with valuation just above R_1 will be willing to forgo their small expected profit from bidding in the first-round auction in favor of receiving a (possibly favorable) shock to their valuation. In the absence of value innovations, a lower second-round reserve price could also be enough to push v^* above R_1.

In this context, the cutoff value will be defined as the bidder that is indifferent between bidding and not bidding in the first round.

$$\begin{aligned}(v^* - R_1)F_1(v^*)^{N-1} &= \delta E[(v_2 - b_2(v_2))F_2(v_2|v_1 < v^*)^{N-1}]F_1(v^*)^{N-1} \\ v^* - R_1 &= \delta E[(v_2 - b_2(v_2))F_2(v_2|v_1 < v^*)^{N-1}] \qquad (3.1)\end{aligned}$$

where the expectation is taken over the second-period valuation conditional on having a value $v_1 = v^*$. Given v^*, the equilibrium bid function is as follows:

$$b_1(v_1) = v_1 - \frac{\int_{v^*}^{v_1} F_1(u)^{N-1}du}{F_1(v_1)^{N-1}} - (v^* - R_1)\frac{F_1(v^*)^{N-1}}{F_1(v_1)^{N-1}} \qquad (3.2)$$

which, conditional on v^*, is identical to the equilibrium bid function in Samuelson (1985).[2]

The second-round equilibrium is standard for a first-price auction with a binding reserve price; note that the bidders update to account for the adversely selected set of competitors:

$$b_2(v_2) = v_2 - \frac{\int_{R_2}^{v_2} F_2(u|v_1 < v^*)^{N-1}du}{F_2(v_2|v_1 < v^*)^{N-1}} \qquad (3.3)$$

3.3.2 Proof of Cutoff Equilibrium

Going forward, I will operate under an additional structural assumption:

[2]In his model, bid preparation costs drive the wedge between the reserve price and v^*.

5. **Multiplicative shock:** $v_2 = v_1\omega$, and $\omega \in [0, \infty)$.

For expository purposes, I initially assume that $E[\omega] = 1$; however, I will present a more general condition on the distribution of ω that is sufficient to guarantee an equilibrium.[3]

The equilibrium strategies have been defined in the previous subsection. Here, I show that these strategies actually admit a cutoff equilibrium under certain conditions. Specifically, I will show that the derivative of the gain from bidding in the first round rather than waiting is positive for all values of v_1; thus, the gain is increasing in the bidder's first-round valuation.

Consider a potential bidder with some valuation $v \geq R_1$. Suppose that all other agents are playing the strategies outlined in Section 3.3. Then the bidder's best expected payoff from bidding in the first auction is

$$\begin{aligned} (v_1 - b_1(v_1))F_1(v_1)^{N-1}, \qquad & \text{if } v_1 \geq v^* \\ (v_1 - R_1)F_1(v^*)^{N-1}, \qquad & \text{if } v_1 \in [R_1, v^*). \end{aligned}$$

The derivative of this payoff with respect to v is $F(v)^{N-1}$ if $v \geq v^*$ and $F_1(v^*)^{N-1}$ if $v \in [R_1, v^*)$.[4] Note that this derivative is always $\geq F(v^*)^{N-1}$. The expected payoff from waiting for a second auction is:

$$\delta F(v^*)^{N-1} \int_{\frac{R_2}{v_1}}^{\infty} (v_1\omega - b_2(v_1\omega))F_2(v_1\omega; v^*)^{N-1} dF_\omega.$$

[3]In Appendices C.1 and C.2, I also establish that the equilibrium exists with any additive shock or in the absence of a shock.

[4]$\frac{\partial}{\partial v_1}(v_1 - b_1(v_1))F_1(v_1)^{N-1} = F_1(v_1)^{N-1} - b_1'(v_1)F_1(v_1)^{N-1} + (N-1)(v_1 - b_1(v_1))F_1(v_1)^{N-2} f_1(v_1)$, which is equal to $F_1(v_1)^{N-1}$ because the last two terms are equal to the first-order condition for a payoff-maximizing bid.

When differentiated with respect to v_1, this gives:

$$
\begin{aligned}
&\delta F(v^*)^{N-1} \int_{R_2 - v_1}^{\infty} \{\omega(1 - {b_2}'(v_1\omega))F_2(v_1\omega; v^*)^{N-1} \\
&\qquad\qquad + \omega(v_1\omega - b_2(v_1\omega))F_2(v_1\omega; v^*)^{N-2} f_2(v_1\omega; v^*)(N-1)dF_\omega\} \\
&\qquad\qquad + \frac{R_2}{v_1^2} \underbrace{(R_2 - b_2(R_2))}_{=0} F(R_2)^{N-1} \\
&= \delta F(v^*)^{N-1} \int_{R_2 - v_1}^{\infty} \omega F_2(v_1\omega; v^*)^{N-1} \\
&\qquad\qquad - \omega {b_2}'(v_1\omega)F_2(v_1 + \omega; v^*)^{N-1} \\
&\qquad\qquad + \omega(v_1 + \omega - b_2(v_1\omega))F_2(v_1 + \omega; v^*)^{N-2} f_2(v_1\omega; v^*)(N-1)dF_\omega\} \\
&= \delta F(v^*)^{N-1} \int_{R_2 - v_1}^{\infty} \omega F_2(v_1\omega; v^*)^{N-1} dF_\omega
\end{aligned}
$$

The final equality holds because the last two lines of the previous expression equal zero: they are ω times the first order condition for optimal bidding in the second-round auction for a given value of $v_2 = v_1\omega$. Now I can confirm that the relative gains from bidding are increasing in v_1:

$$
\begin{aligned}
\frac{d}{dv_1} E[\text{Payoff from bidding in auction 1}] &\geq F(v^*)^{N-1} \\
&> \delta F(v^*)^{N-1} \\
&= \delta F(v^*)^{N-1} E[\omega] \\
&> \delta F(v^*)^{N-1} E\left[\omega | \omega \geq \frac{R_2}{v_1}\right] \left[1 - F_\omega(\frac{R_2}{v_1})\right] \\
&> \delta F(v^*)^{N-1} \int_{\frac{R_2}{v_1}}^{\infty} \omega F_2(v_1\omega; v^*)^{N-1} dF_\omega \\
&= \frac{d}{dv_1} E[\text{Payoff from waiting for auction 2}]
\end{aligned}
$$

Thus, given that a bidder with $v_1 = v^*$ is indifferent between bidding in auction 1 and not bidding, all $v < v^*$ prefer to wait, and all $v > v^*$ prefer to bid in auction 1. Comparing the second line to the fourth line, a more general sufficient condition for the equilibrium arises:

$$1 \geq \underbrace{\delta E\left[\omega|\omega \geq \frac{R_2}{v_1}\right]\left[1 - F_\omega(\frac{R_2}{v_1})\right]}_{\text{“Contribution” of right tail to } E[\omega]} = \delta \int_{\frac{R_2}{v_1}}^{\infty} \omega dF_\omega \geq \delta \int_{\frac{R_2}{v_1}}^{\infty} \omega F_2(v_1\omega; v^*)^{N-1} dF_\omega \tag{3.4}$$

The second term (labeled with the brace below) is the expectation integral of ω truncated below at R_2/v_1. It is increasing in v_1 and approaches $E[\omega]$ as $v_1 \to \infty$.[5] In fact, as $v_1 \to \infty$, $F_2(v_1\omega; v^*)^{N-1} \to 1 \ \forall\omega$ and, therefore:

$$\lim_{v_1 \to \infty} \delta F(v^*)^{N-1} \int_{\frac{R_2}{v_1}}^{\infty} \omega F_2(v_1\omega; v^*)^{N-1} dF_\omega = \delta E[\omega] \tag{3.5}$$

Thus, a more general sufficient condition is that $E[\omega] < 1/\delta$. Essentially, if the first-round bidders' discounted expectation of their second-round valuations is less than their current valuation, there is a cutoff equilibrium. Depending on the specifics of reserve prices, the discount rate, and the value distributions, this cutoff might be below R_1, in which case all bidders with valuations above the first-period reserve price would bid. Interestingly, the validity of the cutoff equilibrium does not directly depend on the second-period reserve price, because the potentially problematic valuations are in the right tail of the distribution.[6]

3.4 Identification

From this point forward, I assume that the increasing symmetric cutoff equilibrium outlined in Section 3.3 exists and is played. The structural objects of interest are the first-round distribution of values (F_1), the second-round distribution of values (F_2), and the distribution of innovations linking the right-truncated version of F_1 with F_2 (i.e., F_ω). In this section, I find that F_1 and

[5] The expression is decreasing in the lower limit of integration: $\frac{\partial}{\partial a} E\left[\omega|\omega \geq a\right]\left[1 - F_\omega(a)\right] = \frac{\partial}{\partial a}\int_a^\infty \omega dF_\omega = -af_\omega(a) < 0$. Since the lower limit is decreasing in v_1, the expression is increasing in v_1.

[6] The actual levels of the equilibrium bid functions and the cutoff value will still depend on the relative values of R_1 and R_2. It is simply the case that the monotonicity of the relative payoffs of bidding versus waiting does not directly depend on the reserve prices.

F_2 are nonparametrically identified for valuations above v^* and R_2, respectively, as is the cutoff value v^*. However, because the values in the second period derive from unobserved first-round values ($v_1 < v^*$), there is no information about the density that can be used to deconvolve the two distributions. Because F_ω determines the link between first- and second-round valuations, it is crucial to simulating any counterfactual scenario that affects the cutoff value v^*.[7] The cutoff value v^* links F_1 to F_2 through the indifference condition and does provide one moment that can be used to characterize F_ω. Still, a distributional assumption on F_1 is necessary to fully identify F_ω using the distribution of second-round pseudovalues ξ_2. In Section 3.8, I discuss potential sources of exogenous variation that could identify F_ω.

3.4.1 Nonparametric Identification

The first-round equilibrium of this model is very similar to that of Samuelson (1985). Xu (2013) shows that the Samuelson model is identified for all values above v^*, as is the value of v^*. I use similar arguments to show that components of the present model are identified.

Let $G^*(b)$ be the distribution of bid data on $[R_1, \bar{b}]$, where $\bar{b}$ is the upper bound of the bid distribution. Also, let $p_1 = 1 - F_1(v^*)$ be the probability that a given bidder bids. Then conditional on bidding (i.e., if $v_1 \geq v^*$), the bidder maximizes:

$$(v_1 - b_1)(p_1 G_1^*(b_1) + (1 - p_1))^{N-1}$$

The first order condition gives the following relationship between values and bids:

$$\xi_1(b_1) = b_1 + \frac{1}{N-1} \frac{p_1 G_1^*(b_1) + (1 - p_1)}{p_1 g_1^*(b_1)} \tag{3.6}$$

[7] Such scenarios include changes in the number of potential competitors, the reserve price policy, and the introduction of an entry fee or subsidy.

These pseudovalues identify the distribution F_1 for values greater than v^*. We know that

$$\xi_1(R_1) = v^* \tag{3.7}$$

$$\text{and } G_1^*(R_1) = 0. \tag{3.8}$$

Inverting the bid function at R_1, I get that:

$$v^* = \xi_1(R_1) = R_1 + \frac{(1-p_1)}{p_1 g_1^*(R_1)(N-1)}$$

which, combined with the zero-profit condition from Equation (3.1), implies that:

$$v^* - R_1 = \frac{(1-p_1)}{p_1 g_1^*(R_1)(N-1)} = \delta E\left[(v_2 - b_2(v_2))F_2(v_2|v_1 < v^*)^{N-1}\right] \tag{3.9}$$

Thus, v^* is identified, and can be estimated with a consistent estimate of the bid density at R_1. This involves nonparametric estimation at the boundary of the observed bids. Possible approaches for estimating this object have been discussed in the empirical auction literature by Xu (2013) and Hickman and Hubbard (forthcoming).

Using the approach of Guerre, Perrigne, and Vuong (2000), I can identify the distribution $F_2(v_2|v_1 < v^*)$ for all $v_2 \geq R_2$:

$$\xi_2(b_2) = b_2 + \frac{1}{N-1}\frac{p_2 G_2^*(b_2) + (1-p_2)}{p_2 g_2^*(b_2)} \tag{3.10}$$

where $p_2 = 1 - F_2(R_2)$. However, $\xi_2(R_2) = R_2$ if and only if $g_2^*(b) \to \infty$ as $b \to R_2$ from above. Thus, Guerre, Perrigne, and Vuong (2000) suggest a transformation, $\tilde{b}_2 = \sqrt{b_2 - R_2}$, which leads to:

$$\xi_2(\tilde{b}_2) = R_2 + \tilde{b}_2^2 + \frac{2\tilde{b}_2}{N-1}\frac{p_2 \tilde{G}_2^*(\tilde{b}_2) + (1-p_2)}{p_2 \tilde{g}_2^*(\tilde{b}_2)} \tag{3.11}$$

Given these pseudovalues ξ_1 and ξ_2, one can follow the nonparametric second stage of Guerre, Perrigne, and Vuong (2000) or the recent semiparametric approach in Aryal, Gabrielli, and Vuong (2014). In this latter paper, a parametric value distribution is assumed. The second step chooses the parameters governing the value distribution to best match the distribution of pseudovalues using GMM with the optimal moments (the score function).

Note that if the researcher assumes that there is only one round of bidding possible, then theory predicts that $b_1(R_1) = R_1$. This leads to the transformation discussed above. However, given the equilibrium predicted by my model, this approach will fail. In most cases, the possibility of a second-round auction result in $b_1(v^*) = R_1$, with $v^* > R_1$. The model must incorporate expectations about the second period into the first-round bidding strategy through this cutoff point. In the Monte Carlo simulations, I demonstrate that the usual estimator will generally perform poorly in recovering first-round valuations. However, as long as the researcher acknowledges that bids reflect the possibility of a second-round auction, F_1 can be consistently estimated nonparametrically.

Unfortunately, the full distribution of ω is not nonparametrically identified using only observed bids.[8] Intuitively, the only information about ω comes through the indifference condition on a bidder with a first-period value of v^*. Given estimates of F_2 and F_1, one could assume a one-parameter distribution for ω and use the indifference moment to estimate it. Alternatively, one could assume a two-parameter distribution and add a constraint, such as assuming that valuations do not change on average (i.e., $E[\omega] = 1$). Suppose that ω is log-normally distributed, with location μ and scale σ. Then the mean-one assumption means that the parameters μ and σ are univariate functions of the variance of the distribution, which could potentially be estimated using the indifference condition for a bidder with the already-identified cutoff value v^*.

3.4.2 The Importance of Estimating F_ω

The calculation of the optimal reserve price clearly illustrates the importance of estimating the relationship between first- and second-round valuations. The second-round optimal reserve price is straightforward; it is simply the standard formula, accounting for the impact of the first-round

[8]In the conclusion, I discuss the possibility of using an instrumental variable approach to gain identification.

reserve price on the second-round value distribution. Taking the first-order condition of the auctioneer's expected revenue and rearranging gives:

$$\begin{aligned}
&\max_{R_2} \quad v_0 F_2(R_2;R_1)^N + N\int_{R_2}^{\bar{v}} [vF_2(v;R_1)^{N-1} - \int_{R_2}^{v} F_2(u;R_1)^{N-1}du]\ f_2(v;R_1)dv \\
&\Rightarrow R_2^*(R_1) = v_0 + \frac{1-F_2(R_2^*;R_1)}{f_2(R_2^*;R_1)}
\end{aligned}$$

However, the first-round is much different. Since the first-round reserve price affects second-round seller profits and endogenously determines the cutoff bidder type, the seller's optimization problem (and first-order condition) is more complicated:

$$\begin{aligned}
&\max_{R_1}\ E_1[\pi_{0,2}(R_1)]F_1(R_1)^N + N\int_{v^*}^{\bar{v}} [v_1F_1(v_1)^{N-1} - \int_{v^*}^{v_1} F_1(u)^{N-1}du - (v^*-R_1)F(v^*)^{N-1}]\ f_1(v_1)dv_1] \\
&\quad \text{s.t.} \quad v^* = R_1 + \delta E[(v_2 - b_2(v_2))F_2(v_2;R_1)^{N-1}]
\end{aligned}$$

where $\pi^*_{0,2}(R_1)$ is the seller's optimal expected second-round payoff, given that R_1 is chosen in the first round. The first source of complication is that the cutoff valuation is endogenously determined, which affects the probability that there are no bids. The second source of complication is that the cutoff value affects the distribution of firms bidding in the case of a second-round auction. Therefore, the seller's no-bid payoff is also an endogenous function of the reserve price choice, which is not normally the case.

The transition probabilities between first-round and second-round valuations are crucial here because they govern the impact of the first-period reserve price on (1) the lowest type willing to bid in the first-period and (2) the distribution of types in the second period, conditional on the cutoff.

3.4.3 A Semiparametric Approach

Adding a constraint, such as $E[\omega] = 1$, is a substantive assumption, and not simply a normalization. Rather than impose this type of constraint, I take a semiparametric approach similar to that developed by Aryal, Gabrielli, and Vuong (2014). I assume parametric distributions for F_1 and F_ω, with associated parameter vectors θ_1 and θ_ω. Conditional on the estimate of v^* and a guess of θ_1 and θ_ω, I can write the likelihood of observing pseudovalues $\vec{\xi_1}$ and $\vec{\xi_2}$. Estimates of θ_1 and θ_ω can be obtained using maximum likelihood or GMM with optimal moments (the score function). The estimation procedure follows these steps, where estimates of variables and distributions are denoted with hats:

1. Calculate the empirical probabilities of submitting a bid, $\hat{p}_1$ and $\hat{p}_2$.
2. Nonparametrically estimate the density and distribution of first- and second-round bids ($\hat{g}_1^*$, $\hat{G}_1^*$, $\hat{g}_2^*$, $\hat{G}_2^*$) using a kernel estimator.
3. Plug these estimated bid densities and distributions into Equations 3.6 and 3.11 to recover the pseudovalues $\vec{\xi_1}$ and $\vec{\xi_2}$ for the first- and second-round auctions.
4. Obtain an estimate of the first-round density of bids at the reserve price ($\hat{g}_1^*(R_1)$) using the one-sided nearest neighbor estimator outlined in Xu (2013).
5. Use this estimate of $g_1^*(R_1)$ to evaluate Equation 3.9 to obtain $\hat{v}^* = \xi_1(R_1)$.
6. Guess candidate parameters θ_1 and θ_ω, and calculate the likelihood of the pseudosample $\{\vec{\xi_1}, \vec{\xi_2}\}$.
7. Search for parameters that minimize the chosen objective function.

As shown in Aryal, Gabrielli, and Vuong (2014), such an approach circumvents the curse of dimensionality with respect to covariates and is similar to parametric assumptions that are often informally made in applying the method proposed in Guerre, Perrigne, and Vuong (2000).

3.5 Monte Carlo Simulation

In this section, I test the performance of the estimator outlined above and compare it to naive estimates that ignore the strategic implications of re-auctioning. A few practical issues need to be addressed. First, there is the matter of bandwidth selection. I use the "rule-of-thumb" bandwidth from Silverman (1986), adjusted for the triweight kernel. Guerre, Perrigne, and Vuong (2000) note that the kernel density estimator used to recover $g^*(\cdot)$ will not perform well near the boundaries of the bid data. In practice, I use the boundary-correction approach of Hickman and Hubbard (forthcoming), and estimate the density at the first-round reserve price, $g^*(R_1)$, using the one-sided nearest neighbor approach in Xu (2013).

I parameterize the model such that the first-round values come from a truncated lognormal distribution with location parameter μ_1 and scale parameter σ_1. The shock to values that occurs between rounds is log-additive, and comes from a truncated lognormal distribution with location parameter μ_ω and scale parameter σ_ω.[9] I also choose the first and second-round reserve prices, R_1 and R_2; the number of potential bidders, N; the number of auctions in the simulated dataset, A. These parameter values are listed in Table 3.1.

The results of the 1000 simulations are presented in Table 3.2. In the first panel, I present the results from the joint estimation process. The distribution of first-period values are very well estimated. The bias is nearly zero, and the standard deviation across replications is quite low. The estimate of the point of truncation v^* via $\xi_1(R_1)$ is pretty close on average, though biased slightly upward. The estimates of the innovation parameters are less accurate. In particular, the median (location) of the distribution is underestimated, while the scale is overestimated. Figure 3.2 compares the density associated with the mean parameter estimates compared with the density from the true data-generating process.

Although this seems a bit off, it is informative to examine the accuracy of the implied second-round distribution, F_2. I calculate and plot the estimated and true densities of $v_2 = v_1\omega$. In Figure 3.3, the densities are based on an untruncated F_1 (i.e., $v^* = 0$). In Figure 3.4, the truncation point is the true value of v^* for the true distribution and the estimated value of v^* for the estimated value. These figures illustrate that the second-round distribution is still being captured fairly well by the estimates of the first-round values and value innovations.

[9] For technical reasons, these distributions are truncated above at $\bar{v} = 200$, a value that is extremely unlikely to arise given the other parameters.

The results also highlight the danger of entirely ignoring the possibility of re-auctioning. I estimate the first-period valuation distribution assuming the "naive" one-period model and list the results for comparison. I find that these estimates are well off the mark; the severity of the problem is clear from the plotted densities in Figure 3.1. The naive model misspecifies the lower bound of the recovered valuations at $R_1 = 12$ instead of near $v^* = 14.15$. Thus, the naive estimator recovers a wider distribution with a smaller median to explain the observed bids.

3.6 Application: Auctions for DNR Timber Contracts

In this section, I present a description of institutions and data related to Michigan DNR timber auctions. In particular, I focus on the extent to which these auction are an appropriate setting to apply my estimation procedure.

3.6.1 DNR Timber Auctions and the Reserve Price Policy

The Michigan DNR is tasked with managing the state forest system with the threefold aim of maintaining ecological integrity, providing recreational opportunities, and supporting the local timber industry and earning revenue through timber sales. The DNR holds roughly 500 auctions each year at various field offices scattered throughout the Northern Lower Peninsula and the Upper Peninsula.

When a stand of trees is ready to be auctioned, the contract terms are made public and 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.[10] The auctions are sealed-bid first-price auctions with public reserve prices. The bids, bidder identities, and number of bids submitted are held confidential until the results are made fully public at the bid opening. The highest bidder wins the contract and is obligated to harvest the

[10]Conversations with DNR officials and estimates in Herrnstadt (2015) suggest that bid preparation costs are quite small in this context.

specified timber before a contract deadline. Failure to fulfill the contract terms results in a financial penalty and possible exclusion from future sales.

There is a reserve price based on the winning bid on nearby timber of the same tree species. The DNR forester then uses a rubric to set the reserve price based on factors related to the cost of cutting and transporting the logs to market, as well as a "market conditions" adjustment to allow for room to bid (Heym, 2013). In the case that the sale "goes no-bid" (i.e., it does not receive any bids), the DNR follows a well-known procedure. First, the sale is made available at the reserve price for about 60 days.

If there are no interested parties (or multiple interested parties), then the sale will be re-auctioned at a reserve price that is roughly 20-40 percent lower than the original. Figure 3.5 shows the distribution of second-round auction reserve prices as a share of the corresponding first-round reserve price. Clearly there is quite a bit of variation, but the 10th and 90th percentiles are 55% and 80%, respectively, so the data are fairly consistent with the 20-40% discount cited in conversations with the DNR. This second-round auction typically takes place 3-9 months after the initial auction, as shown in Figure 3.6.

In Table 3.3, I provide some evidence that suggests firms are bidding differently in second-round auctions than first-round auctions. For first- and second-round auctions, I regress the log of the winning bid on various sets of controls. In Column 1, I control for a rich set of sale characteristics. The R^2 is 0.915, so these are explaining most of the variation in winning bids. Second-round auctions receive bids that are roughly 27 log points lower. This could be driven by three phenomena. First, the reserve prices are lower in second-price auctions, which we know from Figure 3.5. This structural factor will affect bidding behavior. Second, in a second-round auction, bidders know that their competitors are draws from a weaker distribution; this is analogous to $F_2(v_2|v_1 < v^*)$ in the model. This will reduce the aggressiveness with which loggers bid. Third, if there is unobserved auction heterogeneity, the sales that make it to the second round could be adversely selected on this basis. Thus, we would expect a sale with a vector of observables that receives a bid in the first round to garner higher bids than one that is auctioned in the second round.

I can address the first factor directly by including the reserve price in Column 3. To the extent that the DNR incorporate the unobserved auction heterogeneity into their reserve-price setting procedure, this will also help address the third factor. In this specification, the second-round effect is reduced by about 50 percent, but is still 14 log points. When I control for the first-round reserve price in Column 5, this helps control for the truncation process driven by the no-bid first-round auctions. When the first-round reserve price is higher independent of the second-round reserve price, the bids in the second round are slightly higher. This may suggest that bidders recognize that a higher first-round reserve price will result in stronger second-round competition on average. These regressions are largely consistent with the predictions of the theoretical model; structural estimation is necessary to parse out the true impacts of the commitment problem.

3.6.2 Is the Model Appropriate?

The preceding discussion is a simple model of re-auctioning. However, the empirical context is a bit more complicated in two ways. First, the model simplifies by assuming the seller is willing to hold the contract off the market after the second round. This is empirically untrue: Table 3.4 shows the outcomes of each round of auctions in the data. Of the 283 second-round auctions, 199 receive bids, but 52 reach a third round, 8 reach a fourth round, and a single auction reaches a fifth round. Further, on 47 occasions, a no-bid contract is never actually re-auctioned, even if no logger buys them at the reserve price. Although the fact that 19 of these come from the last 3 quarters of my data indicates that truncation may be driving much of this, the model may still need to model the time gap as stochastic. More simply, I could just estimate the probability of not being re-auctioned and incorporate that directly into the loggers' objective function.

Second, the DNR allows firms to purchase the contract at the reserve price for 60 days. If zero or more than one bidder shows interest, then the sale is re-auctioned. In the data, a contract is issued at the reserve price after the fact in roughly one-third of all no-bid auctions.

The model I have outlined predicts that we would never observe firms buying the sale at the reserve price between auctions. If their valuation was above v^*, they would have bid. If their valuation was between R_1 and v^*, they would not have bid above the reserve price in the auction. In fact, the calculus does not change after the auction has resolved itself. When such a bidder has the opportunity to buy the contract at the reserve price, they will compare $(v - R_1)$ to $\delta E[(v_2 - b_2(v_2))F_2(v_2|v_1 < v^*)^{N-1}]$. In fact, this is exactly the indifference condition for a bidder with a valuation of v^*. Having shown that the benefits of waiting are increasing in v_1, it is clear that any loggers involved a no-bid auction would prefer to wait rather than buy the sale at R_1. Even for a logger with a valuation of v^*, the current gains from buying at R_1 still do not outweigh the expected surplus from waiting.

However, suppose that new information arises as a result of concurrent auctions. In particular, suppose that bidders who buy up the no-bid contracts have placed bids in other concurrent auctions. If the marginal value of the no-bid auction is a function of the other auction outcomes because the sales are complements or substitutes, then this could certainly lead to a bidder snapping up the no-bid auction. For instance, suppose a firm only wants either sale 1 or sale 2, but expects to make a bigger surplus on sale 1. Then they lose the auction for sale 1. Sale 2 now looks like a good deal at the reserve price.

In this case, modeling the process of the posted-price sale becomes very important. In practice, the DNR assesses whether there is more than one firm interested in the sale in determining whether to allow a posted-price sale or to re-auction the sale. When there are zero or one interested parties, it is clear how to proceed. When there are multiple interested parties, three possible modeling approaches are to assume that (1) another auction is automatically held, (2) the contract is randomly assigned among all willing to buy it (first-come, first-served), or (3) the contract is assigned to the interested party with the highest value (efficient rationing). The first option is closest to the policy officially outlined, but conversations with DNR officials will hopefully reveal how this situation is handled in practice.

Incorporating interdependent values seems like it could be a promising approach: perhaps learning about other bidders signals leads to updating of one's expected value and results in a purchase above the reserve price. However, this alone cannot rationalize the purchases at R_1, for reasons similar to those just outlined. The equilibrium in the first-price auction with a reserve price and interdependent values involves the following cutoff condition: $b_1(x^*) = R_1$, where $x^* = \inf\{x : E[v_1|X = x, Y_1 < x] \geq R_1\}$, x is the realization of the bidder's private signal X, and Y_1 is the highest opponent's signal. That is, the expected value for a bidder with the cutoff signal x^*, conditional on having the highest signal, is exactly R_1. So even after the auction has resolved itself, the new information is the information the marginal type was already conditioning on to begin with; this outcome already led them to an expected value of exactly R_1. Given that all other bidders refused to bid above the reserve price, x^* remains exactly the cutoff type for which buying the sale at R_1 makes sense economically. Still, this dynamic could be important for accurately recovering valuations from second-round auctions: when the auction receives no bids, loggers would use this information to update their own expected valuation, as well as the distribution of their competitors' signals.

3.7 Incorporating Unobserved Auction Heterogeneity

If I can assume that sales that go no-bid are the same as those that do receive a bid and are observably similar, then there is no problem identifying the valuations. However, if there is unobserved heterogeneity that persists from round one to round two, things are a more difficult. In this section, I outline an approach that would control for unobserved differences across auctions. Descriptive results from the DNR data suggest that this approach would be valid in this context.

It would likely be a poor assumption that the distribution of unobservable auction characteristics are the same ex ante for auctions that receive bids and those that do not. Thus, differences between first-round bids and second-round bids are jointly driven by (1) the truncation of the first-round valuation distribution, (2) the lower reserve price, and (3) different realizations of unobserved heterogeneity.

This would be problematic for the desired counterfactuals. For example, consider observably similar auctions, one that went no-bid (auction 1) and most of which that received bids (auctions

{2...A}). Without unobserved heterogeneity, the only difference between auctions 1 and {2...A} is that bidders in auction 1 received idiosyncratically low values. The optimal reserve policy would call for the same reserve price for these observably identical auctions. However, if the no-bid was driven by unobserved auction heterogeneity, the *distribution* of values is actually lower in auction 1. Thus, the response to a change in the reserve price could be wildly different.

Solving this problem is uniquely complicated in this setting. Since there are no bids observed in the no-bid auction, I cannot apply the now-standard deconvolution approach put forth by Krasnokutskaya (2009). First, it requires multiple bids per auction, but I observe zero bids from the first-round. Second, it identifies the distribution of heterogeneity, but does not allow the econometrician to associate a given observation with a particular realization of the unobservable. Thus, even if I recovered a distribution of unobserved auction heterogeneity separately for the two rounds, I would be observing an advantageously selected distribution for the first round, and an adversely selected distribution for the second round.

Despite the lack of first-round bids in these no-bid auctions, I still observe the reserve price and the auction characteristics. The approach in Roberts (2013) leverages an observable seller choice variable that is determined by unobserved heterogeneity. In his case, it is the auction reserve price. Under some assumptions, the residual from a regression of the reserve price on observable auction characteristics can identify the quantile of the unobserved heterogeneity for a given observation. Assuming that unobserved heterogeneity is perfectly persistent within the same sale across rounds, it can be identified for all auctions using first-round reserve prices.

One assumption that is necessary for this approach to work is that there cannot be any unobserved factors entering into the seller's reserve price-setting equation that do not enter into the bidding equation. One way to test for this is to look at whether, controlling for other observables, the reserve price affects the probability an auction receives no bids. Columns 1-3 of Table 3.5 are logit regressions of a dummy for whether a sale went no-bid or not. Unfortunately in Column 2, the reserve price is significantly related to a higher likelihood of going no-bid.

However, this can be explained by the manner in which the DNR sets their reserve prices. The final reserve price is a benchmark appraisal based on historical winning bids, multiplied by an index of “appraisal factors”. If prices are on a downward trajectory, then the high historical bids would lead to reserve prices that are too high in the present. Thus, higher reserve prices could be mechanically linked to auctions that go no-bid. In Column 3, I separate out the two components and find that the result is indeed driven by the benchmark price. The choice variable of the DNR is actually the appraisal factors, which is not correlated with a higher likelihood of going no-bid.

Finally, the approach hinges on the idea that the DNR and loggers are responding similarly to the unobserved heterogeneity. As seen by comparing Column 4 of Table 3.5 with Column 1 of Table 3.3, bidders and the DNR respond similarly to observable characteristics in setting their appraisal factors and bids, respectively. With the exception of the percentage bid species variables, the significant effects have the same signs in the appraisal factors regression as they did in the winning bid regressions. The fact that both sides of the market respond similarly to observable differences in sales lends some credence to the hypothesis that they respond similarly to *unobserved* differences as well.

3.8 Conclusion and Future Work

This paper outlines the importance of accounting for seller commitment in empirically analyzing auction data. First, I develop a theoretical model that gives the seller an option to re-auction the object once, and incorporates intertemporal shocks to bidder valuations. Second, I discuss identification of key distributions and structural objects and find that a semiparametric approach is be necessary to fully identify the model with only bid data. Third, I propose a semiparametric estimator, demonstrate that it performs well in a Monte Carlo simulation, and discuss a strategy for dealing with unobserved auction heterogeneity. Fourth, I explore Michigan timber auctions as a possible empirical setting.

Future work will apply the estimator to the empirical DNR setting. This will require some modifications to the model. Although simply extending the model to allow for two periods of bidding is an improvement, it ignores the fact that contracts are often re-auctioned multiple times. Further, I will need to formally model the period during which the sale is available at the reserve price, as this is empirically important.

Further, while the semiparametric approach is appealing from a practical standpoint, the method would be more credible if the nonparametric identification of F_ω was established. One possible avenue that I have not fully explored is the use of an instrumental variable. The essential identification problem is that I do not observe first-round bidders with valuations below v^*. However, if there is a variable that exogenously shifts v^*, this will reveal a new portion of the distribution. Assuming that the ω is independent of the instrument and the first-round valuation v_1, comparing the second-round distributions should allow identification of the distribution of ω. One particularly promising candidate is variation in the reserve price that is mechanically driven by bids in past auctions. As noted in Section 3.7, the benchmark price seems to have a large impact on the likelihood of a sale receiving zero bids, conditional on a rich set of other covariates. This indicates that past shocks to winning bids, either through market conditions or idiosyncratic auction outcomes, are indeed shifting v^*. Two other candidates that have been used in past papers for a similar purpose are the number of potential bidders or the distance from potential bidders to the sale.

3.9 Figures

Figure 3.1: Performance of Naive Estimator of F_1

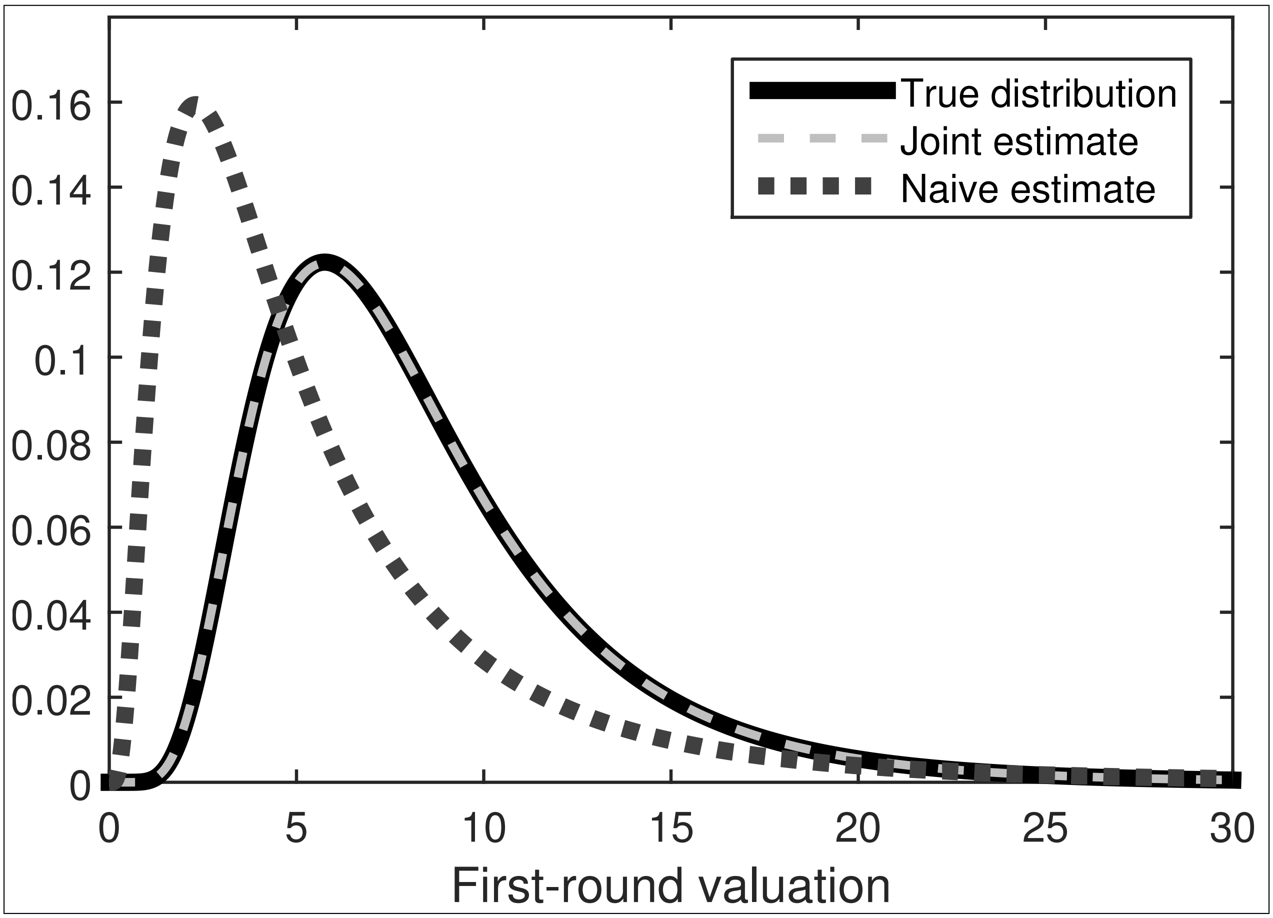

Figure 3.2: Monte Carlo Estimates of F_ω

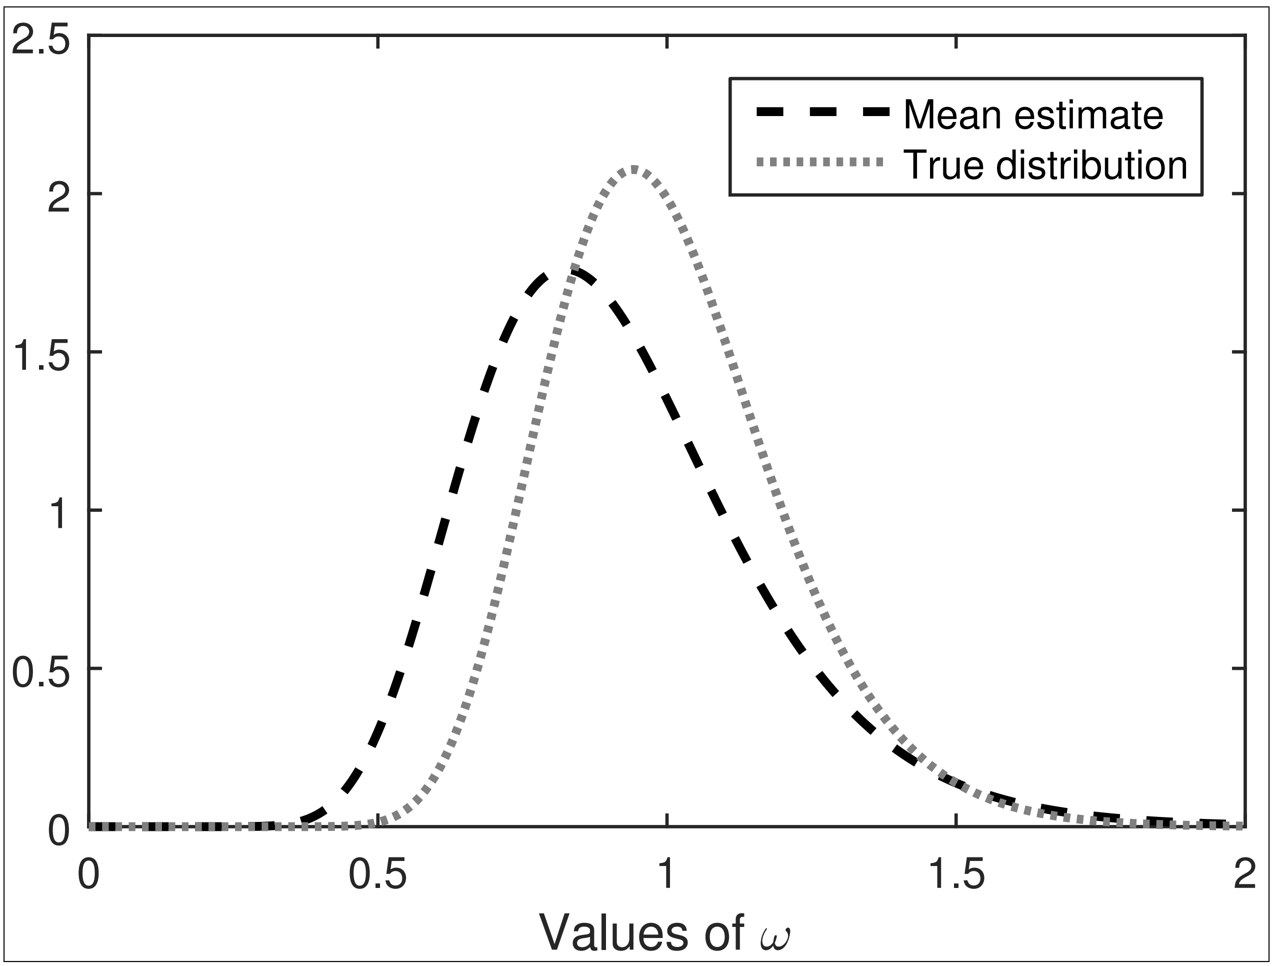

Figure 3.3: Monte Carlo Estimates of F_2 (Untruncated)

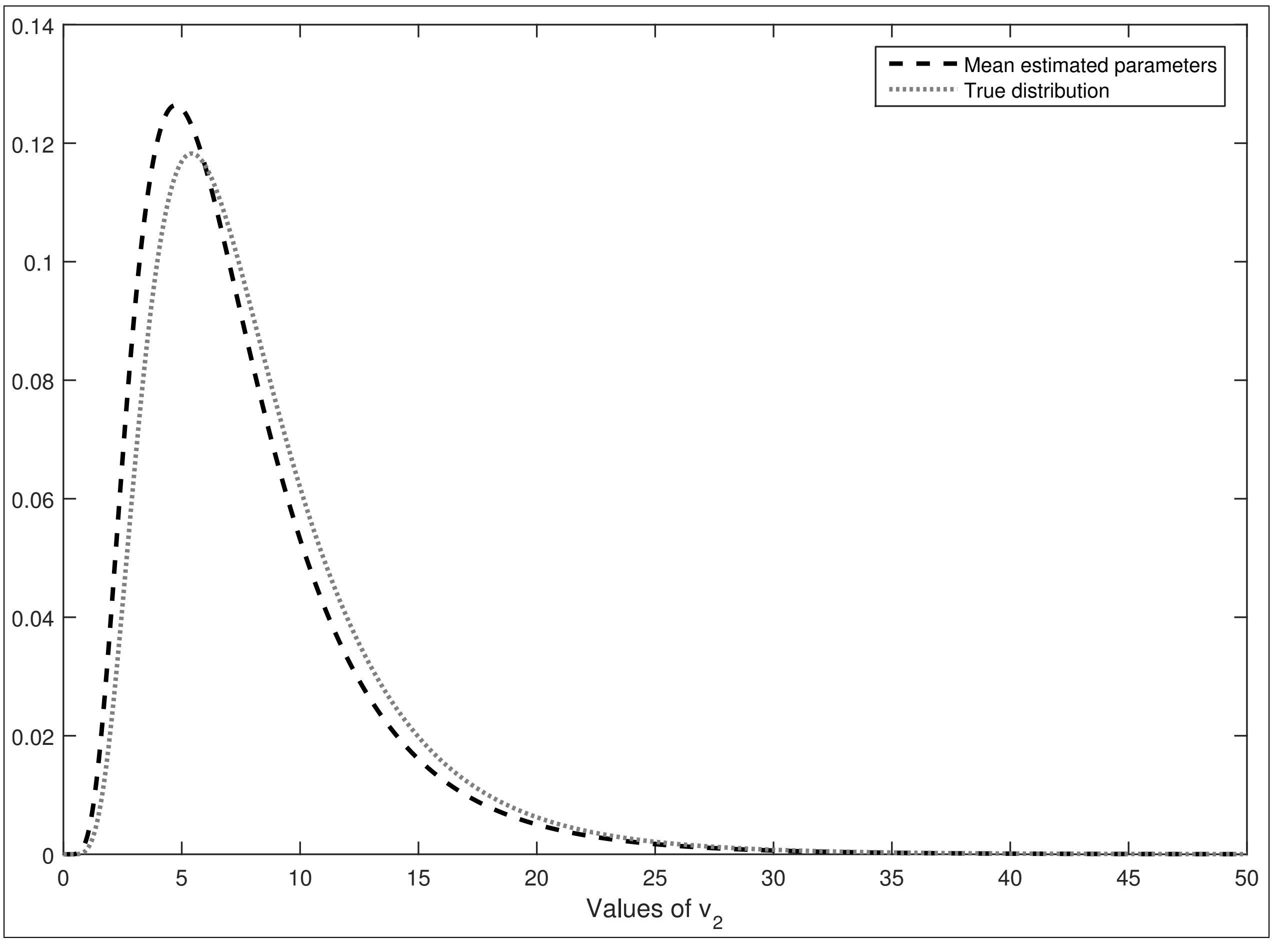

Figure 3.4: Monte Carlo Estimates of F_2 (Truncated at true or mean estimated v^*)

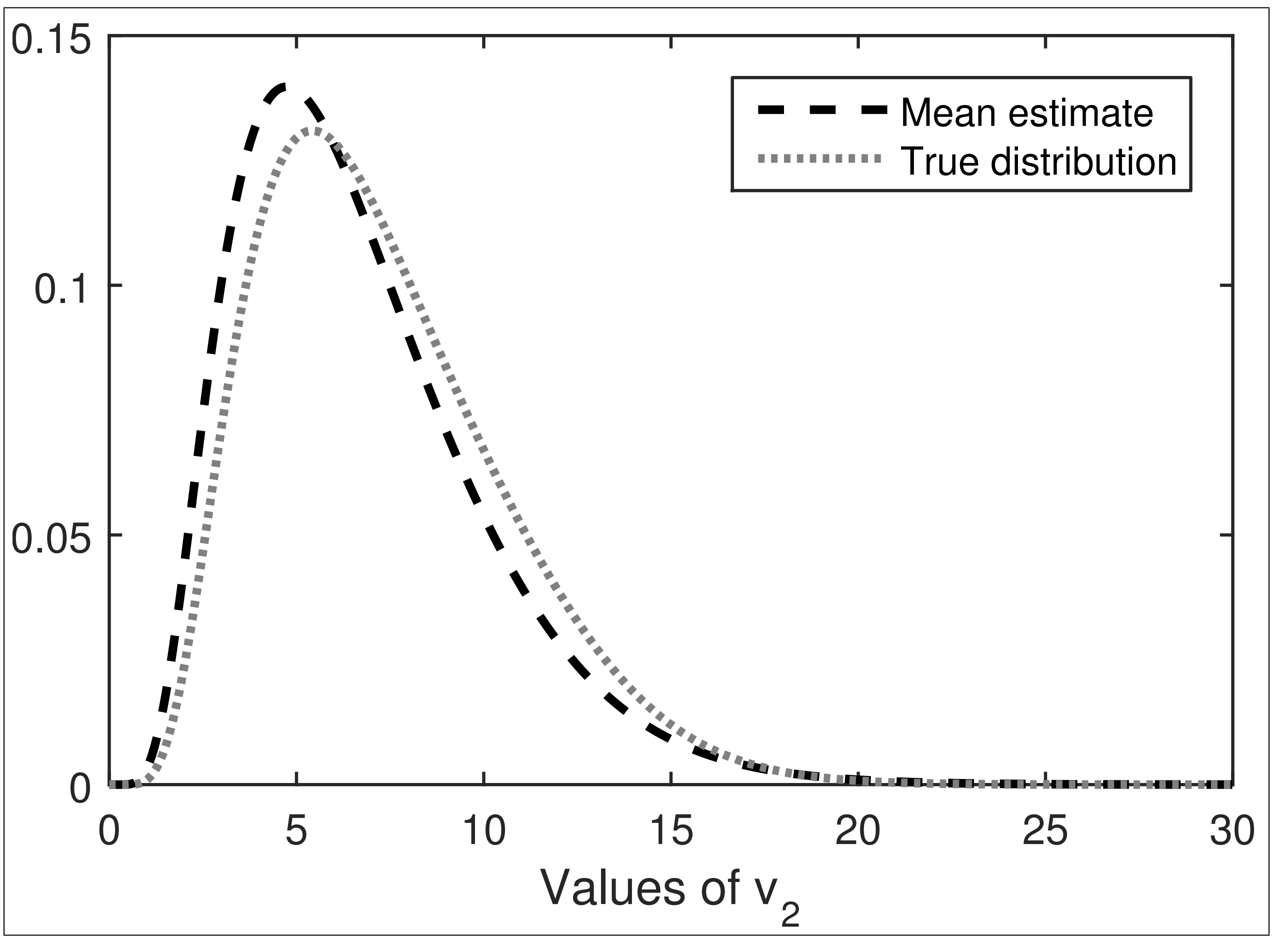

Figure 3.5: DNR reserve prices in second-round auctions relative to first round

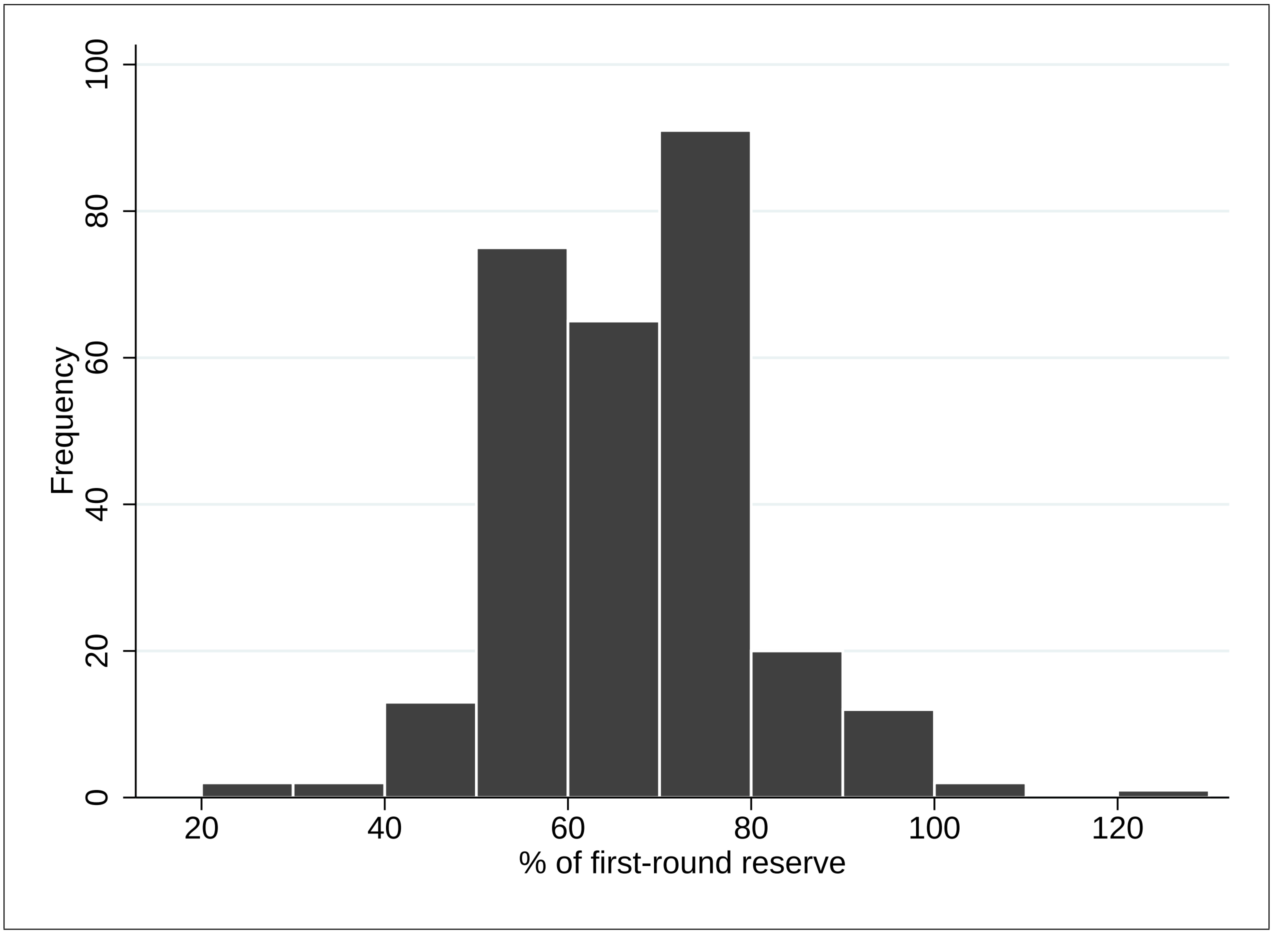

Figure 3.6: Time gap between first and second round

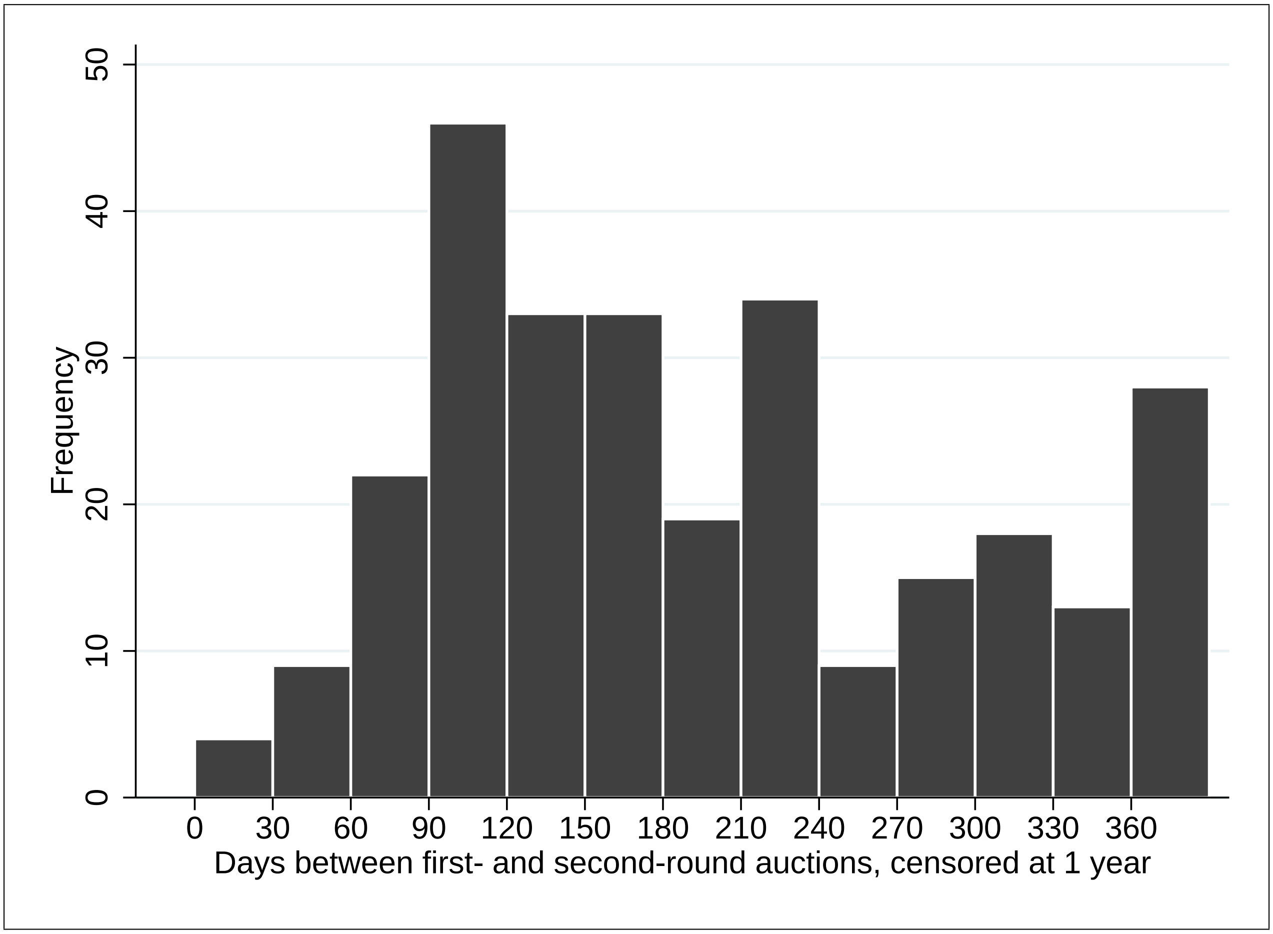

3.10 Tables

Table 3.1: Parameterization of Monte Carlo

Parameter	μ_1	σ_1	μ_ω	σ_ω	R_1	R_2	N	A
Value	2	0.5	-0.02	0.2	12	9	10	500

Table 3.2: Monte Carlo Results

Parameter	μ_1	σ_1	υ^*	μ_ω	σ_ω
Truth	2	0.5	14.15	-0.02	0.2
Joint Estimation					
Mean	2.001	0.501	14.210	-0.123	0.265
SD	0.095	0.050	0.474	0.081	0.028
RMSE	0.095	0.050	0.478	0.131	0.071
Naive Estimation					
Mean	1.464	0.793	12	-	-
SD	0.057	0.030	-	-	-
RMSE	0.539	0.294	-	-	-

Table 3.3: DNR Regressions

VARIABLES	(1)	(2)	(3)	(4)
Second round dummy	-0.273***	-0.216***	-0.137***	-0.478***
	(0.025)	(0.013)	(0.014)	(0.130)
ln(Reserve price)		0.983***	0.754***	0.752***
		(0.004)	(0.013)	(0.013)
ln(First-round reserve price)				0.034***
				(0.012)
ln(Total timber volume)	1.071***		0.279***	0.279***
	(0.011)		(0.015)	(0.015)
Share Softwood: sawlogs	1.978***		0.225***	0.223***
	(0.067)		(0.054)	(0.054)
Share Hardwood: sawlogs	1.358***		0.104***	0.106***
	(0.038)		(0.032)	(0.032)
Share Hardwood: pulpwood	0.108***		-0.082***	-0.083***
	(0.024)		(0.017)	(0.017)
ln(Acres)	-0.019		0.022***	0.022***
	(0.013)		(0.008)	(0.008)
Species-product HHI	0.751***		0.249***	0.248***
	(0.029)		(0.021)	(0.021)
Percent bid species	0.622***		0.227***	0.229***
	(0.060)		(0.044)	(0.044)
Constant	2.772***	0.588***	0.853***	0.867***
	(0.070)	(0.040)	(0.062)	(0.062)
Observations	5,353	5,353	5,353	5,353
R-squared	0.917	0.934	0.959	0.959
Major species dummies	Y	N	Y	Y
Quarter-by-year dummies	Y	N	Y	Y
Management Unit dummies	Y	N	Y	Y

Notes: *** p<0.01, ** p<0.05, * p<0.1. Robust standard errors are in parentheses. Dependent variable is the log of the winning bid. Sample includes only first- and second-round auctions.

Table 3.4: Round-by-Round Contracting Outcomes

	Round					
Outcome	1	2	3	4	5	Total
Bid Received	5136	199	34	7	1	5377
Purchased at Reserve	161	25	6	0	0	192
Not Purchased	319	59	12	1	0	391
Total	5616	283	52	8	1	5960

Table 3.5: Exploring unobserved heterogeneity assumptions

VARIABLES	(1) No-bid Logit	(2) No-bid Logit	(3) No-bid Logit	(4) Determinants of Factors
ln(Reserve price)		0.420**		
		(0.198)		
ln(Appraisal Factors)			-0.003	
			(0.350)	
ln(Benchmark Valuation)			0.619***	-0.032***
			(0.229)	(0.010)
ln(Total timber volume)	-0.633***	-1.094***	-1.266***	0.091***
	(0.118)	(0.245)	(0.263)	(0.012)
Share Softwood: sawlogs	0.831	-0.140	-0.495	0.243***
	(0.649)	(0.765)	(0.805)	(0.038)
Share Hardwood: sawlogs	1.237**	0.598	0.372	0.181***
	(0.481)	(0.574)	(0.587)	(0.025)
Share Hardwood: pulpwood	0.927***	0.855***	0.780***	-0.041***
	(0.272)	(0.280)	(0.288)	(0.013)
ln(Acres)	-0.025	0.010	-0.017	-0.062***
	(0.127)	(0.128)	(0.130)	(0.007)
Species-product HHI	-0.322	-0.606	-0.699*	0.107***
	(0.364)	(0.382)	(0.388)	(0.016)
Percent bid species	-1.211*	-1.411**	-1.499**	0.014
	(0.707)	(0.717)	(0.718)	(0.032)
Constant	0.077	-0.968	-1.729	-0.279***
	(1.056)	(1.163)	(1.262)	(0.050)
Observations	5,491	5,491	5,491	5,615
R-squared				0.315
Major species dummies	Y	Y	Y	Y
Quarter-by-year dummies	Y	Y	Y	Y
Management Unit dummies	Y	Y	Y	Y

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Robust standard errors are in parentheses. In Columns 1-3, the dependent variable is whether the sale receives a bid or not. In Column 4, the dependent variable is the logarithm of the appraisal factors. Sample includes only first- and second-round auctions.

APPENDIX A.

Appendices for Chapter 1

A.1 Examining seasonality in restrictions

Certain seasons are more likely than others to be restricted. Figure A.1 shows the average months of restrictions in each season. The average sale has nearly one month of restrictions in the spring and summer, while it is restricted for only one week in fall and winter each. Conditional on having any restrictions, the average number of months restricted is about 1.5 months in spring and summer each, and 0.4 months in fall and winter each. Figure A.2 shows that there is also substantial variation in the number of months restricted in each season as the total annual number increases.

One potential concern is that there is a particular season that loggers prefer to have available (i.e., a "good" season). If that good season is only restricted when most other months are restricted, this could drive the sort of results that I produce. To examine this possibility, I estimate the following specification:

$$\begin{aligned} ln(winning\ bid_a) = & \sum_{s=1}^{4} \alpha^s MonthsRestricted_a^s \quad + \\ & \sum_{s=1}^{4} \pi^s MonthsRestricted_a^s * MonthsRestricted_a^{-s} + \beta X_a + \varepsilon_a \end{aligned}$$

where $MonthsRestricted_a^s$ is the number of months during season s that auction a is restricted, and $MonthsRestricted_a^{-s}$ is the number of months in seasons *other* than s that auction a is restricted. The results are presented in Table A.1. If restricting a particular season has an impact on bids even when there are essentially no other restrictions, this would suggest that the reduced-form effect is driven by a good season. Figure A.3 plots the marginal effects by month, as a function of the number of months restricted in other seasons. In all seasons, the marginal effect of an additionally month is very close to zero when other seasons are lightly restricted. This suggests that there is no good season driving the effect on bids.

Of the four seasons, only fall has a marginal effect that is statistically significantly difference from zero. Most of the variation in restrictions beyond 6 months is a tradeoff between fall and winter months. Thus, I cannot rule out that winter is a particularly costly season to cut. Unfortunately, the data are not rich enough to determine whether the effect of heavy restrictions arises because loggers particularly dislike having to harvest during winter, or because of the loss of flexibility emphasized by the loggers and DNR foresters. In the exposition of the paper, I emphasize the flexibility explanation. Although the policy design implications of the two mechanisms would be different, my analyses of effects on bids, costs, and surplus still reflect the impact of seasonally restricting a contract for most of the year.

A.2 IV estimation of reduced-form results

The identification of the effects implied by the reduced-form regressions is discussed in the main body of the text. However, in this section, I attempt to perform an instrumental variables analysis as a robustness check to see if some omitted variables may be biasing the reduced-form regressions that establish the equilibrium effects.

Conversations with the DNR revealed that the assignment of appraising foresters to different sales within a management unit is done without regard to individual skills or preferences; i.e., conditional on the management unit, forester assignment should be quasi-random. Following the discussion in Maestas, Mullen, and Strand (2013), I generate an instrumental variable based on the mean number of months of restrictions that a given forester places on their sales observed in the sample. For each auction, the instrumental variable is the mean of the potentially endogenous "Months Restricted" variable for all of that forester's observations except the observation of interest. This "leave-one-out" approach is necessary to prevent the observation from influencing its own constructed instrument. I only include foresters that appraise at least 5 sales in my data, leaving 124 foresters over 5098 sales (4647 of which receive bids). Figure A.4 shows how many sales my sample foresters appraised.

Figure A.5 shows the distribution of the instrument. There is substantial variation; however, it is primarily in the 0-4 month range. In the OLS specification, the results primarily arise within the 5-10 month range. This is likely to attenuate any IV results, but the analysis could still inform as to whether the estimates among less-regulated are obviously tainted by endogeneity. As the randomization takes place at the management unit level, we want to actually look at how much variation there is within management unit. That is, much of the variation in Figure A.5 could be due to cross-unit differences. I demean the instrument within-management unit (Figure A.6) and scale by the within-management unit standard deviation (Figure A.7)

The results are shown in Table A.2 with standard errors clustered on the forester. The instrument is quite strong, with first stage F-statistics of over 100. The point estimate of the coefficient of interest is similar to the OLS estimate, but is considerably less precise. The lack of precision makes it difficult to draw rigorous conclusions. However, given that the instrument lies mostly in the 0-4 month range, this is not particularly surprising. The nonlinear OLS effects presented before suggest that most of the action is in the right tail; thus, it is unsurprising that this instrument attenuates the linear effect. Still, the treatment effect implied by the TSLS point estimate of 4 months of restrictions (-1.6 percent) roughly corresponds to the treatment effect along that portion of the distribution in the main specification.

A.3 Derivation of bid likelihood

Given the bid function the distribution of values, and the entry threshold, the bid likelihood follows easily. Let the inverse bid function be denoted $b^{-1}(b)$. Then the CDF of observed bids will simply be equal to the CDF of the value associated with the bids through the inverse bid function. Thus, the bid density can be derived by the standard transformation using the Jacobian:

$$
\begin{aligned}
G(b) &= F(b^{-1}(b)) \\
g(b) &= f(b^{-1}(b))b^{-1\prime}(b) \\
g(b) &= \frac{f(b^{-1}(b))}{b'(v)}
\end{aligned}
$$

Given the closed-form solution for the equilibrium bid function, this density can be expressed as:

$$
g(b) = \frac{F(b^{-1}(b))^N}{(N-1)\left[\int_{v^*}^{b^{-1}(b)} F(u)^{N-1}du + F(v^*)^{N-1}\right]}
$$

The likelihood of observing a given bid conditional on a vector of parameters is:

$$
\ell_a(b_{ia}|\theta) = \begin{cases} F(v^*), & \text{bidder } i \text{ did not participate} \\ g(b), & R \leq b \leq b(\bar{v}) \\ 0, & b > b(\bar{v}) \end{cases}
$$

The likelihood of observing a given vector of bids (or non-participation) $\mathbf{b_a}$ is the within-auction product of these bid-specific likelihoods. To calculate the likelihood of a given observation, I average the likelihoods across simulation draws.

Although the upper bound of the bid support, $b(\bar{v})$ varies depending on the specific simulation draw of θ, it does not actually vary with guesses of the parameters Γ. Because Γ essentially determines the weights of a mixture distribution over θ, the support is actually independent of Γ in the limit: bids anywhere in the interval $[R, \bar{v})$ could always arise given enough simulation draws. In practice, I take 1000 draws per data point, which proves sufficient to ensure a non-zero likelihood for every observation. Further, Monte Carlo simulations based on this importance sampling MSL procedure perform quite well.

A.4 Figures and Tables

Figure A.1: Average months of restrictions, by season

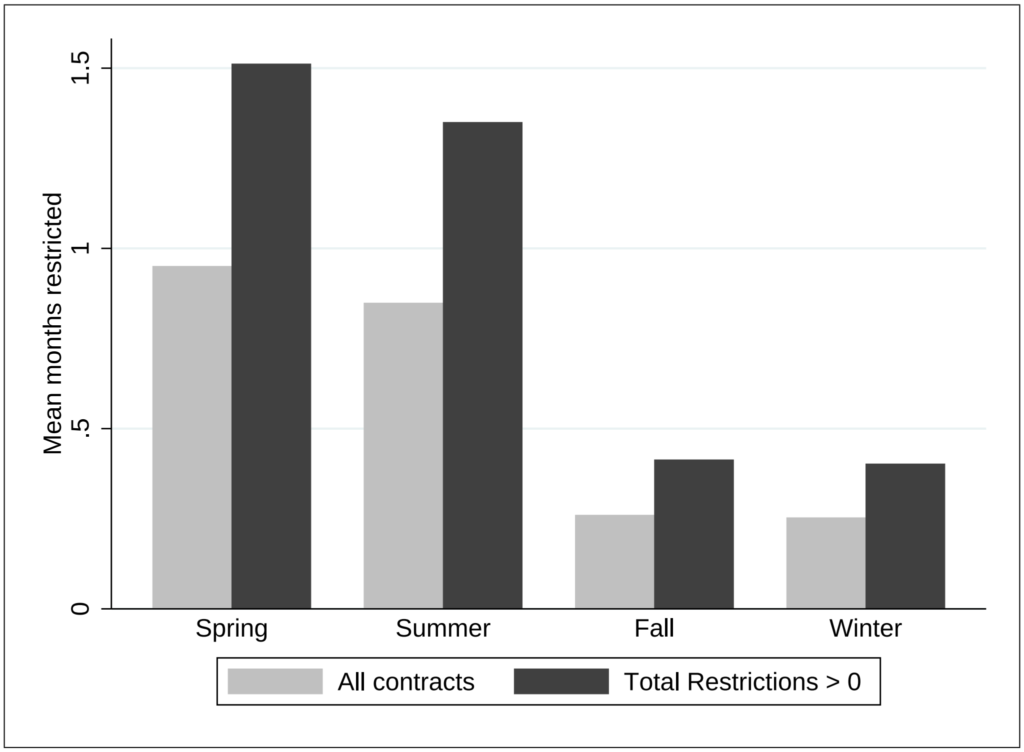

Figure A.2: Restrictions in various seasons, by total months restricted

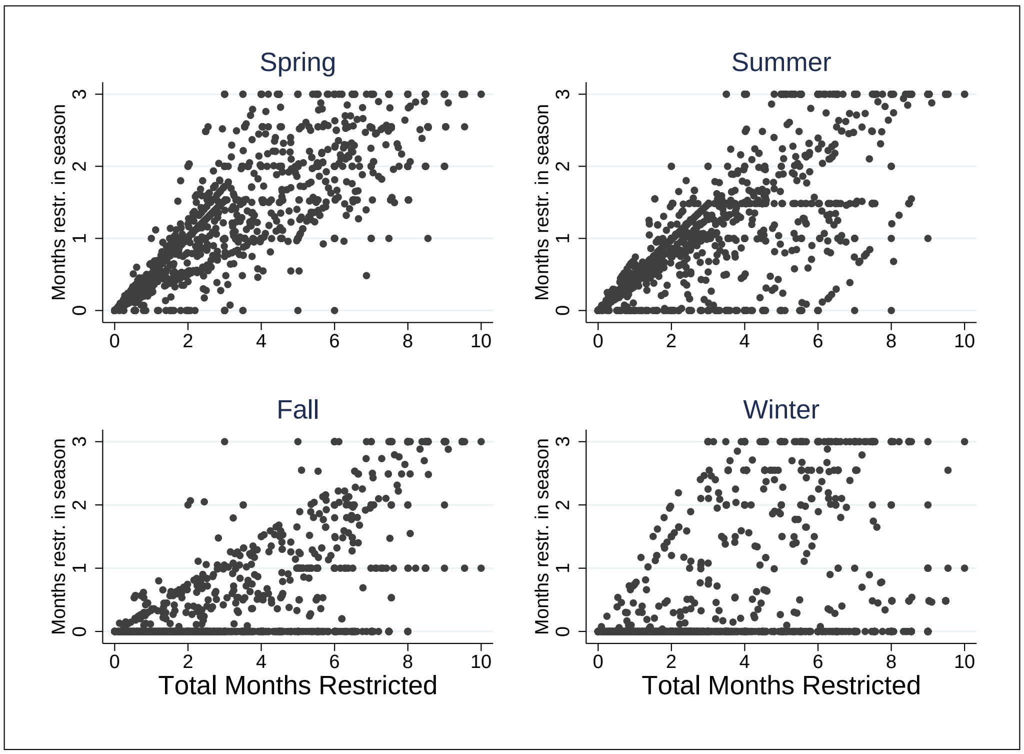

Figure A.3: Marginal effect of an additional month in a given season

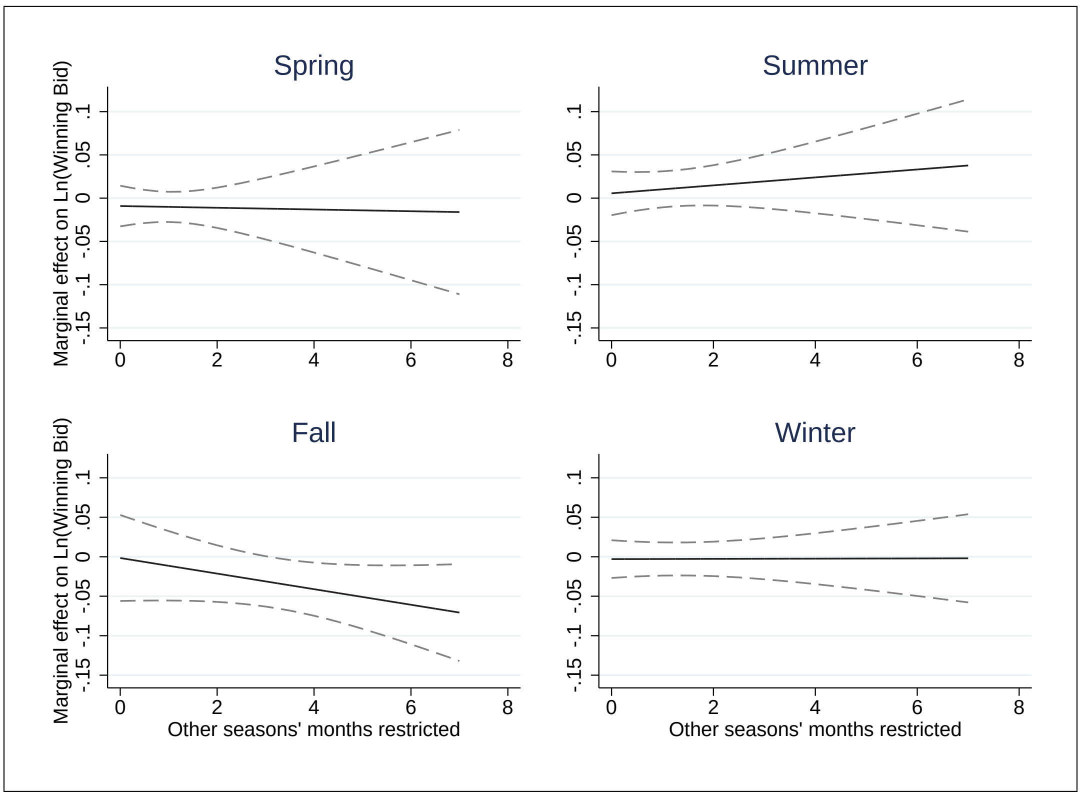

Figure A.4: Number of sales appraised by each forester

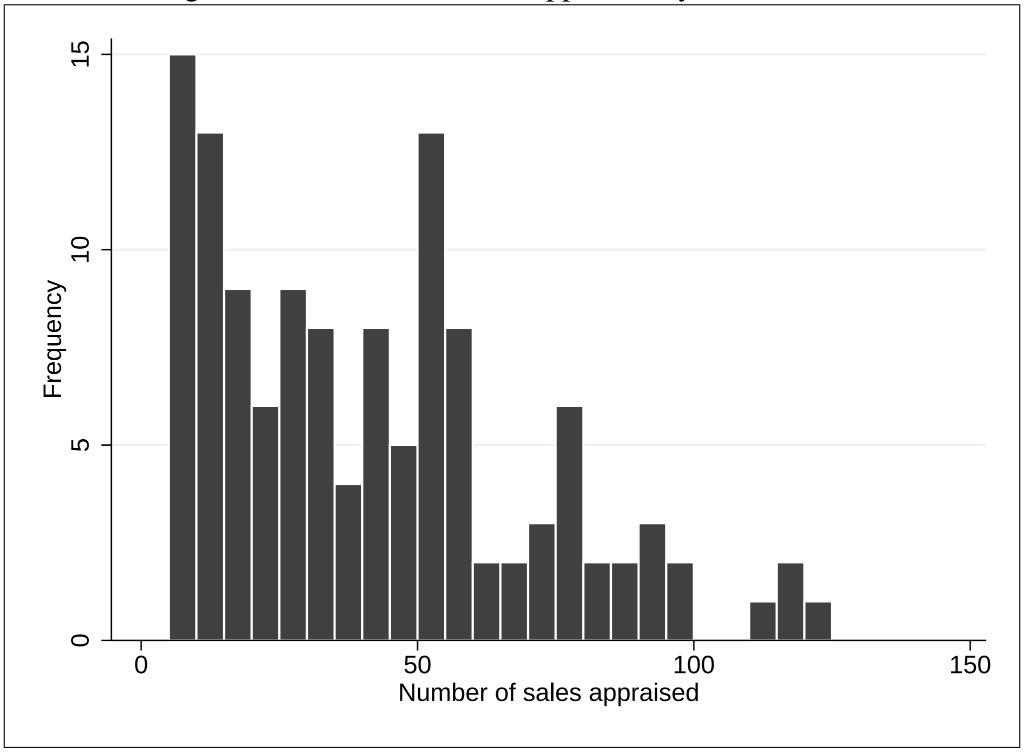

Note: Each observation is a forester. Excludes foresters with fewer than 5 observations.

Figure A.5: Distribution of instrumental variable

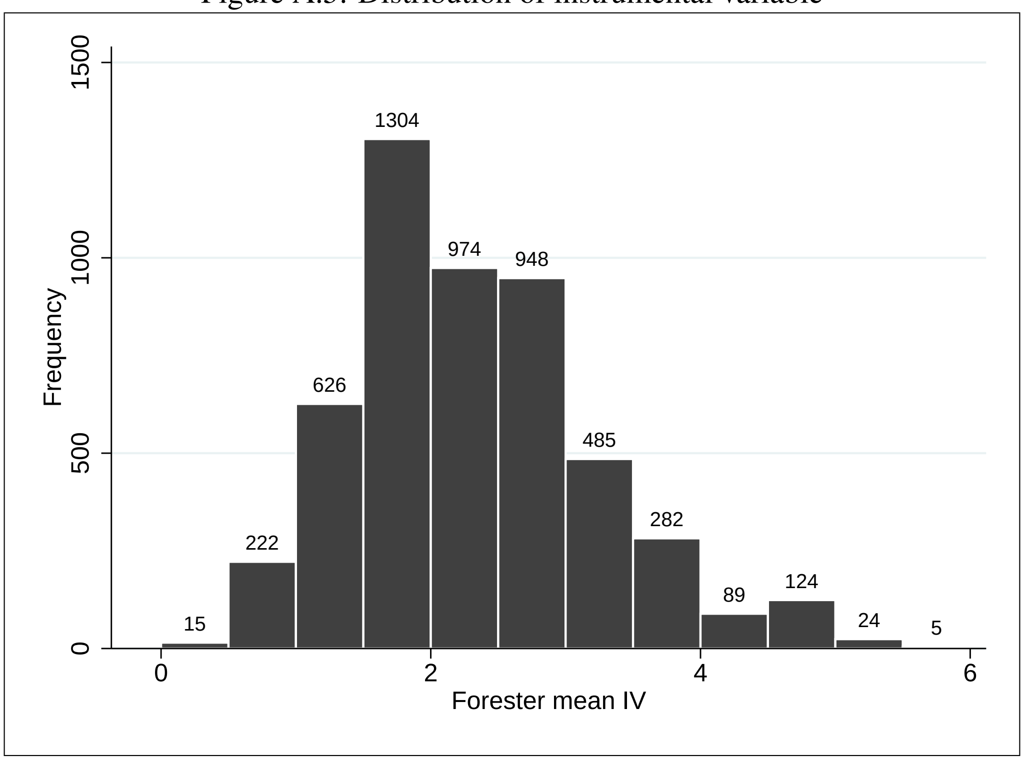

Note: Excludes sales appraised by foresters with fewer than 5 observations.

Figure A.6: Distribution of instrumental variable, demeaned within MU

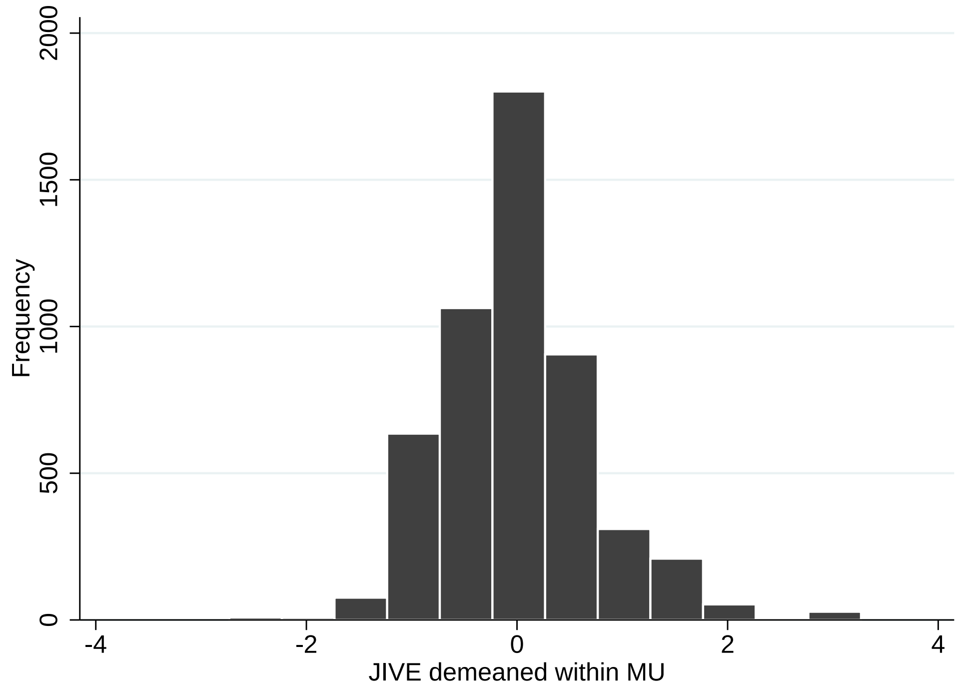

Note: Excludes sales appraised by foresters with fewer than 5 observations.

Figure A.7: Distribution of instrumental variable, standardized within MU

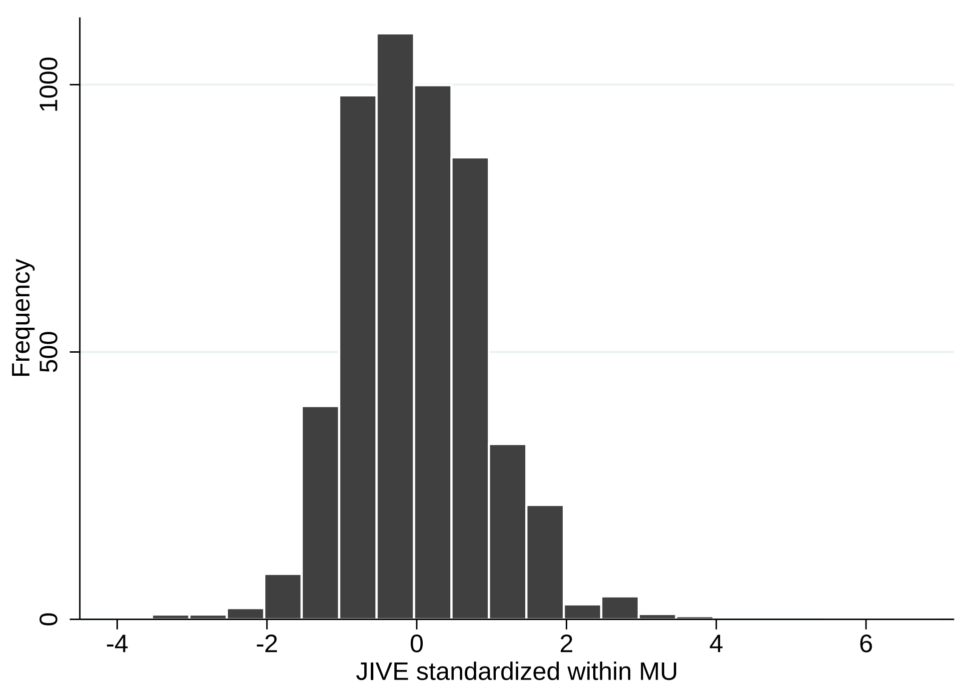

Note: Excludes sales appraised by foresters with fewer than 5 observations.

Table A.1: Linear regressions, effects of different seasons

VARIABLES	(1) Ln(win bid)	(2) Ln(win bid)	(3) Ln(win bid)	(4) Ln(win bid)
Spring Months Restr.	-0.001	-0.008	-0.023	-0.009
	(0.018)	(0.017)	(0.016)	(0.012)
X Other Months Restr.	0.004	-0.002	0.000	-0.001
	(0.011)	(0.010)	(0.009)	(0.008)
Summer Months Restr.	0.023	0.024	0.009	0.006
	(0.020)	(0.018)	(0.017)	(0.013)
X Other Months Restr.	-0.001	0.005	0.004	0.005
	(0.009)	(0.008)	(0.007)	(0.006)
Fall Months Restr.	-0.049	-0.004	0.001	-0.002
	(0.047)	(0.039)	(0.034)	(0.028)
X Other Months Restr.	-0.010	-0.016*	-0.011	-0.010
	(0.012)	(0.010)	(0.009)	(0.007)
Winter Months Restr.	-0.010	-0.016	-0.006	-0.003
	(0.016)	(0.015)	(0.014)	(0.012)
X Other Months Restr.	0.001	0.004	0.003	0.000
	(0.006)	(0.005)	(0.005)	(0.005)
Share Softwood: sawlogs	1.896***	1.846***	1.704***	1.659***
	(0.115)	(0.100)	(0.094)	(0.088)
Share Hardwood: sawlogs	1.515***	1.389***	1.183***	1.208***
	(0.045)	(0.046)	(0.050)	(0.046)
Share Hardwood: pulpwood	0.405***	0.204***	0.179***	0.157***
	(0.032)	(0.030)	(0.032)	(0.028)
Upper peninsula	0.346***	0.300***	0.294***	0.301***
	(0.029)	(0.026)	(0.028)	(0.020)
DNR cost factors	1.134***	1.112***	1.110***	0.805***
	(0.081)	(0.077)	(0.068)	(0.061)
Log acres		0.049***	0.036***	0.039***
		(0.010)	(0.009)	(0.008)
Species-product HHI		0.631***	0.676***	0.691***
		(0.042)	(0.042)	(0.038)
Percent bid species		0.724***	0.804***	0.684***
		(0.094)	(0.089)	(0.070)
Constant	3.042***	2.025***	1.918***	2.376***
	(0.063)	(0.100)	(0.097)	(0.087)
Observations	4,750	4,750	4,750	4,750
R-squared	0.424	0.528	0.576	0.647
Major species dummies	-	-	X	X
Quarter dummies	-	-	-	X
Year dummies	-	-	-	X

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

Table A.2: TSLS regressions

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6) OLS	(7) OLS
Months restricted	-0.0296	-0.0205	-0.0212	-0.0189	-0.0040	-0.0105***	-0.0083***
	(0.023)	(0.018)	(0.019)	(0.019)	(0.016)	(0.002)	(0.002)
Pine bark beetle note			0.0507	0.0475	0.0286	0.0403**	0.0311*
			(0.041)	(0.041)	(0.032)	(0.018)	(0.017)
Beech scale note			-0.0557	-0.0611	-0.0286	-0.0273	-0.0485
			(0.042)	(0.042)	(0.036)	(0.035)	(0.032)
Share Softwood: sawlogs		1.8104***	1.8156***	1.8107***	1.7437***	1.7567***	1.7143***
		(0.118)	(0.120)	(0.121)	(0.118)	(0.119)	(0.112)
Share Hardwood: sawlogs		1.3358***	1.3185***	1.3266***	1.3384***	1.3342***	1.2502***
		(0.068)	(0.072)	(0.070)	(0.063)	(0.062)	(0.056)
Share Hardwood: pulpwood		0.1554***	0.1427***	0.1426***	0.1381***	0.1372***	0.1641***
		(0.030)	(0.030)	(0.030)	(0.028)	(0.028)	(0.025)
Log acres		0.0309***	0.0316***	0.0321***	0.0393***	0.0390***	0.0454***
		(0.010)	(0.010)	(0.010)	(0.008)	(0.008)	(0.008)
Percent bid species		0.6888***	0.6928***	0.6942***	0.6268***	0.6211***	0.6084***
		(0.107)	(0.109)	(0.107)	(0.087)	(0.084)	(0.081)
Species-product HHI		0.7509***	0.7456***	0.7453***	0.7318***	0.7301***	0.6759***
		(0.049)	(0.049)	(0.049)	(0.050)	(0.050)	(0.045)
DNR cost factors							0.9042***
							(0.061)
Constant	4.5610***	2.8495***	2.8448***	2.8095***	3.0510***	3.0670***	2.3186***
	(0.077)	(0.143)	(0.144)	(0.147)	(0.118)	(0.106)	(0.109)
Observations	4,647	4,647	4,647	4,647	4,647	4,647	4,646
R-squared						0.617	0.664
Major species dummies	X	X	X	X	X	X	X
Quarter dummies	-	-	-	X	X	X	X
Year dummies	-	-	-	-	X	X	X
Management Unit dummies	X	X	X	X	X	X	X
First stage F-Stat	143.5	150.3	128.7	126.9	118.9		

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered by forester ID.

APPENDIX B.

Appendices for Chapter 2

B.1 Robustness to missing data

As noted earlier in the paper, the algorithm that Google Insights uses to distribute data publicly induces a censoring problem. There is no suitable instrument that varies at the week-state level that we could use for an exclusion restriction in a selection model. Hence, we carefully analyze the patterns of missing observations and rerun our main specification on several subsamples. First, Figure B.1 shows year-by-year maps depicting the number of weeks of missing data in each of the lower 48 states. After the first couple of lower search volume years, the missing observations are largely confined to the northern Rocky Mountain states (which are sparsely populated). The most severe case is Wyoming, which never appears in our sample.

In Table B.1, we restrict our sample to state-years in which there are no missing observations. That is, this sample consists only of the state-years on the maps that appear in the lightest shade of red. For convenience, the corresponding columns from Panel B of Table 2.3 are interspersed: these are the columns in which “Full sample” is labeled “Y”. The results are broadly similar; the main exception is that the effect of an unusually warm winter week is no longer significant. However, the effect of unusually little snow remains.

In Table B.2, we construct a balanced panel for the years 2007-2011. We determined that this set of years gave us the best balance between panel length (5 years) and inclusion of states (32). The set of states covers a broad swath of the country, both politically and geographically.

The columns in which the panel is unbalanced are the full sample from 2007-2011 to facilitate comparison. Again, the results of the balanced panel analysis are broadly comparable to those with the unbalanced panels, whether the unbalanced sample starts in 2004 or 2007.

B.2 Other Robustness Checks

Our main specifications include year-month and state-month of year fixed effects. The impact of weather on search has been identified off of within-state variation, controlling for state-specific seasonality and nationwide variation for the month. This is our preferred specification because it allows us to effectively compare observations with similar local conditions, while controlling for broad trends in weather and search intensity. However, we present the results using other sets of fixed effects in Table B.4.

Although the coefficients do change in magnitude across the specifications, they are remarkably robust. By comparing magnitudes, we can explore the relative importance of local and national variation in weather. Internet searches related to climate change may be driven by a combination of local and national weather. In addition, weather is likely to be correlated in the cross-section. If there is a heatwave in Iowa, it is likely that Ohio will be experiencing unusually hot weather as well. In this case, part of the search response in Iowa will be due to personal experience with the unusual weather and local media coverage. However, some of the response will also be due to interactions with family and friends in Ohio and regional or national news coverage of the broader heatwave. Our preferred specification includes year-month fixed effects and does not exploit national variation in weather. By comparing estimates with different sets of fixed effects, we can gauge the relative importance of local versus national trends.

In Column 4, we reproduce our preferred specification, with year-month fixed effects. These control for short-run nationwide events, such as a major heatwave or drought. Relative to the specification in column 3 which only include year fixed effects, the coefficients on the temperature deviation variables fall in magnitude by roughly 20%. This suggests that a part of our estimated weather-search effect may be driven by spillovers from national events and trends.

As we move through the columns, we add various richer combinations of fixed effects. We present our most flexible specification in Column 7. Here we allow for unobservables at the year-week and state-week of year level. Even after netting out all national variation in search intensity and weather, we get very similar results to our preferred specification.

In Figure B.3, we plot the year-week fixed effects from Column 7 in Panel A. Panel B shows the weekly average deviation in maximum temperature across the sample. Clearly, some level of nationwide time fixed effect is important for proper estimation. The four red lines signify peaks in the fixed effects corresponding to important climate-related news events: Hurricane Katrina, the release of the 2007 IPCC report[1], the *Massachusetts v. EPA* Supreme Court Clean Air Act ruling, and the "Climategate" scandal of late 2009.

Looking at Panel B, it is clear that Hurricane Katrina happened in a week that was unusually warm across the county, while Climategate (coincidentally) was followed by an unusually cold week. Controlling for such correlations ensures that we do not pick up search activity driven by contemporaneous nationwide news coverage and weather events.

Finally, our main specification uses state-level variation in weather and search intensity. Because the weather stations are not weighted by nearby population or search activity, we may measure weather with measurement error. In the case of small states such as Rhode Island, a state-level analysis might be a reasonable aggregation; however, a state like California is climatically, politically, and economically diverse. Luckily, the Google Insights tool also reports search intensity at the city level for major U.S. metropolitan areas. As a check, we also perform our analysis on the 25 largest U.S. cities as grouped by Google Insights. Table B.5 presents estimates from regressions that include year-month and city-month of year fixed effects; that is, they are analogous to our preferred state-level estimates from Table 2.3. The results reveal that the city-level relationships are quite similar to those found at the state level. That is, even when we restrict our data to a set of metropolitan areas, the relationship between weather and search intensity holds.

[1]This report stated that recent climate change is anthropogenic with greater than 90% confidence.

B.3 Figures and Tables

Figure B.1: Missing observations, 2004-2007

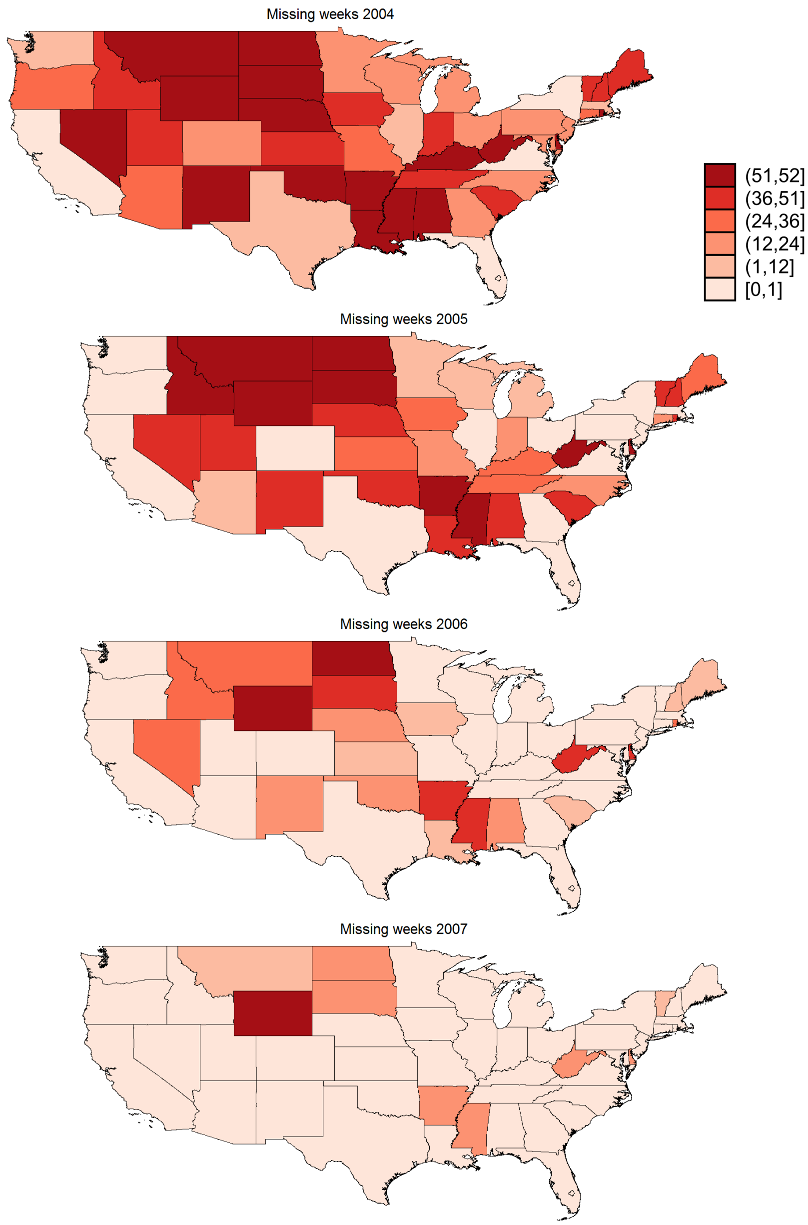

Figure B.2: Missing observations, 2008-2011

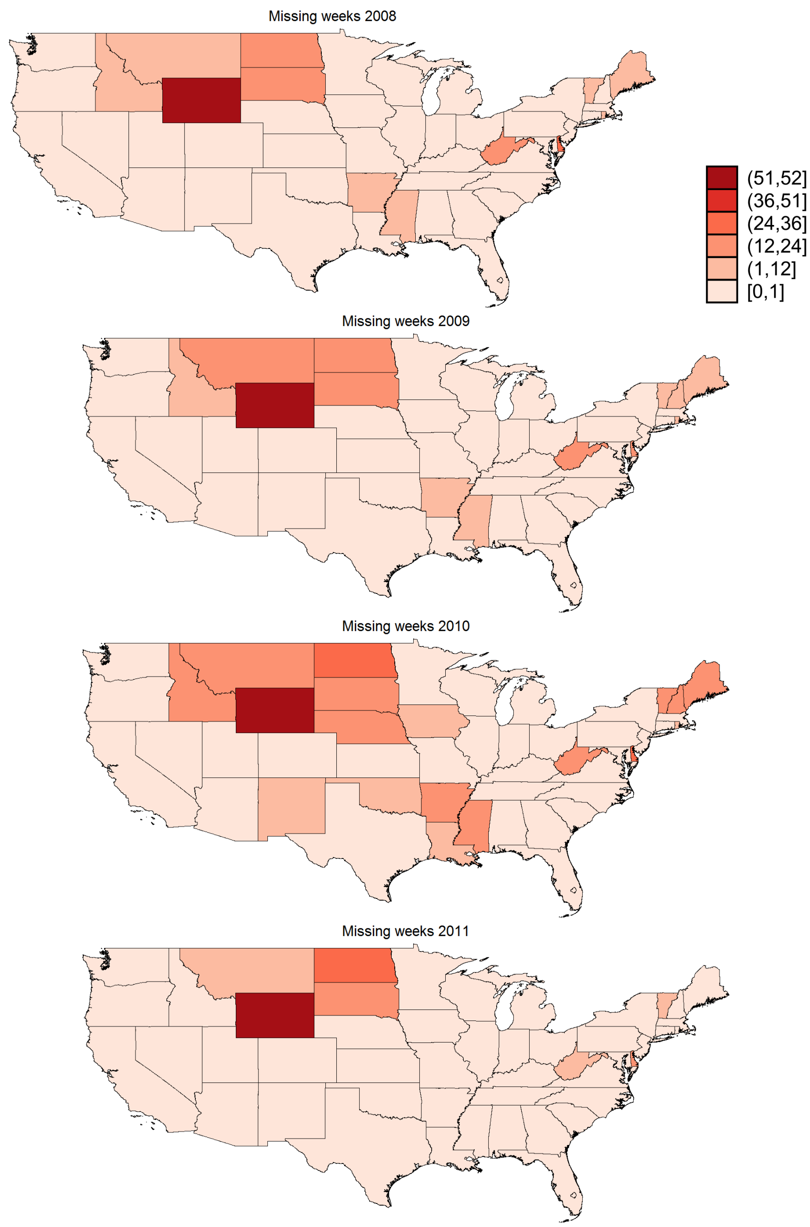

Figure B.3: Year-week fixed effects

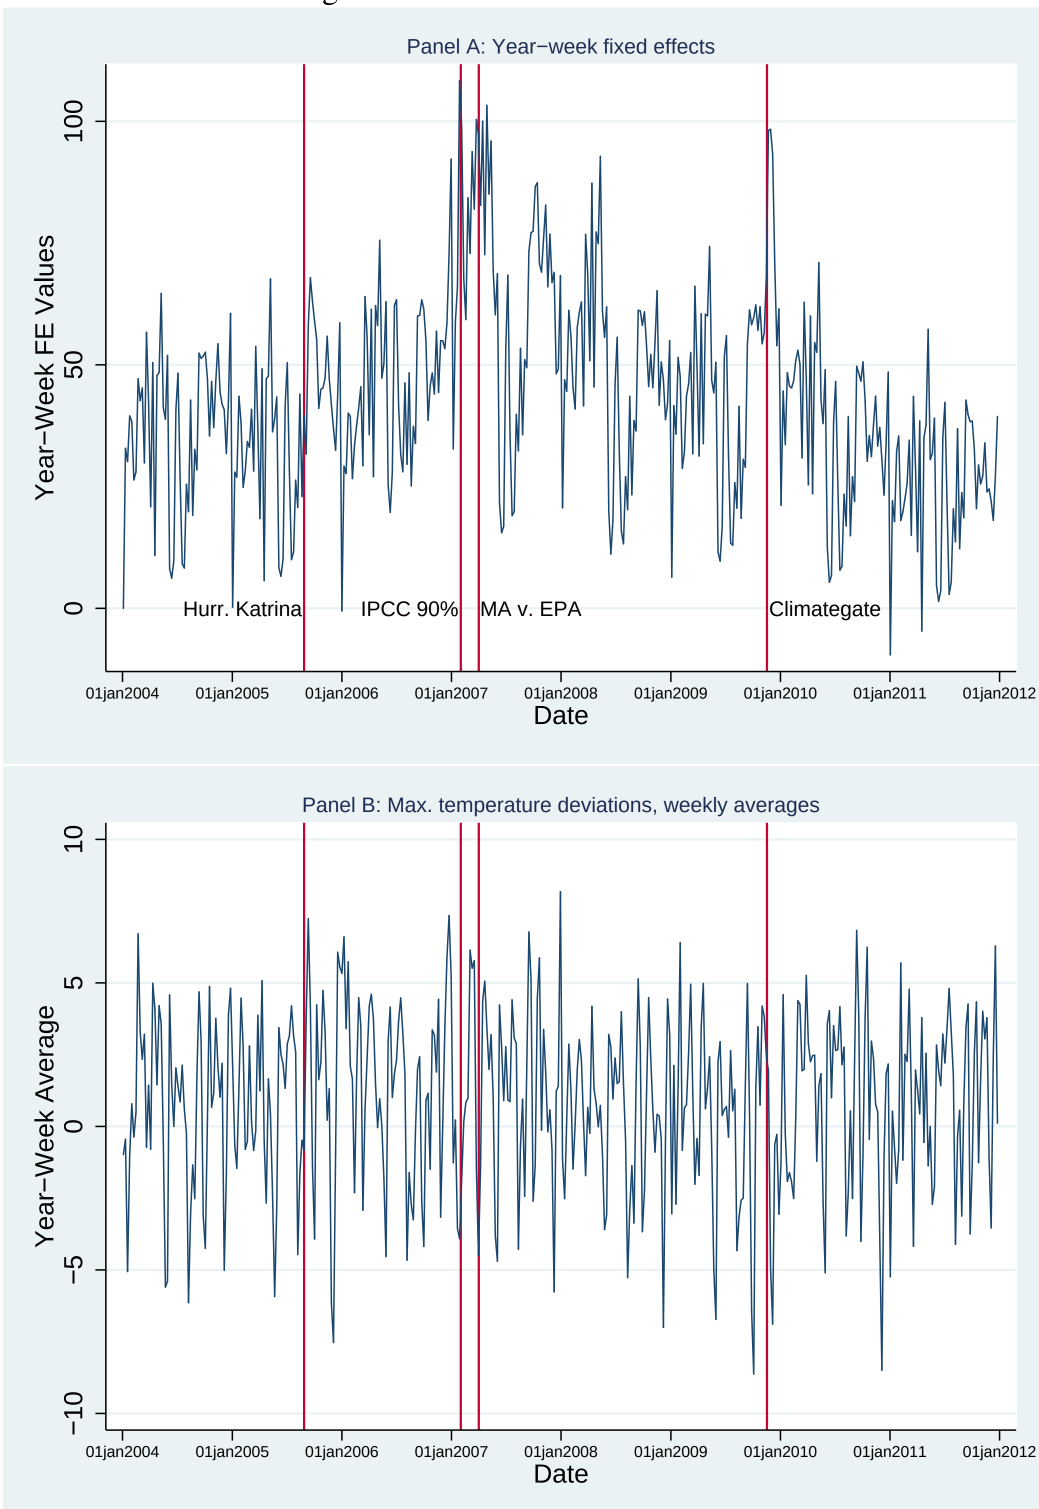

Table B.1: Robustness to Missing Data

VARIABLES	(1) indexCC	(2) All Seasons	(3) Winter	(4) Winter	(5) Spring	(6) Spring	(7) Summer	(8) Summer	(9) Fall	(10) Fall
Pos dev, Max Temp, deg. C	0.292***	0.262***	0.547***	0.214	-0.110	0.0454	0.707***	0.725***	0.0822	0.179
	(0.0785)	(0.0873)	(0.155)	(0.161)	(0.142)	(0.157)	(0.163)	(0.126)	(0.178)	(0.147)
Neg dev, Max Temp, deg. C	0.806***	0.792***	1.634***	1.469***	-0.0256	0.0782	0.322*	0.263*	0.246*	0.458***
	(0.101)	(0.0978)	(0.172)	(0.162)	(0.139)	(0.136)	(0.192)	(0.144)	(0.133)	(0.141)
Pos dev, Precip., mm	0.0570	0.128*	0.844***	0.886***	-0.0955	-0.0291	0.0176	0.0600	-0.0796	0.0254
	(0.0733)	(0.0729)	(0.160)	(0.197)	(0.133)	(0.124)	(0.0737)	(0.0801)	(0.110)	(0.105)
Neg dev, Precip., mm	0.122	0.294**	1.048***	1.426***	-0.311	-0.226	-0.285**	-0.240*	-0.159	-0.0382
	(0.124)	(0.122)	(0.244)	(0.291)	(0.203)	(0.199)	(0.139)	(0.138)	(0.161)	(0.135)
Pos dev, Snowfall, mm	0.00341	0.0847***	-0.0472	0.0254	0.0235	0.119	7.961*	8.946	0.132*	0.113
	(0.0289)	(0.0197)	(0.0294)	(0.0231)	(0.0661)	(0.0739)	(4.055)	(5.469)	(0.0700)	(0.0846)
Neg dev, Snowfall, mm	0.285***	0.347***	0.256**	0.339***	-0.0393	0.0522	11.32	22.77	0.774**	0.516
	(0.0883)	(0.0870)	(0.101)	(0.0975)	(0.117)	(0.146)	(29.21)	(33.76)	(0.317)	(0.481)
Pos dev, Snow Depth, mm	-0.00705	-0.0151*	-0.00348	-0.0127	-0.0325*	-0.0367	0.0272	-0.0352	0.0193	0.0484
	(0.00866)	(0.00813)	(0.00940)	(0.00830)	(0.0192)	(0.0266)	(0.0367)	(0.0213)	(0.0549)	(0.0517)
Neg dev, Snow Depth, mm	0.0446**	0.0651*	0.0761***	0.0767*	-0.0670	-0.00581	0.540	0.111	0.300**	0.479*
	(0.0202)	(0.0349)	(0.0229)	(0.0410)	(0.0599)	(0.0222)	(0.527)	(0.132)	(0.122)	(0.251)
Constant	21.41***	17.52***	16.19***	12.47***	27.03***	40.56***	90.71***	106.4***	60.04***	136.0***
	(1.662)	(1.683)	(1.874)	(2.775)	(1.218)	(2.825)	(2.086)	(0.893)	(1.416)	(4.906)
Observations	16,546	12,302	4,269	3,025	4,320	3,105	3,485	3,106	4,472	3,066
R-squared	0.763	0.791	0.696	0.720	0.781	0.816	0.783	0.776	0.741	0.762
Full Sample	Y	N	Y	N	Y	N	Y	N	Y	N

Notes: *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the Google search index. Columns for which the full sample is not used include only year-state combinations for which there are no missing observations. All regressions also include year * month FE and state * month of year FE. Standard errors are clustered at the state level.

Table B.2: Robustness to Balanced Panel 2007-2011

VARIABLES	(1) All Seasons	(2) All Seasons	(3) Winter	(4) Winter	(5) Spring	(6) Spring	(7) Summer	(8) Summer	(9) Fall	(10) Fall
Pos dev, Max Temp, deg. C	0.281***	0.243**	0.631***	0.246	-0.0830	0.0201	0.593***	0.731***	0.0133	0.0898
	(0.0878)	(0.105)	(0.199)	(0.186)	(0.183)	(0.180)	(0.192)	(0.157)	(0.187)	(0.153)
Neg dev, Max Temp, deg. C	0.895***	0.755***	1.819***	1.453***	-0.136	-0.168	0.270	0.359*	0.325*	0.545***
	(0.115)	(0.102)	(0.202)	(0.146)	(0.179)	(0.178)	(0.202)	(0.188)	(0.186)	(0.152)
Pos dev, Precip., mm	0.0458	0.196**	1.132***	1.000***	-0.219	0.0704	-0.00924	0.0884	-0.192	-0.00819
	(0.103)	(0.0834)	(0.193)	(0.210)	(0.177)	(0.132)	(0.0940)	(0.0874)	(0.159)	(0.107)
Neg dev, Precip., mm	0.103	0.273*	1.401***	1.466***	-0.559**	-0.418*	-0.252	-0.0986	-0.403	-0.152
	(0.161)	(0.152)	(0.307)	(0.346)	(0.243)	(0.229)	(0.187)	(0.191)	(0.243)	(0.150)
Pos dev, Snowfall, mm	-0.00246	0.0550**	-0.0537	-0.00857	0.00861	0.0553	8.107*	9.565	0.0896	0.157**
	(0.0323)	(0.0215)	(0.0350)	(0.0302)	(0.0747)	(0.0483)	(4.234)	(6.178)	(0.0718)	(0.0695)
Neg dev, Snowfall, mm	0.174	0.277***	0.0762	0.227**	-0.0245	-0.0434	13.72	3.248	0.661*	0.732*
	(0.168)	(0.0983)	(0.186)	(0.110)	(0.117)	(0.159)	(31.28)	(17.97)	(0.387)	(0.392)
Pos dev, Snow Depth, mm	-0.00663	-0.0145	-0.00292	-0.0148	-0.0258	-0.0145	0.0615	-0.0290*	-0.00461	-0.0155
	(0.00770)	(0.01000)	(0.00939)	(0.00972)	(0.0185)	(0.0265)	(0.0600)	(0.0157)	(0.0569)	(0.0336)
Neg dev, Snow Depth, mm	0.0253	0.0114	0.0722**	0.00879	-0.123	0.00638	3.511*	1.829***	0.186	0.342***
	(0.0386)	(0.0195)	(0.0318)	(0.0198)	(0.0842)	(0.0252)	(1.848)	(0.475)	(0.164)	(0.120)
Constant	65.52***	15.78***	28.44***	52.52***	89.87***	87.03***	90.38***	28.16***	68.90***	44.14***
	(1.449)	(1.043)	(1.857)	(1.813)	(1.870)	(1.353)	(1.555)	(0.624)	(1.411)	(1.046)
Observations	12,179	8,270	3,112	2,007	3,207	2,099	2,701	2,099	3,159	2,065
R-squared	0.761	0.786	0.669	0.672	0.767	0.811	0.763	0.691	0.749	0.758
Balanced panel	N	Y	N	Y	N	Y	N	Y	N	Y

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Dependent variable is the Google search index. Columns with a balanced panel are a balanced panel for of states that are not missing any observations from 2007-2011: AL, AZ, CA, CO, CT, DC, FL, GA, HI, IL, IN, KS, KY, MA, MD, MI, MN, MO, NC, NJ, NV, NY, OH, OR, PA, SC, TN, TX, UT, VA, WA, WI. Columns without a balanced panel include all observations from 2007-2011. All regressions also include year * month FE and state * month of year FE. Standard errors are clustered at the state level.

Table B.3: Regressions by Month

	(1) Jan	(2) Feb	(3) Mar	(4) Apr	(5) May	(6) Jun	(7) Jul	(8) Aug	(9) Sep	(10) Oct	(11) Nov	(12) Dec
Pos dev, Max Temp, deg. C	0.281	0.692***	-0.106	-0.151	0.0591	0.667**	0.910***	0.420	0.824**	-0.306	-0.0129	0.671**
	(0.226)	(0.223)	(0.316)	(0.216)	(0.280)	(0.252)	(0.218)	(0.268)	(0.327)	(0.262)	(0.227)	(0.309)
Neg dev, Max Temp, deg. C	1.227***	1.267***	-0.123	-0.485**	0.518	-0.354	0.580**	0.545***	0.104	-0.295*	1.289***	2.413***
	(0.236)	(0.263)	(0.229)	(0.216)	(0.360)	(0.427)	(0.225)	(0.195)	(0.225)	(0.157)	(0.404)	(0.271)
Pos dev, Precip., mm	0.0401	0.635*	-0.0783	-0.583	0.202	0.284*	-0.0904	-0.0537	0.000470	-0.0383	-0.0269	1.325***
	(0.233)	(0.335)	(0.157)	(0.368)	(0.193)	(0.167)	(0.0909)	(0.117)	(0.108)	(0.185)	(0.219)	(0.250)
Neg dev, Precip., mm	0.611	0.553	-0.357	-0.378	-0.253	-0.626**	-0.183	0.177	-0.0636	0.133	-0.777**	1.342***
	(0.368)	(0.391)	(0.354)	(0.491)	(0.382)	(0.249)	(0.239)	(0.260)	(0.246)	(0.321)	(0.322)	(0.387)
Pos dev, Snowfall, mm	-0.0866*	0.00124	-0.0327	0.116	1.177***	9.337**	-8.101	428.3*	9.039	-0.0698	-0.0617	0.0763
	(0.0472)	(0.0385)	(0.0892)	(0.177)	(0.333)	(4.277)	(63.94)	(225.4)	(5.795)	(0.200)	(0.112)	(0.0739)
Neg dev, Snowfall, mm	-0.0488	0.426**	-0.0149	-0.543	3.300*	12.12	-196.7***	168.8***	3.339	0.972	0.884**	0.581***
	(0.154)	(0.196)	(0.117)	(0.400)	(1.720)	(28.06)	(71.64)	(51.09)	(3.065)	(0.765)	(0.361)	(0.205)
Pos dev, Snow Depth, mm	0.0334***	0.00822	-0.0334*	-0.00929	-0.0396	0.00800	0.0619**	0.199***	-6.603	0.258	-0.0344	-0.0820***
	(0.00958)	(0.00920)	(0.0182)	(0.0454)	(0.0369)	(0.0488)	(0.0298)	(0.0483)	(4.131)	(0.183)	(0.0518)	(0.0288)
Neg dev, Snow Depth, mm	0.124***	0.0562**	-0.0478	-0.188	-0.176	0.667	-0.940***	2.757***	6.880***	0.340***	0.284*	0.0347
	(0.0382)	(0.0224)	(0.0435)	(0.136)	(0.248)	(0.557)	(0.175)	(0.292)	(0.652)	(0.0763)	(0.143)	(0.0343)
Constant	19.25***	27.05***	26.75***	34.58***	16.61***	18.31***	15.35***	23.27***	24.25***	28.48***	34.15***	16.45***
	(2.106)	(1.569)	(1.474)	(1.311)	(1.680)	(1.705)	(1.448)	(1.720)	(1.110)	(1.597)	(1.606)	(2.477)
Observations	1,419	1,349	1,423	1,448	1,449	1,122	1,178	1,185	1,346	1,594	1,532	1,501
R-squared	0.725	0.775	0.785	0.792	0.739	0.761	0.820	0.767	0.749	0.772	0.699	0.630

Notes: *** p<0.01, ** p<0.05, * p<0.1. Dependent variable is the Google search index. All regressions include state and year FE. Standard errors are clustered at the state level.

Table B.4: Sensitivity to the Inclusion of Various Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Pos dev, Max Temp, deg. C	0.808**	0.912***	0.381***	0.292***	0.483***	0.367***	0.229***
	(0.317)	(0.0859)	(0.0858)	(0.0785)	(0.0803)	(0.0684)	(0.0785)
Neg dev, Max Temp, deg. C	1.899***	1.505***	1.050***	0.806***	0.887***	0.537***	0.420***
	(0.293)	(0.119)	(0.107)	(0.101)	(0.114)	(0.102)	(0.112)
Pos dev, Precip., mm	-0.220	0.132	0.134	0.0570	0.0572	-0.0257	-0.0853
	(0.309)	(0.0862)	(0.0843)	(0.0733)	(0.107)	(0.0782)	(0.0841)
Neg dev, Precip., mm	-0.469	0.294*	0.206*	0.122	0.102	0.0286	-0.0812
	(0.616)	(0.167)	(0.116)	(0.124)	(0.127)	(0.108)	(0.117)
Pos dev, Snowfall, mm	-0.0131	0.111***	-0.0293	0.00341	-0.0476	-0.00165	-0.00185
	(0.0848)	(0.0302)	(0.0345)	(0.0289)	(0.0427)	(0.0364)	(0.0388)
Neg dev, Snowfall, mm	1.492***	0.962***	0.419***	0.285***	0.436***	0.281***	0.294***
	(0.382)	(0.124)	(0.0851)	(0.0883)	(0.0868)	(0.0727)	(0.0858)
Pos dev, Snow Depth, mm	0.0609***	0.00432	-0.00291	-0.00705	0.00938	0.00877	0.0123
	(0.0200)	(0.00911)	(0.0102)	(0.00866)	(0.0125)	(0.0104)	(0.0108)
Neg dev, Snow Depth, mm	0.168***	0.0460**	0.0434**	0.0446**	0.0547**	0.0596**	0.0640**
	(0.0392)	(0.0203)	(0.0200)	(0.0202)	(0.0246)	(0.0247)	(0.0244)
Constant	33.93***	25.04***	25.99***	21.41***	16.13***	12.31***	12.78***
	(2.452)	(0.936)	(0.854)	(1.662)	(0.874)	(1.796)	(2.969)
Observations	16,546	16,546	16,546	16,546	16,546	16,546	16,546
R-squared	0.072	0.597	0.695	0.763	0.737	0.811	0.837
Year FE	-	X	X	-	X	-	-
Year-Month FE	-	-	-	X	-	X	-
Year-Week FE	-	-	-	-	-	-	X
State FE	-	X	-	-	-	-	-
State-MOY FE	-	-	X	X	-	-	-
State-WOY FE	-	-	-	-	X	X	X

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the state level. Dependent variable is the Google search index.

Table B.5: Asymmetric effects of weather deviations: 25 Largest Cities

	(1) All Seasons	(2) Winter	(3) Spring	(4) Summer	(5) Fall
Pos dev, Max Temp, deg. C	0.477***	0.598***	0.187	0.948***	0.280**
	(0.0904)	(0.170)	(0.177)	(0.223)	(0.128)
Neg dev, Max Temp, deg. C	0.715***	1.464***	-0.111	0.0826	0.466***
	(0.106)	(0.175)	(0.160)	(0.162)	(0.153)
Pos dev, Precip., mm	-0.0145	0.0889	-0.0809	0.0787	0.0558
	(0.0527)	(0.141)	(0.115)	(0.0700)	(0.0634)
Neg dev, Precip., mm	0.214*	0.666***	0.0560	0.0118	-0.0851
	(0.119)	(0.176)	(0.260)	(0.156)	(0.215)
Pos dev, Snowfall, mm	-0.0283	-0.0140	-0.0877**	-0.220	0.0630*
	(0.0189)	(0.0247)	(0.0346)	(0.490)	(0.0333)
Neg dev, Snowfall, mm	0.231**	0.197**	-0.0450	12.94	1.906***
	(0.0933)	(0.0885)	(0.167)	(21.82)	(0.652)
Pos dev, Snow Depth, mm	0.000300	-0.00916	0.00472	0.00507**	-0.00198
	(0.00510)	(0.00944)	(0.00405)	(0.00185)	(0.0138)
Neg dev, Snow Depth, mm	0.0259*	0.0316**	0.0212	-0.000867	-0.135
	(0.0139)	(0.0146)	(0.0213)	(0.0357)	(0.144)
Constant	17.65***	9.426***	25.53***	25.70***	40.16***
	(1.245)	(1.704)	(2.472)	(1.006)	(1.594)
Observations	8,868	2,201	2,248	2,072	2,347
R-squared	0.738	0.635	0.791	0.658	0.687

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the city level. Dependent variable is the Google search index. Regression includes 25 largest metro areas, as grouped by Google: Atlanta, Baltimore, Boston, Chicago, Dallas/Fort Worth, Denver, Detroit, Houston, Los Angeles, Miami, Minneapolis/St. Paul, New York, Orlando, Philadelphia, Phoenix, Pittsburgh, Portland, Sacramento, San Antonio, San Diego, San Francisco, Seattle/Tacoma, St. Louis, Tampa, Washington (DC). All regressions also include year-month FE and city-month of year FE.

APPENDIX C.

Appendices for Chapter 3

C.1 Proof of cutoff equilibrium with an additive value shock

Suppose that $v_2 = v_1 + \omega$, where $\omega \sim F_\omega$ is a firm-specific shock independent of v_1. Let $F_2(v; v^*)$ be the distribution of v_2, given that only firms with $v_1 < v^*$ would be bidding in the second auction.

Consider some valuation $v \geq R_1$. Then the best expected payoff from bidding in the first auction is

$$
\begin{aligned}
(v_1 - b_1(v_1))F(v_1)^{N-1}, &\qquad \text{if } v_1 \geq v^* \\
(v_1 - R_1)F(v^*)^{N-1}, &\qquad \text{if } v_1 \in [R_1, v^*)
\end{aligned}
$$

The derivative of this payoff with respect to v is $F(v)^{N-1}$ if $v \geq v^*$ and $F(v^*)^{N-1}$ if $v \in [R_1, v^*)$. Note that this derivative is always $\geq F(v^*)^{N-1}$.

The expected payoff from waiting is given by

$$
\delta F(v^*)^{N-1} \int_{R_2 - v_1}^{\infty} (v_1 + \omega - b_2(v_1 + \omega))F_2(v_1 + \omega; v^*)^{N-1} dF_\omega
$$

Differentiating this with respect to v gives:

$$\begin{aligned}
&\delta F(v^*)^{N-1} \int_{R_2-v_1}^{\infty} \{(1 - b_2{}'(v_1+\omega))F_2(v_1+\omega; v^*)^{N-1} \\
&\qquad + (v_1 + \omega - b_2(v_1+\omega))F_2(v_1+\omega; v^*)^{N-2} f_2(v_1+\omega; v^*)(N-1)dF_\omega\} \\
&\qquad + \underbrace{(R_2 - b_2(R_2))F(R_2)^{N-1}}_{=0} \\
&= \delta F(v^*)^{N-1} \int_{R_2-v_1}^{\infty} \{F_2(v_1+\omega; v^*)^{N-1} \\
&\qquad - b_2{}'(v_1+\omega)F_2(v_1+\omega; v^*)^{N-1} \\
&\qquad + (v_1 + \omega - b_2(v_1+\omega))F_2(v_1+\omega; v^*)^{N-2} f_2(v_1+\omega; v^*)(N-1)dF_\omega\} \\
&= \delta F(v^*)^{N-1} \int_{R_2-v_1}^{\infty} F_2(v_1+\omega; v^*)^{N-1} dF_\omega
\end{aligned}$$

The final equality holds because the last two lines of the previous expression equal zero: they are the first order condition for optimal bidding in the second-round auction for a given value of $v_2 = v_1 + \omega$. Using this new expression, I can confirm that the relative gains from bidding are increasing in v_1:

$$\begin{aligned}
\frac{d}{dv_1}E[\text{Payoff from bidding in auction 1}] &\geq F(v^*)^{N-1} \\
&> \delta F(v^*)^{N-1} \\
&> \delta F(v^*)^{N-1}(1 - F_\omega(R_2 - v_1)) \\
&= \delta F(v^*)^{N-1} \int_{R_2-v_1}^{\infty} f_\omega(\omega) d\omega \\
&> \delta F(v^*)^{N-1} \int_{R_2-v_1}^{\infty} F_2(v_1+\omega)^{N-1} dF_\omega \\
&= \frac{d}{dv_1}E[\text{Payoff from waiting for auction 2}]
\end{aligned}$$

Thus, given that a bidder with $v_1 = v^*$ is indifferent between bidding in auction 1 and not bidding, all $v < v^*$ prefer to wait, and all $v > v^*$ prefer to bid in auction 1.

C.2 Proof of cutoff equilibrium with no shocks to values

In this subsection, I will prove that a version of the strategies outlined above is an equilibrium in the case of constant valuations (i.e., $v_2(v_1, \omega) = v_1$).

Suppose that bidder strategies are the following: In auction 1, a firm bids above zero only if it has a valuation $v \geq v^*$. These bids are of the form

$$b_1(v) = v - \frac{\int_{v^*}^{v} F(u)^{N-1} du}{F(v)^{N-1}} - (v^* - R_1)\frac{F(v^*)^{N-1}}{F(v)^{N-1}}.$$

Otherwise the firm bids zero.

In the second auction, it must be that only firms with valuations $v < v^*$ are competing. These firms bid as follows:

$$b_2(v) = v - \frac{\int_{R_2}^{v} F_2(u)^{N-1} du}{F_2(v)^{N-1}},$$

where $F_2(v) = F(v|v < v^*)$ is the distribution of values once it is known that no bidder has a valuation above the cutoff. In the case that a bidder at or above the cutoff strays from equilibrium and bids in the second round auction, she bids according to the above expression evaluated at v^*.

The cutoff v^* is the type that is exactly indifferent between bidding in auction 1 and waiting for auction 2:

$$(v^* - R_1)F(v^*)^{N-1} = \delta F(v^*)^{N-1}\left(v^* - v^* + \frac{\int_{R_2}^{v^*} F_2(u)^{N-1} du}{F_2(v^*)^{N-1}}\right)$$
$$\Rightarrow (v^* - R_1)F(v^*)^{N-1} = \delta F(v^*)^{N-1}\left(\int_{R_2}^{v^*} F_2(u)^{N-1} du\right)$$

because $F_2(v^*) = 1$.

Consider some bidder with value $v > v^*$. For this to be an equilibrium, this bidder must prefer to bid in auction 1, rather than hoping for auction 2 to arrive. If the bidder bids, they get the following expected payoff:

$$(v - b_1(v))F(v)^{N-1} = \int_{v^*}^{v} F(u)^{N-1}du + (v^* - R_1)F(v^*)^{N-1}$$

If the bidder waits for the second auction, it arrives with probability $F(v^*)^{N-1}$. This bidder does best by placing the same bid that a bidder with $v = v^*$ would. The expected payoff is then:

$$\begin{aligned}\delta F(v^*)^{N-1}F_2(v^*)^{N-1}(v - b_2(v^*)) &= \delta F(v^*)^{N-1}\left(v - v^* + \frac{\int_{R_2}^{v^*} F_2(u)^{N-1}du}{F_2(v^*)^{N-1}}\right)F_2(v^*)^{N-1} \\ &= \delta F(v^*)^{N-1}\left((v - v^*) + \int_{R_2}^{v^*} F_2(u)^{N-1}du\right)\end{aligned}$$

Differentiating the relative gain from bidding in the first auction establishes that the gain is increasing in v:

$$\begin{aligned}\frac{d}{dv}\Psi_1(v) &= \frac{d}{dv}\int_{v^*}^{v} F(u)^{N-1}du + (v^* - R_1)F(v^*)^{N-1} - \delta F(v^*)^{N-1}\left((v - v^*) + \int_{R_2}^{v^*} F_2(u)^{N-1}du\right) \\ &= F(v)^{N-1} - \delta F(v^*)^{N-1} > 0.\end{aligned}$$

Since $\Psi_1(v^*) = 0$, $\Psi_1(v) > 0$ for any $v > v^*$, proving the result.

Establishing a similar result for $v < v^*$ follows similar logic. For any $v < R_1$, the result is trivial. For $v \in [R_1, v^*]$, this type of bidder will do best when bidding in auction 1 by submitting the reserve price. So the relative gain from bidding in the first auction is given by:

$$\begin{aligned}\Psi_2(v) &= (v - R_1)F(v^*)^{N-1} - \delta F(v^*)^{N-1}(v - b_2(v))F_2(v)^{N-1} \\ &= (v - R_1)F(v^*)^{N-1} - \delta F(v^*)^{N-1}(v - b_2(v))F_2(v)^{N-1} \\ &= (v - R_1)F(v^*)^{N-1} - \delta F(v^*)^{N-1}\int_{R_2}^{v} F_2(u)^{N-1}du\end{aligned}$$

Differentiating Ψ_2 with respect to v reveals that this function is increasing in the bidder's valuation:

$$\frac{d}{dv}\Psi_2(v) = F(v^*) - F(v^*)\delta F_2(v)^{N-1} > 0.$$

Since $\Psi_2(v^*) = 0$, $\Psi_2(v) < 0$ for any $v < v^*$, proving the result.

BIBLIOGRAPHY

ACKERBERG, D. (2009): "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*.

ARAL, K. D., D. R. BEIL, AND L. N. V. WASSENHOVE (2014): "Total-Cost Procurement Auctions with Sustainability Audits to Inform Bid Markups," *Working Paper*.

ARYAL, G., M. F. GABRIELLI, AND Q. VUONG (2014): "Semiparametric Estimation of First-Price Auction Models," *Working Paper*.

ATHEY, S., D. COEY, AND J. LEVIN (2013): "Set-Asides and Subsidies in Auctions," *American Economic Journal: Microeconomics*.

ATHEY, S., P. CRAMTON, AND A. INGRAHAM (2002): "Setting the Upset Price in British Columbia Timber Auctions," Discussion paper, Market Design, Inc. and Criterion Auctions,.

ATHEY, S., J. LEVIN, AND E. SEIRA (2011): "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions," *Quarterly Journal of Economics*.

BAJARI, P., AND G. LEWIS (2011): "Procurement Contracting with Time Incentives: Theory and Evidence," *Quarterly Journal of Economics*.

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

BHATTACHARYA, V., J. W. ROBERTS, AND A. SWEETING (2014): "Regulating Bidder Participation in Auctions," *NBER Working Paper 19352*.

BRIANS, C., AND M. WATTENBERG (1996): "Campaign Issue Knowledge and Salience: Comparing Reception from TV Commercials, TV News and Newspapers," *American Journal of Political Science*.

BROWN, J., J. HASTINGS, E. MANSUR, AND S. VILLAS-BOAS (2008): "Reformulating Competition? Gasoline Content Regulation and Wholesale Gasoline Prices," *Journal of Environmental Economics and Management*.

BROWN, R., M. KILGORE, J. COGGINS, AND C. BLINN (2012): "The impact of timber-sale tract, policy, and administrative characteristics on state stumpage prices: An econometric analysis," *Forest Policy and Economics*.

Brown, R., M. Kilgore, J. Coggins, C. Blinn, and C. Pfender (2010): "Assessing State Timber Sale Policies, Programs, and Stumpage Price Drivers," Department of Forest Resources, University of Minnesota.

Burguet, R., and J. Sákovics (1996): "Reserve Prices without Commitment," *Games and Economic Behavior*.

Burstein, P. (2003): "The Impact of Public Opinion on Public Policy: A Review and an Agenda," *Political Research Quarterly*.

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

Chodorow-Reich, G. (2014): "The Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008-2009 Financial Crisis," *The Quarterly Journal of Economics*.

Choi, H., and H. Varian (2009): "Predicting the Present with Google Trends," *Mimeo, Google Inc.*

Deryugina, T. (2013): "How do people update? The effects of local weather fluctuations on beliefs about global warming," *Climatic Change*.

Egan, P. J., and M. Mullin (2012): "Turning Personal Experience into Political Attitudes: The Effect of Local Weather on Americans' Perceptions about Global Warming," *Journal of Politics*.

Fabra, N., and M. Reguant (2014): "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*.

Gentry, M., and C. Stroup (2014): "Entry and Competition in Takeover Auctions," *Mimeo*.

George Edwards, W. M., and R. Welch (1995): "Explaining Presidential Approval: The Significance of Issue Salience," *American Journal of Political Science*.

Groeger, J. (2014): "A Study of Participation in Dynamic Auctions," *International Economic Review*.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*.

Haile, P. A., and E. Tamer (2003): "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*.

Hamilton, L. C., and M. D. Stampone (2013): "Blowin' in the Wind: Short-term Weather and Belief in Anthropogenic Climate Change," *Weather, Climate, and Society*.

Hansen, J., M. Sato, and R. Ruedy (2012): "Perception of Climate Change," *Proceedings of the National Academy of Sciences*.

Herrnstadt, E. (2015): "Conservation Versus Competition? Environmental Objectives in Government Contracting," *Working Paper*.

HEYM, D. (2013): Private Communication,.

HICKMAN, B. R., AND T. P. HUBBARD (forthcoming): "Replacing Sample Trimming with Boundary Correction in Nonparameteric Estimation of First-Price Auctions," *Journal of Applied Econometrics*.

HOLT, C. A. (1980): "Competitive Bidding for Contracts under Alternative Auction Procedures," *Journal of Political Economy*.

HUBBARD, T. P., AND H. J. PAARSCH (2009): "Investigating bid preferences at low-price, sealed-bid auctions with endogenous participation," *International Journal of Industrial Organization*.

HUSSAIN, A., AND D. N. LABAND (2005): "The tragedy of the political commons: Evidence from U.S. Senate roll call votes on environmental protection," *Public Choice*.

JOFRE-BONET, M., AND M. PESENDORFER (2003): "Estimation of a Dynamic Auction Game," *Econometrica*.

KAHN, M. E. (2002): "Demographic Change and the Demand for Environmental Regulation," *Journal of Policy Analysis and Management*.

KAHN, M. E., AND M. J. KOTCHEN (2011): "Business Cycle Effects on Concern About Climate Change: The Chilling Effect of Recession," *Climate Change Economics*.

KALT, J. P., AND M. A. ZUPAN (1984): "Capture and Ideology in the Economy Theory of Politics," *American Economic Review*.

KRASNOKUTSKAYA, E. (2009): "Identification and Estimation in Procurement Auctions Under Unobserved Auction Heterogeneity," *Review of Economic Studies*.

KRASNOKUTSKAYA, E., AND K. SEIM (2011): "Bid Preference Programs and Participation in Procurement," *American Economic Review*.

LEVIN, D., AND J. L. SMITH (1994): "Equilibrium in Auctions with Entry," *American Economic Review*.

LEVITT, S. D. (1996): "How Do Senators Vote? Disentangling the Role fo Voter Preferences, Party Affiliation and Senator Ideology," *American Economic Review*.

LI, T., AND X. ZHENG (2012): "Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions," *Journal of Econometrics*.

MACKES, K. (2004): "Sawmill Short Course: Log Inputs – Measurement and Conversion Factors," Discussion paper, Colorado State University.

MADESTAM, A., D. SHOAG, S. VEUGER, AND D. YANAGIZAWA-DROTT (2013): "Do Political Protests Matter? Evidence from the Tea Party Movement," *The Quarterly Journal of Economics*.

MAESTAS, N., K. J. MULLEN, AND A. STRAND (2013): "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt," *American Economic Review*.

MANUEL, K. M., AND L. E. HALCHIN (2013): "Environmental Considerations in Federal Procurement: An Overview of the Legal Authorities and Their Implementation," Discussion paper, Congressional Research Service.

MARION, J. (2007): "Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions," *Journal of Public Economics*.

MCAFEE, R. P., D. C. QUAN, AND D. R. VINCENT (2002): "How to Set Minimum Acceptable Bids, with an Application to Real Estate Auctions," *Journal of Industrial Economics*.

MCAFEE, R. P., AND D. VINCENT (1997): "Sequentially Optimal Auctions," *Games and Economic Behavior*.

OWEN, A. L., E. CONOVER, J. VIDERAS, AND S. WU (2012): "Heat Waves, Droughts, and Preferences for Environmental Policy," *Journal of Policy Analysis and Management*.

PAARSCH, H. J. (1997): "Deriving an estimate of the optimal reserve price: An application to British Columbian timber sales," *Journal of Econometrics*.

PELTZMAN, S. (1976): "Towards a More General Theory of Regulation," *Journal of Law and Economics*.

PORTER, R. (1995): "The Role of Information in U.S. Offshore Oil and Gas Lease Auctions," *Econometrica*.

RILEY, J., AND W. SAMUELSON (1981): "Optimal Auctions," *American Economic Review*.

ROBERTS, J., AND A. SWEETING (2010): "Entry and Selection in Auctions," *Working Paper, Duke University*.

——— (2013): "Bailouts and the Preservation of Competition," *NBER Working Paper 16650*.

ROBERTS, J. W. (2013): "Unobserved Heterogeneity and Reserve Prices in Auctions," *RAND Journal of Economics*.

RYAN, S. (2012): "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*.

SAMUELSON, W. F. (1985): "Competitive Bidding with Entry Costs," *Economics Letters*.

SILVERMAN, B. (1986): *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London.

SKRETA, V. (2013): "Optimal Auction Design Under Non-Commitment," *Working Paper*.

SMITH, V. K., R. H. VON HAEFEN, AND W. ZHU (1999): "Do Environmental Regulations Increase Construction Costs for Federal-Aid Highways? A Statistical Experiment," *Journal of Transportation and Statistics*.

SOMAINI, P. (2014): "Competition and Interdependent Costs in Highway Procurement," *Working Paper*.

STEPHENS-DAVIDOWITZ, S. (2013): "The Effects of Racial Animus on Voting: Evidence Using Google Search Data," *Mimeo, Harvard University*.

STIGLER, G. J. (1971): "The Theory of Economic Regulation," *Bell Journal of Economics and Management Science*.

UNITED STATES DEPARTMENT OF DEFENSE (2008): "Department of Defense Green Procurement Program Strategy," Discussion paper, .

VARTIAINEN, H. (2011): "Auction Design without Commitment," *Journal of the European Economic Association*.

XU, P. (2013): "Nonparametric Estimation of Entry Cost in First-Price Procurement Auctions," *Journal of Applied Econometrics*.