

A Computational Framework for Analyzing Dynamic Auctions: The Market Impact of Information Sharing*

John Asker, Chaim Fershtman, Jihye Jeon, and Ariel Pakes
(Columbia, Tel-Aviv, Boston, and Harvard Universities)

September 23, 2019

Abstract

This paper develops a computational framework to analyze dynamic auctions and uses it to investigate the impact of information sharing among bidders. We show that allowing for the dynamics implicit in many auction environments enables the emergence of equilibrium states that can only be reached when firms are responding to dynamic incentives. The impact of information sharing depends on the extent of dynamics and provides support for the claim that information sharing, even of strategically important data, need not be welfare reducing. Our methodological contribution is to show how to adapt the experience-based equilibrium concept to a dynamic auction environment and to provide an implementable boundary-consistency condition that mitigates the extent of multiple equilibria.

Keywords: *Experience-Based Equilibria, Dynamic Procurement Auctions, Information Sharing.*

*JEL Codes: D43, K21, L41, C63, C73. We would like to thank numerous seminar audiences for their comments and questions. We are particularly grateful to Gautam Gowrisankaran and Mark Satterthwaite for extensive comments. El Hadi Caoui and Chuqing Jin provided excellent research assistance. Financial Assistance from the US-Israel Binational Science Foundation is greatly appreciated. Contact details: John Asker: johnasker@econ.ucla.edu; Chaim Fershtman: fersht@post.tau.ac.il; Jihye Jeon: jjeon@bu.edu; and Ariel Pakes: apakes@fas.harvard.edu. The usual caveat applies.

1 Introduction

This paper develops a computational framework to analyze dynamic auctions and then applies it to illustrate the possible implications of different rules for information exchange in that setting.

Dynamic auctions are sequential auctions in which the state of the bidders, and therefore their evaluation of the good that is auctioned, changes endogenously depending on the outcomes of prior auctions. The value of winning an auction to produce aircraft or ships depends on the backlog or the order book of the firm. Similarly, the value of winning a highway repair project or a timber auction depends on whether the inputs currently under the control of the firm are fully committed for the following period. The fact that the auction is dynamic implies a rich set of strategic incentives. For example, a firm may choose to allow a competitor’s state to transition to a point where that competitor bids less aggressively in order to win a subsequent auction at a lower bid.

A central feature of this environment is that competing firms may not have complete information about each other’s state variables, at every point in time. Empirically, this information asymmetry seems an important feature of many industries. Indeed, the fact that some firms are observed to make an effort to share information, at times illegally, underscores this general point (see the discussion below of the antitrust treatment of information sharing in the US and EU).¹

We provide a framework for analyzing dynamic auctions that allows for serially correlated asymmetric information, which implies that a competitor’s prior bids are signals of his current states. We use the framework to examine how the extent of information sharing impacts competition in a dynamic sequence of procurement auctions.² The analysis sheds light on the extent to which dynamic considerations can color the way antitrust regulators, procurement agencies, and other policy agencies approach the regulation of information sharing. The specific model we investigate is loosely based on the description of timber auctions in Baldwin, Marshall, and Richard (1997), although, to keep the model simple, we make many departures from the precise institutional features described therein. Having this specific empirical example in mind eases much of the exposition.

In each period, two firms can bid for the right to harvest a lot of timber in a first-price sealed-bid auction. Each firm has a stock of timber that it already has the right to harvest (its inventory). This stock is private information, and its evolution is the source of dynamics. To compete in the auction, firms must pay a participation fee and

¹The account of timber auctions in Baldwin, Marshall, and Richard (1997) motivates many of our modeling choices. Baldwin, Marshall, and Richard note several dimensions of private information, ranging from production costs (called “overrun” in the industry), to the amount of timber in a lot (particularly in old-growth forests) and the harvesting outcomes realized on timber lots located on private lands.

²The closest model to ours is the one estimated in the innovative contribution of Jofre-Bonet and Pesendorfer (2003). This framework is further extended in Groeger (2014), Saini (2013), Balat (2015), and Jeziorski and Krasnokutskaya (2016). Jofre-Bonet and Pesendorfer’s model (and those that follow) has private information that is conditionally independent across states. That is, conditional on (observed) state variables, knowing the private information of a rival last period provides no information as to the private information of the rival this period, which is not the case in our model. This means that the competitor’s prior period bid is a signal on its current state, and that information sharing has persistent value across periods.

submit a bid simultaneously. A firm may also choose to not participate. The winner of the auction, if any, receives the right to harvest the lot, and discovers how much harvestable material it contains. A harvest with a random outcome then occurs, which depletes the stock of timber each firm has in inventory.

Our benchmark model has full revelation of the state variable every T periods. That is, each firm observes the stock of unharvested timber of its competitor every T periods. Information sharing is modeled as shrinking the time interval between full-revelation periods so that we can investigate the possible implications of different rules of information exchange. We also investigate a model in which firms decide whether to share information. Voluntary information sharing involves firms making a choice every T periods as to whether to reveal in every period for the next T periods. For voluntary information sharing to occur over the next T periods, all firms must want to share information. Finally, we compare the results from these models with those we obtain from a model with myopic firms.

The numerical analysis of this game illustrates how information sharing can affect bidding behavior at a given state by increasing the precision of the firm's perceptions about its competitors' states. This, in turn, shapes the desirability, and therefore the likelihood of being in different states. An important point to bear in mind is that, conditional on the information they have, firms bid to maximize the expected NPV of their individual profits, rather than industry profits, in the model. On one hand, an increase in information increases the intensity of bidding and decreases profits in most (but not all) states. On the other hand, because firms have more precise information about when their competitor will be more aggressive, they are able to spend a greater fraction of the time in states where bidding is less aggressive. These states are the ones in which both firms' inventory is higher. The net effect is that information sharing leads to an increase average profits as well as an increase in the total sales of the auctioned timber.

Through this channel increasing information increases the value of firms. However, in our voluntary information-exchange game, firms have difficulty committing to exchange information and most often choose not to share. In addition, we find that in a model with myopic firms, the extent of information sharing has negligible effects.

This paper also has a methodological contribution. A framework for analyzing dynamic auctions in which a competitor's past behavior has a direct effect on a firm's perceptions about its competitor's likely action must allow for serially correlated asymmetric information. Fershtman and Pakes (2012) considered the numerical analysis of dynamic games with serially correlated asymmetric information, and we provide the modification required to use it in order to analyze dynamic auctions. Perhaps more importantly, we extend their notion of a restricted experience-based equilibrium by adding a consistency requirement on the boundary of the recurrent class of states—an extension that is possible to use in all dynamic games with asymmetric information. The boundary-consistency condition refines the set of computable equilibria (or, equivalently, mitigates potential multiple equilibria issues), and we provide intuition for when and why it can be used. We also show how to compute and test for boundary-consistent equilibria.

The results comprise an example of what can happen when firms share information based on the computational output from one parameterization. To that extent, the results provide a form of a possibility result. Setting methodological contributions

aside, we feel that the nature of the possibility result is important, in that in our setting information exchange is essentially welfare-neutral, despite having a real impact on firms' conduct (as noted in the paragraphs above). As such, this example illustrates the conceptual issues that may need to be confronted, and the level of care needed, in policy work or antitrust enforcement in this area.

This paper is organized as follows. In the remainder of this section, we discuss the related literature and provide a brief review of the role of information sharing in antitrust policy. Section 2 describes our baseline model, and then the information-sharing and the voluntary-information-sharing variants of the model. In section 3, computational details are described. A reader not concerned with computational details can skip this section and proceed directly to section 4, which discusses the numerical analysis, focusing on the competitive impact of information sharing. Section 5 concludes.

Related Literature

Our paper is closely related to the literature on the numerical analysis of dynamic oligopolistic games that uses the Ericson and Pakes framework (1995; for a survey of this literature, see Doraszelski and Pakes, 2007). Recent applications of this methodology to questions related to antitrust policy include Besanko, Doraszelski, and Kryukov (2014) on predatory pricing, and Mermelstein, Nocke, Satterthwaite, and Whinston (2014) on mergers. Within this literature, the closest papers to ours are Saini (2013) and particularly Jeziorski and Krasnokutskaya (2016). Both papers apply the Markov perfect equilibrium concept to auction settings, exploring the optimal procurement policy given capacity-constrained suppliers and subcontracting, respectively.³

Jeziorski and Krasnokutskaya (2016) find that reducing the relevance of private information (through subcontracting) lowers information rents and profits. This result contrasts with our finding that although reducing private information decreases profits for a given state, it induces firms to spend more time in states with less intense bidding, thereby leading to an increase in average per-period profits. Our paper's setting differs from Jeziorski and Krasnokutskaya's in at least two ways. First, in Jeziorski and Krasnokutskaya, private information (auction participation costs and marginal costs) is independent across firms and time, whereas in our setting, it is correlated across firms and time and depends on the firm's own and competitors' actions. Second, we focus on the effect of reducing the amount of private information among firms, whereas Jeziorski and Krasnokutskaya focus on the effect of private information becoming "less important" as subcontracting allows firms to modify unfavorable cost draws and control future costs by mitigating backlog accumulation.

As noted, our paper differs from this literature in that our focus is on information asymmetry, as in Fershtman and Pakes (2012). While that paper focuses on capital-accumulation games, we consider a more complex structure where, because we are modeling an auction, the evolution of a firm's state depends not only on its own action (its bid), but also on the bids of its competitors. We also introduce and operationalize a boundary-consistency condition that narrows the set of computed equilibria and can be rationalized either by prior information or experimentation.

Within the auction literature, Maskin and Riley (2000) consider asymmetric auctions and show that sealed bidding tends to favor weaker bidders, while in an open

³Both these papers build on Jofre-Bonet and Pesendorfer (2003).

auction, the bidder with the highest value win. Athey, Levin, and Seira (2011) extend the framework to a repeated auction. They consider a theoretical model of a repeated auction and then use data on timber auctions to conduct an empirical analysis of the effect of the type of auction (open or sealed bid) on the firms' participation and bidding.⁴

Our paper also relates to the empirical literature on bidding collusion. The literature uses several approaches for examining whether an auction is competitive or collusive. See, for example, Porter and Zona (1993, 1999), Baldwin, Marshall, and Richard (1997), Pesendorfer (2000), Bajari and Ye (2003), and Asker (2010). Aoyagi (2003) considers collusion in a repeated auction when bidders are allowed to communicate with each other before each auction. In another paper, Athey and Bagwell (2008) consider collusion between competitors in a repeated homogenous-good-Bertrand market in which costs (types) are private information and evolve over time according to an exogenous Markov process. In contrast to the environment considered here, the evolution of costs (types) in that model is unaffected by the actions of any player. We do not model collusion explicitly in our model, but we do examine the effect of information exchange regarding the firms' inventories on the firms' participation and bidding behavior.

The policy implications of our paper relate also to the extensive literature on information sharing in oligopoly; see Clarke (1983), Gal-Or (1985, 1986), Shapiro (1986), and Kirby (1988). For a survey of this literature, see Kuhn and Vives (1995). More recent empirical work includes Doyle and Snyder (1999) and Luco (2017). More recent theoretical work involving dynamic oligopoly models includes Overgaard and Møllgaard (2008) and Kubitz and Woodward (2019).

Information exchange and antitrust policy

The application in this paper is to information sharing between bidders, in which bidders share information as to their state. Hence, we consider the sharing of strategically valuable information as distinct from an explicit price-fixing or bid-rigging agreement. Though explicit agreements to fix prices are *per se* violations of the antitrust laws, the legal treatment of information sharing among competitors is less clear.⁵ The legality of an exchange of price information is determined in part by the extent to which the audience is restricted. Clearly, a merchant who posts prices in a public display is communicating price information to competitors but is not in violation of the law. More problematic is the communication of price information between competitors in a way that consumers do not have access to.⁶ U.S. courts apply a rule of reason test to decide whether the exchange of price information constitutes an unreasonable restraint of

⁴Bonatti, Cisternas, and Toikka (2017) have a related paper that considers the evolution of a dynamic Cournot game.

⁵The canonical statement of the *per se* nature of price fixing under section 1 of the U.S. Sherman Act is *United States v. Socony-Vacuum Oil* 310 U.S. 150 (1940). Information sharing also tends to fall within the scope of section 1 of the Sherman Act. See the majority decision in *United States v. Container Corp.* 393 U.S. 333 (1969).

⁶In *Container Corp.*, the U.S. Supreme Court held that, despite any agreement on pricing, the exchange of information about specific prices offered to specific customers was a violation of the antitrust laws. This case created confusion as to whether *per se* treatment applied to information sharing. This confusion was clarified in *United States v. Citizens & Southern National Bank* 422 U.S. 86., which explicitly adopted a rule of reason approach. In doing so, the court appealed to the idea that price exchange facilitated price

trade.⁷ Factors that are taken into account include the level of market concentration, the fungibility of the products, the nature of the information exchanged, its timeliness and specificity, and whether the information is made publicly available.⁸

U.S. courts take a sympathetic view of the sharing of non-price information, recognizing that efficiencies are more likely from the sharing of information regarding production processes and costs. For instance, the Supreme court in the 1925 *Maple Flooring Manufacturers* decision held that:

“... corporations which openly and fairly gather and disseminate information as to the cost of their product, the volume of production, ..., stocks of merchandise on hand, ... without however reaching or attempting to reach any agreement or any concerted action with respect to prices or production or restraining competition do not thereby engage in unlawful restraint of commerce...”⁹

Contemporary guidance from the FTC and DoJ states, “The sharing of information relating to price, cost, output, customers, or strategic planning is more likely to be of competitive concern than the sharing of less competitively sensitive information.”¹⁰ This suggests a somewhat more nuanced view in modern times. The EU, by contrast, has tended to take a harsher view of both price and non-price information-sharing agreements. The exchange of information relating to future prices is considered a restriction of competition by object (equivalent to a *per se* offense in the U.S.).¹¹ This may include non-price strategic information.

Our application illustrates that a harsh (*per se*) approach to the sharing of information can be misguided, in that bidders can engage in information sharing that is welfare neutral. As a result, there is no welfare-based justification for enforcement in this setting.

2 A Model of a Dynamic Auction

We consider a model with n firms in the market and no entry into or exit from the industry. Each firm can harvest and sell a portion of their stock of lumber at a fixed price in each period. The actual quantity that can be sold in each period depends on a firm-specific random outcome of a harvesting process from a stock of timber that has not yet been harvested, and is private information. The stock will be increased if the firm wins a procurement auction, which occurs every period. The procurement auction is a simple first-price sealed-bid auction. Participation in the procurement auction is costly, and participation decisions are public information observed by all

stabilization (a form of price fixing).

⁷In this context, an unreasonable restraint would be one that synthesizes or facilitates a cartel-like pricing structure. Information exchange may also constitute a facilitating practice in inferring the existence of an explicit price-fixing conspiracy.

⁸A modern discussion of the judicial approach taken can be seen in the decision of Justice Satomayor, while sitting as a judge on the Second Circuit Court of Appeals, in *Todd v. Exxon Corp* 275 F.3d 191 (2001).

⁹See *Maple Flooring Manufacturers’ Assn. v. United States* 268 U.S. 563 (1925).

¹⁰See FTC/DoJ’s April 2000 *Antitrust Guidelines for Collaborations Among Competitors*, page 15.

¹¹See the EU 2011 *Guidelines on the applicability of Article 101 of the Treaty on the Functioning of the European Union to horizontal co-operation agreements* and *Dole Food Company et al. v. Commission*.

firms. However, the amount of lumber per lot won in the auction is random and observed only by the winning firm.

There are two types of periods: periods with full information exchange and periods without information sharing. In our baseline model, full information exchange occurs every T periods. This assumption has a number of possible rationales and keeps the information set finite.¹²

We begin by describing the timing of the events that occur within a period. Then, we describe the overall structure of the game. Next, we define the equilibrium conditions, explain our computational procedure, and then provide and compare results from models with different amounts of information sharing.

Timing

1. Each firm brings into the period a stock of timber that can be harvested ($\omega_{i,t}$).
2. Every period begins with the announcement of a first-price sealed-bid auction.
3. Firms observe the realization of their stochastic participation fee. We assume $F_{it} \sim U[F_l, F_h]$. The realization is not observed by rival firms.
4. Each firm decides whether to participate in the auction. All the firms that decide to participate submit their bids simultaneously. At the time of bidding, participation decisions of rival firms are not observable.
5. The rules of the auction define an increment \underline{b} . Bids must be multiples of this increment. Hence, bids must be elements of the set $\mathcal{B} \equiv \{\underline{b}, 2\underline{b}, 3\underline{b}, \dots, \bar{b}\}$.
6. The highest bid wins. If high bids are tied, then the winner is decided randomly, with each tied bid having an equal chance of winning. We denote the probability of winning by firm i by $p^w(b, b_{-i})$, where b is the firm's bid and b_{-i} are competing bids. The winning bid, the identity of the winner, and the participants in the auction become public information.
7. If information exchange occurs, it does so at this point. If it is a period of information exchange (which occurs every T periods), then $\omega_{i,t}$ of all the firms is revealed. Otherwise, the new public information revealed in the period is; who participated in the auction, denoted as \mathbf{p}_t (a vector), who won the auction at period t , denoted by i_t^* , and the winning bid b_t^* . We denote the new public information as $\xi_t^n \equiv [i_t^*, b_t^*, \mathbf{p}_t]$. In a period of information exchange, the new public information is $[i_t^*, \boldsymbol{\omega}_t]$, the identity of the firm that won the auction and the observed state $\boldsymbol{\omega}_t \equiv \{\omega_{i,t}\}$.¹³

¹²See Fershtman and Pakes (2012) for a list of ways to keep the information set finite. Information revelation every T periods is convenient for us, because it allows us to directly compare equilibria based on variation in T . We can justify our structure by assuming that a regulator imposes mandatory periodic information revelation, or assuming the existence of a trade group that facilitates the sharing of information every T periods. In an empirical setting, the issue that would arise that is analogous to setting T here is how to define the state space. One way to do this is to let the data guide the selection by asking which candidate elements of the state space best explain observed actions. If, after selecting the observed states, serial correlation exists in the residuals, one might want to allow, in addition, a serially correlated unobservable. An example of this data-based approach to state-space specification in dynamic models can be found in Blundell, Gowrisankaran, and Langer (2018).

¹³Note that in a period of information revelation, the winning bid and the participation decision of that

8. The winner discovers the amount of timber on the plot it won. This amount is given by $\theta + \eta_t$, where θ is the average amount and η_t is an i.i.d. (across time) discrete random variable. η_t is not observed by the competing (losing) firms. The timber in stock ($\omega_{i,t}$) is updated accordingly. There is a random realization of the ability to extract, $e + \epsilon_{i,t}$, where $\epsilon_{i,t}$ is a discrete random variable with probabilities $p(\epsilon_{i,t})$. The draws on $\epsilon_{i,t}$ are independent over firms and not observed by competitors.¹⁴
9. Harvest is made and each firm sells all its harvested timber at a unit price of \$1. Thus, a firm's per-period revenue is given by $\min\{\omega_i + \mathbb{I}_{\{i=i^*\}}(\theta + \eta), e + \epsilon_i\}$, where $\mathbb{I}_{\{i=i^*\}}$ is an indicator function that takes the value of 1 if i wins the auction, and zero otherwise.¹⁵ The quantity harvested by firm i is not observable by other firms.¹⁶
10. Note that if $b = \emptyset$ signifies no participation, the above implies that at the time of bidding, the firm knows that its profits will be

$$\pi(b, F_i, \omega_i, \eta, \epsilon_i) =$$

$$\left[\mathbb{I}_{\{i=i^*\}} \left(\min\{\omega_i + (\theta + \eta), e + \epsilon_i\} - b_i \right) + (1 - \mathbb{I}_{\{i=i^*\}}) \min\{\omega_i, e + \epsilon_i\} \right] - \{b \neq \emptyset\} F_i,$$

and that $\mathbb{I}_{\{i=i^*\}}$ will depend on b_{-i} , the competitors' bids. The minimum reflects the fact that the amount sold cannot be larger than either the amount in stock (which is $\omega_i + (\theta + \eta)$ if the firm won the auction or ω_i if not) or the amount harvested ($e + \epsilon_i$).

11. After profits are realized, all firms update their private ω_i .

Firms' Strategy Space

In general, the strategy space could include everything observed from the history of the game. Most of the early applied literature focused on equilibria with strategies that depend only on variables that are either "payoff" or "informationally" relevant. The payoff-relevant variables are defined, as in Ericson and Pakes (1995) or Maskin and Tirole (2001), to be those variables that are not current controls and affect the current profits of at least one of the firms. In a game with asymmetric information, observable variables that are not payoff relevant will affect behavior if they are informationally relevant. A variable is informationally relevant if and only if some player can improve its expected discounted value of net cash flows by conditioning on the variable, even if no firms' strategy depended on it; for more details see, Fershtman and Pakes (2012). That paper also shows that in models with periodic revelation of information there exists

period do not enter the public information because they are payoff and informationally irrelevant. They do not provide any additional signal on the ω of the firms, because these ω 's are revealed in that period.

¹⁴The game form could be simplified by collapsing these η_t and $\epsilon_{i,t}$ into one random variable (most simply by removing η_t altogether). Such a simplification could maintain private information about the evolution of each firm's ω . By contrast, the specification we adopt stays close to the motivating timber-auction example at the expense of a little more complication. It also injects additional informational asymmetry at the point at which a firm wins an auction. This specification is broadly consistent with the descriptions of timber auctions in Baldwin, Marshall and Richard (1997).

¹⁵Here, and in what follows, we drop time subscripts, except where they add clarity.

¹⁶Otherwise, the observable harvested quantity may serve as a signal regarding ω_i .

an equilibrium that only conditions on the revealed information and the information that has become available since the revelation. We focus on this equilibrium in the remainder of the paper.¹⁷

The information set of firm i consists of public and private information. The public information at the beginning of period t , denoted by ξ_t consists of: $\tau_t \in [1, \dots, T]$, the time since the last information exchange, $\omega_{t-\tau_t}$, the last revealed ω vectors, and the τ_t -period history of winning bids, winner identities, and participant identities. Formally, $\xi_t = \{\tau_t, \omega_{t-\tau_t}, \xi_{t-1}^n, \dots, \xi_{t-\tau_t}^n\}$.¹⁸ Information revelation occurs when $\tau_t = T$ (which is period $\tau_t = 0$ for the next cycle). The private information at the point in time decisions are made includes $\omega_{i,t}$ and $F_{i,t}$. However, since $F_{i,t}$ is i.i.d. and enters the value function linearly, it does not have an independent effect on future values, whereas the other state variables do. As a result, continuation values are determined by $J_{i,t} = (\omega_{i,t}, \xi_t)$, but both $J_{i,t}$ and $F_{i,t}$ are needed to determine bids.

Strategies

A firm's strategy has two elements: the participation strategy and the bidding strategy. We denote firm i strategy as $b(J_i, F_i) \rightarrow \{\mathcal{B} \cup \emptyset\}$, where $b = \emptyset$ signifies no participation. The subscript i is added only where doing so helps differentiate between the firm's b_i (equivalently, b) and competing bids, b_{-i} .¹⁹

2.1 The Value Function

We let $V(J_i, F_i)$ be the value of the game for player i given information (J_i, F_i) . We have

$$V(J_i, F_i) = \max \left\{ W(\emptyset|J_i), \max_{b \in \mathcal{B}} [W(b|J_i) - F_i] \right\}, \quad (1)$$

where (i) $W(\emptyset|J_i)$ is the value of the game if the firm decides not to participate in the auction in that period, and (ii) $W(b|J_i)$ is the value when the firm participates and bids $b \in \mathcal{B}$.

Now consider the value of the game when firm i participates in the auction and bids $b \in \mathcal{B}$. For every possible J_i , we define $p^w(b|J_i)$ to be the player's perception about the probability of winning the auction when it bids b . Recall that i^* is the winning firm,

¹⁷The equilibrium is defined formally in section 2.2.

¹⁸Note that for a period with information revelation, the public information includes only the identity of the winner in the auction and not the winning bid or the participant identities, as these variables are not informationally relevant.

¹⁹If a static version of our game is considered (for example, by imposing that players are myopic), the existence of a pure-strategy equilibrium is somewhat unclear. Bidders in our setting have a "type" that consists of ω_i and F_i : ω_i is discrete, while F_i is continuous. Importantly, in our model players do not observe the entry cost F_{-i} of their rivals or the rivals' inventory state ω_{-i} . (Note that even in a period with information exchange, the players observe only the inventory level from the beginning of the period after their participation and bidding decision have been done.) Depending on the parametrization, F_i and ω_i may introduce enough "mixing" to ensure the existence of pure strategy equilibria. It is clear, however, that as the support of F_i approaches a point mass, or as the grid of ω_i 's becomes sufficiently coarse relative to the bid grid, problems with the existence of pure strategy equilibria will arise. In our numerical analysis in section 4, we choose the bid grid to be $\{0.5, 1, 1.5, 2\}$. Restricting the numbers of bids firms can facilitate computation, since the algorithm requires updating and saving the continuation value for each possible bid in every state that players visit. See section 3.1 for more detail on the computation algorithm.

so $p(\xi'|\xi, \omega_i, b, i = i^*)$ is the firm's perception of the probability distribution of future public information given its current public (ξ) and private (ω_i) information should it win the auction with bid b . Analogously, $p(\xi'|\xi, \omega_i, b, i \neq i^*)$ is the firm's perception of the probability distribution of future public information, should the firm lose the auction. The equilibrium assumption below will restrict these perceptions.

Recall that the firm receives revenue from selling its harvested lumber, where the amount harvested is the minimum of the extraction ability of the firm ($e + \epsilon_i$) and the firm's actual stock ($\omega_i + (\theta + \eta)$ if it wins, and ω_i otherwise). Thus, the firm's expectation of current-period revenue (which excludes F_i) is

$$\pi^e(b|J_i) = \sum_{\epsilon_i, \eta} \left[p^w(b|J_i) \left(\min\{\omega_i + \theta + \eta, e + \epsilon_i\} - b \right) + [1 - p^w(b|J_i)] \min\{\omega_i, e + \epsilon_i\} \right] p(\epsilon_i) p(\eta). \quad (2)$$

Letting β be the discount factor, the firm's expectation of its continuation value conditional on bidding $b \in \mathcal{B}$ is

$$\begin{aligned} W(b|J_i) &= \pi^e(b|J_i) + \\ & p^w(b|J_i) \beta \sum_{\epsilon_i, \eta, \xi', F'_i} V(\omega'(\omega, \eta, \epsilon_i), \xi', F'_i) p(\xi'|\xi, \omega_i, b, i = i^*) p(F'_i) p(\eta) p(\epsilon_i) \\ & + (1 - p^w(b|J_i)) \beta \sum_{\epsilon_i, \xi', F'_i} V(\omega'(\omega, \epsilon_i), \xi', F'_i) p(\xi'|\xi, \omega_i, b, i \neq i^*) p(F'_i) p(\epsilon_i), \end{aligned} \quad (3)$$

where $\omega'(\omega, \eta, \epsilon_i)$ is the updated ω_i when the firm does win the auction, and is a function of the random outcomes of the size of the lot won (η) and the harvesting decision (ϵ_i); that is, $\omega'(\omega, \eta, \epsilon_i) = \max\{0, \omega_i - (e + \epsilon_i) + \theta + \eta\}$. When the firm does not win the auction, its updated ω is a function of the initial ω and the random outcome of the harvesting process, ϵ_i ; that is, $\omega'(\omega_i, \epsilon_i) = \max\{0, \omega_i - (e + \epsilon_i)\}$. Note that the continuation value when a firm does not participate in the auction, or $W(\emptyset|J_i)$, is obtained by setting $p^w(\emptyset|J_i) = 0$ in equations (2) and (3).

2.2 Restricted Experience-Based Equilibrium

We now derive for this game the conditions of a restricted experience-based equilibrium (a REBE), as defined in Fershtman and Pakes (2012). This derivation requires us to define three objects and then detail the restrictions that the equilibrium notion places on the values of those objects. To ease the notation in the definition below, we let \mathbf{s} be the set consisting of the payoff- and informationally relevant states of all the firms. That is, $\mathbf{s} = (J_1, \dots, J_n)$, where all the J_i have the same public component ξ . So $\mathbf{s} = (\omega_1, \dots, \omega_n, \xi)$. We will say that $J_i = (\omega_i, \xi)$ is a component of \mathbf{s} if it contains the information set of one of the firms whose information is included in \mathbf{s} . We define the set of possible states as $\mathcal{S} = \{\mathbf{s} : (\omega_1, \dots, \omega_n) \in \Omega^n(\omega), \xi \in \Omega(\xi)\}$.

Definition of a REBE: A restricted experience-based equilibrium consists of

1. A set \mathcal{R} that is a subset of the state space (i.e. $\mathcal{R} \subset \mathcal{S}$).
2. Bidding and participation strategies, $b^*(J_i, F_i)$, for all $F_i \in [F_l, F_h]$ and for every J_i that is a component of any $\mathbf{s} \in \mathcal{S}$.

3. A set of numbers $\mathcal{W} \equiv \{W^*(b|J_i)_{b \in \mathcal{B} \cup \emptyset}\}$, for every J_i that is a component of any $\mathbf{s} \in \mathcal{S}$, that have an interpretation as the firm's perceptions of the expected discounted values of current and future cash flows conditional on its information set should it bid b or not participate in the auction (i.e. where $b = \emptyset$).

For these objects to define a REBE, they must satisfy the following three conditions.

C1: \mathcal{R} is a recurrent class. The Markov process generated by any initial condition $\mathbf{s}_0 \in \mathcal{R}$, and the transition kernel generated by $\{b^*(J_i, F_i)\}_{J_i \in \mathcal{S}, F_i \in [F_l, F_h]}$ has \mathcal{R} as a recurrent class; that is, with probability 1, any subgame starting from an $\mathbf{s}_0 \in \mathcal{R}$ will generate sample paths that are within \mathcal{R} forever.

C2: Optimality of strategies. Conditional on $\mathcal{W} \equiv \{W^*(b|J_i)_{b \in \mathcal{B} \cup \emptyset}\}_{J_i \in \mathcal{S}}$, the strategies are optimal. That is,

$$b^*(J_i, F_i) = \arg \max_{b \in \mathcal{B} \cup \emptyset} [W^*(b|J_i) - \{b \neq \emptyset\} F_i].$$

C3: Consistency of values on \mathcal{R} . Consistency requires that the perception of discounted values, generated by every possible choice at every J_i that is a component of an $\mathbf{s} \in \mathcal{R}$, equals the expected discounted value of returns generated by that choice from that J_i , where expectations are taken using the distribution of the outcomes from that J_i generated by the policies in C2. This distribution is the empirical distribution of outcomes that the firm would actually observe given equilibrium play, so we denote them by $\mu^E(\cdot|\cdot)$. Formally, for every $b \in \mathcal{B} \cup \emptyset$, $W^*(b|J_i)$, the equilibrium evaluations satisfy

$$W^*(b|J_i) = \pi^E(b|J_i) + \beta \sum_{\epsilon, \eta, F_i} V(J'_i, F_i) \mu^E(J'_i|b, J_i) p(\eta) p(\epsilon) p(F_i), \quad (4)$$

where J'_i is the realization of the following period's $J_i \equiv (\omega_i, \xi)$,

$$\pi^E(b|J_i) = \sum_{\epsilon_i, \eta} \left[\mu_w^E(b|J_i) \left(\min\{\omega_i + \theta + \eta, e + \epsilon_i\} - b \right) + [1 - \mu_w^E(b|J_i)] \min\{\omega_i, e + \epsilon_i\} \right] p(\epsilon_i) p(\eta),$$

$\mu_w^E(b|J_i)$ is the empirical probability of winning if the firm bids b at J_i , or

$$\mu_w^E(b|J_i) = \sum_{J_{-i}, F_{-i}} Pr(i = i^* | b, b_{-i}^*(J_{-i}, F_{-i})) \mu^E(J_{-i} | J_i) p(F_{-i}),$$

$\mu^E(J'_i|b, J_i)$ is the empirical probability of transiting from J_i to J'_i if the firm plays b

$$\mu^E(J'_i|b, J_i) = \sum_{J_{-i}} \mu^E(J'_i|b, J_i, J_{-i}) \mu^E(J_{-i} | J_i)$$

where $\mu^E(\tilde{J}_{-i} | \tilde{J}_i)$ is the empirical probability of the competitors' state \tilde{J}_{-i} conditional on the firm's own state \tilde{J}_i given by

$$\mu^E(\tilde{J}_{-i} | \tilde{J}_i) = \frac{\mu^E(\tilde{J}_{-i}, \tilde{J}_i)}{\mu^E(\tilde{J}_i)} \equiv \frac{\sum_{t=0}^{\infty} \mathbb{I}\{(J_{-i,t}, J_{i,t}) = (\tilde{J}_{-i}, \tilde{J}_i)\}}{\sum_{t=0}^{\infty} \mathbb{I}\{J_{i,t} = \tilde{J}_i\}},$$

and $\mu^E(J'_i|b, J_i, J_{-i})$ is the empirical probability of J'_i given (b, J_i, J_{-i}) , which is obtained differently in revelation and no-revelation periods. In a no-revelation period, the new information revealed is $\xi_t^n \equiv [i_t^*, b_t^*, \mathbf{p}_t]$ and

$$\mu^E(J'_i|b, J_i, J_{-i}) = \sum_{F_{-i}} \Pr(\omega'_i, \xi_t^n | b, b_{-i}^*(J_{-i}, F_{-i}), J_i, J_{-i}) p(F_{-i}),$$

whereas in a revelation period ω_t is revealed so $J'_i = (\omega'_i, \omega_t, i_t^*)$ and

$$\mu^E(J'_i|b, J_i, J_{-i}) = \sum_{F_{-i}} \Pr(\omega'_i, \omega_t, i_t^* | b, b_{-i}^*(J_{-i}, F_{-i}), J_i, J_{-i}) p(F_{-i}). \spadesuit$$

Two aspects of this definition influence the discussion that follows. First, when a bid b results in a transition to a point outside the recurrent class (which implies $b \neq b^*(J_i)$), $\mu^E(J'_i|b, J_i, J_{-i})$ can be computed, but $V(J'_i, F_i)$ is not constrained by condition C3, as C3 only applies to points in the recurrent class. We come back to this point in section 2.3 where we introduce the notion of boundary consistency. Second, Fershtman and Pakes (2012) deal with a capital-accumulation game where the distribution of ω' depended only on the firm's own policy, J_i , and the primitives of the problem. Then, the distribution of J'_i given (J_i, b) for $b \neq b^*(J_i)$ can be computed from the data generated by the equilibrium and knowledge of the primitives. Since the bid of a firm's competitors affects the distribution of ω'_i in an auction, our problem is not a capital-accumulation game. In this non-capital-accumulation game, the only information that is in the data is the distribution of J'_i given $(J_i, b^*(J_i))$, and some experimentation may be required for the agent to learn the distribution of J'_i given $(J_i, b \neq b^*(J_i))$. We come back to in this point in our discussion of computation in section 3.

It is helpful to clarify the conceptual and computational differences between a REBE and a Markov perfect Bayes equilibrium (an MPBE).²⁰ If we were to use an MPBE, we would have to define the players' beliefs about the types of their competitors (i.e., their current ω_{-i}) for every possible information set. In addition, we would need to formalize how those beliefs are updated and ensure their consistency with the equilibrium strategies. This would be a formidable task in a setup with a large number of states.²¹ A REBE only requires firms to have perceptions regarding the expected returns from actions at states (the W 's), and those perceptions need only be consistent with outcomes at states that are actually part of equilibrium play. We conjecture that any set of outcomes supported by an MPBE will be able to be supported by a REBE but the reverse is not the case; a REBE requires less restrictive conditions than an MPBE.

²⁰As discussed in Fershtman and Pakes (2012), the self-confirming equilibrium of Fudenberg and Levine (1993) also has a somewhat similar flavor to REBE in various dimensions.

²¹In our numerical analysis with two firms, the number of states visited in computation for the baseline specification is greater than 7 million. Requiring beliefs to be derived to support an MPBE, even for just this subset of the state space, would require considering beliefs as to the probability of facing at least 10 different ω_{-i} 's for each state. At least for potential algorithms that we can envisage, this alone increases the computation by at least an order of magnitude (additionally, recall that an MPBE requires consistent beliefs in every point in the state space). Moreover, formulating a process to ensure these beliefs were consistent with Bayes rule, where possible, given the strategies, is something we understand to be beyond the scope of the current literature other than in very specific model formulations (see Bonatti et al. (2017) and Board and Meyer-Ter-Veyn (2018)).

Instead of requiring specifications for beliefs about the types of other players, a REBE assumes that players use their own experience (outcomes from previous periods) to form an estimate of expected returns from different possible actions. For example, in our environment each firm has experience in participating in the periodic auctions, and based on its experience forms a perception of the probability of winning with bid b at the information set J_i . At equilibrium these perceptions are given by the empirical distribution generated by the firm’s experience of winning at J_i . Were it not to rely on experience, the firm would need to have a perception regarding the probabilities of the types of their rivals, use the equilibrium notion to construct a distribution of the bids of other firms given their perceptions of the competitors’ likely types, and then use that distribution to construct the probability of winning for every possible bid b . Part of the REBE’s appeal is that the use of experience simplifies both the firm’s and the analyst’s problem. In a REBE, it is past experience that restricts (and, at least at an $s \in \mathcal{R}$, actually determines) the players’ perceptions about future values. In the MPBE, past experience plays no direct role.

Another difference between a REBE and an MPBE is a REBE fully specifies equilibrium conditions only for the recurrent class, or on \mathcal{R} . The reason is that a firm’s experience will only lead to accurate perceptions for states that are visited repeatedly, that is, for the states in \mathcal{R} . In the MPBE, the equilibrium specifies equilibrium conditions and beliefs for all possible states including those that are never visited as a result of equilibrium play. As explained in the next subsection, the beliefs about off-equilibrium play may be important determinants of the implications of the model. In that subsection we consider a condition (boundary consistency) that restricts perceptions of values at certain points outside of \mathcal{R} in a REBE in a way that we think would be appropriate for many applied settings.

Before leaving this subsection, we note an additional difference between the conditions satisfied by a REBE and the conditions satisfied by an MPBE. Continuation values are expectations, and more than one distribution function can lead to the same expectation. Even at points in \mathcal{R} a REBE does not require equilibrium perceptions of outcomes to be consistent with a Bayesian posterior on competitors’ actions. In a sense, this difference is a technical one, because it is the continuation values per se that determine actions, not the perceptions that lead to them. However, this difference enables one to compute a REBE using a learning algorithm that is much simpler than the methods likely required to compute an MPBE (see section 3 for details). As a result, a REBE might both provide a closer approximation to actual behavior and enable researchers to analyze a broader set of issues.

Appendix A computes and compares a REBE with an MPBE in a simple textbook-style model in which policies and values can be derived analytically. This simple example may be helpful to readers who want to familiarize themselves with REBE in a more familiar game form.

2.3 Strengthening REBE: Boundary Consistency

It will be helpful to distinguish between two (mutually exclusive and exhaustive) subsets of states in \mathcal{R} . If optimal policies are followed by all agents at any state in \mathcal{R} , then the transition from that state will be to a state in \mathcal{R} with a probability of 1. However, there may be some states in \mathcal{R} where if a firm chose a feasible, though not optimal, policy, the state might transition to a point not in \mathcal{R} . We refer to such states

as *boundary states*. We refer to states in \mathcal{R} that would only transition to other points in \mathcal{R} regardless of which policies are followed as *interior states*.

The perceptions at points outside of \mathcal{R} are not tied down by equilibrium condition C3. So the perceptions of what would happen were a feasible policy followed from a boundary point need not be consistent with what would actually happen were that policy followed.²² As a result, different perceptions of what would happen if a non-optimal feasible policy were chosen can lead to different recurrent classes. This is a source of multiplicity of the equilibria that can be generated by a REBE that is absent from the equilibria generated by an MPBE. This subsection considers selecting a subset of the REBE equilibria that can be generated in this way through using a restriction that is both likely to be appropriate for many (though not all) applied problems, and can relatively easily be used in analyzing them (see subsection 3.2.2).²³

If firms have prior knowledge or experiment with off-the-equilibrium path policies at boundary points, then we might expect off-the-equilibrium path behavior at boundary points to satisfy some restrictions. This section provides one such restriction: that the actual value of off-the-equilibrium-path play from a boundary point is equal to the perceived value of off-the-equilibrium-path play at those points (recall that condition C1 guarantees that these perceptions are defined on all of \mathcal{S}). We call this restriction a boundary-consistency condition, as it, together with condition C2, ensures that playing any feasible policy at a boundary point would generate a perceived discounted value that is not greater than that of the optimal policy. Note that to impose this condition, we need to only calculate discounted values for profits along sample paths before they re-enter the recurrent class (if they do re-enter) as we can use C3 above to evaluate the periods thereafter.

To formalize our condition, we need to define the set of actions that could be taken from points in the recurrent class that would generate outcomes that are not in the recurrent class. To this end, let $\text{supp}[p_{s'}(\cdot|b_i, b_{-i}^*, \mathbf{s})]$ be the support of the probability distribution over next-period states generated by actions (b_i, b_{-i}^*) and initial state $\mathbf{s} = (J_i, J_{-i})$. The boundary set of couples (b, \mathbf{s}) , which we denote by B , is the set of action-state combinations such that if (i) $\mathbf{s} = (J_i, J_{-i}) \in \mathcal{R}$, (ii) action b is taken by i , and (iii) equilibrium actions are taken by the other firms, then the support of the probability distribution for \mathbf{s}' has a point that is not in the recurrent class. Formally,

$$B \equiv \tag{5}$$

$$\{(b \in \mathcal{B} \cup \emptyset), (J, J_{-i}) = \mathbf{s} \in \mathcal{R}) : \exists F_{-i} \text{ s.t. } \text{supp}[p_{s'}(\cdot|b, b^*(J_{-i}, F_{-i}), s)] \cap (\mathbf{s}' \notin \mathcal{R}) \neq \emptyset\},$$

where \mathcal{B} is the set of bids and \emptyset is the empty set.

The additional condition that needs to be satisfied for the one-period deviation to actually yield an outcome with a value that is less than the value of the optimal play is stated formally in C4 below. In this condition, we use γ to index periods since the off-equilibrium-path policy is played.

²²To see this point, note that $V(J'_i, F_i) = \max\{W(\emptyset|J'_i), \max_{b \in \mathcal{B}}[W(b|J'_i) - F_i]\}$. C3 requires that $W^*(b|J'_i)$ be consistent with current-period profits and $V(J'_i, F_i)$, but does not require similar consistency for $W(b|J'_i)$ or $W(\emptyset|J'_i)$ if $J'_i \notin \mathcal{R}$.

²³Of course, in an empirical problem the data are likely to generate restrictions on these perceptions; see the discussion in Pakes (2016).

C4: Boundary Consistency. $\forall (b, J_i)$ component of $(b, \mathbf{s}) \in B$ (as defined in (5)) and for every F_i ,

$$W(b^*|J_i) - \{b^*(J_i, F_i) \neq \emptyset\}F_i \geq \sum_{J_{-i}, F_{-i}} \left[\pi(b, b_{-i}^*(F_{-i}, J_{-i}), J_i, F_i) + \sum_{\gamma=1}^{\infty} \beta^\gamma \sum_{\mathbf{s}_\gamma, \mathbf{F}_\gamma} \pi(b_i^*(J_i, F_i), b_{-i}^*(F_{-i}, J_{-i}), J_i, F_i) p(\mathbf{s}_\gamma | \mathbf{s}_{\gamma-1}, b^*, \mathbf{F}_\gamma) p(\mathbf{F}_\gamma) \right] p(F_{-i}) \mu^E(J_{-i}|J_i), \spadesuit$$

where $\pi(b, b_{-i}^*(F_{-i}, J_{-i}), J_i, F_i) \equiv \pi^E(b|J_i) - \{b \neq \emptyset\}F_i$, and $\pi(b_i^*(J_i, F_i), b_{-i}^*(F_{-i}, J_{-i}), J_i, F_i)$ is defined analogously. $\mathbf{F} = (F_i, F_{-i})$, and $p(\mathbf{s}_\gamma | \mathbf{s}_{\gamma-1}, b^*, \mathbf{F}_\gamma)$ is the probability of reaching state \mathbf{s}_γ at time γ given that at time $\gamma - 1$, the state is $\mathbf{s}_{\gamma-1}$, participation fees are \mathbf{F}_γ , and the players play the equilibrium strategies b^* .²⁴ Thus, the expression on the right side of the inequality in C4 represents the expected discounted values of current and future cash flows if the firm bids b in the current period and plays equilibrium strategies b^* in all subsequent periods; and its competitors play equilibrium strategies in the current and future periods.

Definition of a Boundary-Consistent REBE: We call an equilibrium that satisfies C1 to C4 a boundary-consistent REBE.

For any sample path (i.e. any $\{\mathbf{s}_\gamma\}_{\gamma=1}^{\infty}$), we define $\gamma_R = \min_\gamma \{\gamma : (\mathbf{s}_\gamma) \in \mathcal{R}\}$ as the number of periods from which the firm first re-enters the recurrent class after leaving it. Then, we can replace

$$\sum_{\gamma=\gamma_R}^{\infty} \beta^\gamma \sum_{\mathbf{s}_\gamma, F_i} \pi(b_i^*(J_i, F_i), b_{-i}^*(F_{-i}, J_{-i}), J_i, F_i) p(\mathbf{s}_\gamma | \mathbf{s}_{\gamma-1}, b^*, \mathbf{F}) p(F_i)$$

in C4 with $\beta^{\gamma_R} \sum_{F_i} V(\mathbf{s}_{\gamma_R}, F_i) p(F_i)$. We provide a formal test for the existence of boundary-consistent policies below. The fact that we can replace the infinite sum in C4 with $\beta^{\gamma_R} \sum_{F_i} V(\mathbf{s}_{\gamma_R}, F_i) p(F_i)$ eases the computational burden of the test.

As a final matter, this boundary-consistency notion can be further strengthened by extending it to points in the recurrent class that are one step (or more) removed from the set of boundary points B . In the appendix, an example is provided in which boundary consistency rules out a REBE, and this one-step extension of boundary consistency would rule out another. In the computational analysis of the dynamic auction modeled here, we found only one equilibrium that failed boundary consistency (see section 4.2), but, as explained in section 3, that finding is likely a result of tradeoffs we made in how we computed our equilibria.

2.4 Information Sharing

We study the role of information sharing between firms participating in a sequence of procurement auctions. In our benchmark case, information is shared every T periods. Between these periods, firms do not observe the evolution of their competitors' states; however, they do observe the public information, which may help in predicting their competitors' behavior. We then compare our baseline model with two models that allow

²⁴The probability distribution $p(\mathbf{s}_\gamma | \mathbf{s}, b^*, \{\mathbf{F}_\tau\}_{\tau=1}^{\gamma-1})$ is derived recursively, with $p(\mathbf{s}_1 | b^*, \mathbf{s}) = \sum_{F_{-i}} p(\mathbf{s}_1 | b_i, b_{-i}^*(J_{-i}, F_{-i}), \mathbf{s}) p(F_{-i})$, and for $\gamma > 1$, $p(\mathbf{s}_\gamma | b^*, \mathbf{s}_{\gamma-1}) = \sum_{\mathbf{F}} p(\mathbf{s}_\gamma | b^*, \mathbf{s}_{\gamma-1}, \mathbf{F}) p(\mathbf{F})$.

for information exchange at more frequent intervals. The only difference among them is the extent of information sharing as we do not allow for any additional mechanism that facilitates coordination among firms. We also assume that when information is exchanged firms reveal their true state.²⁵

2.4.1 Information Exchange (*IE*)

The first information-sharing model has a mandatory information exchange every period. Therefore, in the IE model, at the beginning of each period, a firm's private information includes its own current stock, $\omega_{i,t}$, and the public information includes the identity of the winner and the ω -vector from the last period: $[i_{t-1}^*, \omega_{t-1}]$. Shrinking the interval at which information gets revealed is straightforward to implement (both for the analyst and the firms) within our framework relative to alternative approaches of modeling information sharing (e.g. explicitly giving firms with more accurate beliefs about competitors' states, $\omega_{-i,t}$ in each period). Formally, we compute the model already described with the constraint that $T = 1$. We denote this model as *IE*.

2.4.2 Voluntary Information Exchange (*VIE*)

In the second information-sharing model, we adjust the baseline model such that in the period with a forced information exchange occurs, firms also make a decision on whether to share information in every period for the next $T - 1$ periods. If one of the firms does not wish to share information, no voluntary information sharing occurs over the next $T - 1$ periods, and in the T^{th} period firms choose whether they wish to share information in the subsequent $T - 1$ periods.²⁶ We call this model the *VIE* model and describe it in more detail now.

The period index, $\tau = 0, 1, \dots, T - 1$, designates the time from the period of mandatory information exchange, such that information exchange occurs at $\tau = 0 (\equiv T)$. At $\tau = 0$, each firm also needs to decide if it wishes to be part of an information-exchange scheme. The decision of whether to share information, $\tilde{R}_i \in \{0, 1\}$, is made simultaneously with the participation and bidding decision. $\tilde{R}_i = 1$ denotes that firm i wishes to share information. Information is actually exchanged, denoted by $R = 1$, only when $\tilde{R}_i = \tilde{R}_{-i} = 1$.

The timing of the game is adjusted so that the sequence described in section 2 changes as follows. If $\tau = 0$, step 4 is replaced with

Each firm finds its cost of participating in the auction ($F_{i,t}$). All the firms then simultaneously; i) decide whether to participate in the auction and if so submit a bid, and ii) decide whether to reveal information. If both firms agree to reveal information, there is information exchange over the next T periods and the vol-

²⁵Truthful revelation may require careful design of the incentives surrounding the agreement. To explore this area in the context of explicit cartels in auction markets, see, for example, Graham and Marshall (1987), McAfee and McMillian (1992), and Mailath and Zemsky (1991).

²⁶Requiring information sharing to last for $T-1$ periods (rather than, for instance, allowing a new choice to be made every period) means: (a) the state space does not expand unduly (if choices were made every period, then we would need to keep track of every outcome and choice in each period between forced revelations); and (b) the structure of the game deviates minimally from the baseline, easing exposition, comparisons to the baseline, and implementation.

untary information exchange state R is set to 1. R is 0 otherwise. At the time of bidding, participation decisions of rival firms are not observable.

For $\tau > 0$, we replace step 7 with

Information exchange occurs at this point. If $R = 1$, $\omega_{i,t}$ of all the firms is revealed in addition to the new public information (i.e. who won the auction). If $R = 0$, the new public information revealed in the period is the same as in the baseline model that is $\xi_t^n = [i_t^*, b_t^*, \mathbf{p}_t]$.

In the *VIE* game, the firms' information set is different from in the *B* game in that the public information also includes the most recent information-sharing indicator, or $R \in \{0, 1\}$. That is, for every $\tau > 0$ the public memory, ξ_t includes also the last information-revelation status $R \in \{0, 1\}$ as this indicates if, in the remaining $T - \tau$ periods, the firms will exchange information. On the other hand, we do not keep \tilde{R}_i or \tilde{R}_{-i} in memory, since they are not informational relevant due to the fact that in the period in which decisions on \tilde{R} are made there is forced information revelation.

The information-exchange decision: At periods when $\tau = 0$, firms need to decide if they wish to exchange information in the next T periods. In those periods, we let $\tilde{R} \in [0, 1]$ indicate the decision over whether to exchange information ($\tilde{R} = 1$) or not ($\tilde{R} = 0$). We define the value of the game for player i given information (J_i, F_i, \tilde{R}) as

$$\tilde{V}(J_i, F_i, \tilde{R}) = \max \left\{ \max_{b \in \mathcal{B}} (W(b, \tilde{R}|J_i) - F_i), W(\emptyset, \tilde{R}|J_i) \right\}$$

where $W(b, \tilde{R} = 1|J_i)$ and $W(b, \tilde{R} = 0|J_i)$ are the firm's perceptions of the expected discounted value of current and future cash flows, given the choice of bid and the choice to reveal information in the next T periods, conditional on the firm's information set. The firm submits $\tilde{R} = 1$ if and only if $\tilde{V}(J_i, F_i, \tilde{R} = 1) \geq \tilde{V}(J_i, F_i, \tilde{R} = 0)$. The actual exchange state, our R , has $R = 1$ if and only if $\tilde{R}_i = \tilde{R}_{-i} = 1$.

When $\tau = 0$, $W(b, \tilde{R} = 0, J_i)$ is analogous to $W(b, J_i)$ in equation (3). When $\tau = 0$ and $\tilde{R} = 1$, there is a probability of moving into different R states that depends on the perceptions of whether the competitor will choose to reveal. We let $p(R = 1|J_i, \tilde{R} = 1)$ be the firm's perception of that probability given $\tilde{R}_i = 1$ and J_i . We use this perception combined with equation (3) to form $W(b, \tilde{R}_i = 1|J_i)$. For $\tau > 0$, the dynamics are similar to the *B* case when $R = 0$, and are similar to the dynamics of the *IE* case when $R = 1$.

Definition of a REBE for the *VIE* case: The definition of a REBE for the *VIE* case is analogous to that for the baseline and *IE* cases but with the differences we now consider. In the *VIE* in periods with $\tau > 0$, the public information ξ includes the outcome of the last voluntary information exchange; that is, $R \in \{0, 1\}$. At $\tau = 0$, the optimal policies are given by

$$\begin{aligned} \tilde{R}^*(J_i, F_i) &= \arg \max_{\tilde{R} \in \{0, 1\}} \left[W(b^*(J_i), \tilde{R}|J_i) - \{b^*(J_i) \neq \emptyset\} F_i \right], \\ b^*(J_i, F_i) &= \arg \max_{b \in \{\mathcal{B} \cup \emptyset\}} \left[W(b, \tilde{R}^*(J_i)|J_i) - \{b \neq \emptyset\} F_i \right]. \end{aligned}$$

Before going to our results, we explain the computational algorithm we use to obtain them. A reader who is not interested in the computational algorithm should proceed to section 4.

3 Computation and testing

This section provides a reinforcement-learning algorithm that computes a REBE for our baseline model. We then provide a test for boundary consistency of a computed REBE.

The algorithm models players as having perceptions on the value that is likely to result from the different actions available to them at each state. The players choose the actions that are optimal given those perceptions and the realized participation fees. The realizations of random variables whose distributions are determined by the chosen actions and the current state lead to a current profit and a new state. Players use this profit, together with their perceptions of the continuation values they assign to the new state, to update their perceptions of the values of the starting state. They then proceed to choose an optimal policy for the new state that maximizes the perception of the value from that state. This process continues iteratively.

As explained in Fershtman and Pakes (2012), the reinforcement-learning algorithm described above is an algorithm that firms could actually use to learn the values associated with various actions. If the game is a capital-accumulation game structured such that the transition probabilities for the private-information component of a firm's state depend only on the given firm's own policies, and those policies do not influence the evolution of public information, then the firm would learn the distribution of future states conditional on all of its possible actions. This is not necessarily the case when the game is not a capital-accumulation game, such as the sequence of auctions we consider here. The reason is that in a general game a firm might never know what the evolution of its state would have been if it had played an action off the equilibrium path, even if that action would have kept the firm in the recurrent class with a probability of 1. For example, in the auction game we consider here, a firm that wins the auction at an optimal bid will not learn from repeated equilibrium play what would have happened if it had bid a lower value (since in this auction game firms do not observe the non-winning bids of competitors).

We could perturb the algorithm to maintain the analogy with learning by forcing firms to experiment with different policies at each state (as in Fudenberg and Levine (1998)). This approach would, however, increase the complexity of the algorithm. A less computationally burdensome way of computing a REBE is to use knowledge that the computer has in its memory, but the firm does not have, to update the values associated with all policies (even those the firm does not take).²⁷

We begin this section by outlining the computational algorithm for an arbitrary set of initial conditions and providing a test of whether the output of the algorithm constitutes a REBE. We then discuss how one can test whether the output of the algorithm satisfies boundary consistency, that is, with the stronger notion of equilibrium that

²⁷This is a general point that applies equally to games without private information. The computational approach adopted in Besanko et al. (2014) provides a helpful comparison.

ensures that feasible, though non-optimal, actions at the boundary points are indeed non-optimal.

3.1 The Algorithm

The algorithm consists of an iterative procedure and subroutines for calculating initial values and profits. We begin with the iterative procedure. Each iteration, indexed by k , starts with a location that is a state of the game (the information sets of the players) $\mathbf{s}^k = [J_1^k, \dots, J_n^k]$, and has objects in memory, $M^k = \{M^k(\mathbf{s}) : \mathbf{s} \in \mathcal{S}\}$. Each iteration updates both the location and the memory. The rule for when to stop the iterations consists of a test of whether the equilibrium conditions defined in the last section are satisfied. We begin with the basic algorithm and then move on to testing. We consider ways to increase the efficiency of the basic algorithm in the results section.

Memory: The elements of $M^k(\mathbf{s})$ specify the objects in memory at iteration k for information set J , and hence the memory requirements of the algorithm. Often, there will be more than one way to structure the memory, each with different advantages. Here, we focus on a simple structure that will always be available (though not necessarily always efficient; see Fershtman and Pakes, 2012).

$M^k(\mathbf{s})$ contains a counter, $h^k(\mathbf{s})$, that keeps track of the number of times we have visited \mathbf{s} prior to iteration k . If $h^k(\mathbf{s}) > 0$, it also contains

$$\{W^k(b|J_i)\}_{b \in \mathcal{B} \cup \emptyset}.$$

If $h^k(\mathbf{s}) = 0$, nothing is in memory at location \mathbf{s} . When we need to evaluate policies at an \mathbf{s} at which $h^k(\mathbf{s}) = 0$, we have an initiation procedure that sets

$$\{W^k(b|J_i)\}_{b \in \mathcal{B} \cup \emptyset} = \{W^0(b|J_i)\}_{b \in \mathcal{B} \cup \emptyset}.$$

The choice of initial conditions, the $\{W^0(b|J_i)\}_{b \in \mathcal{B} \cup \emptyset}$, is discussed below.

Updating at \mathbf{s}^k : We find the values in memory associated with different b for each firm at location \mathbf{s}^k , take a random draw on F_i , and determine the optimal bid as

$$b^*(J_i^k, F_i) \equiv \operatorname{argmax}_{b \in \mathcal{B} \cup \emptyset} [W^k(b|J_i^k) - \{b \neq \emptyset\}F_i].$$

These bids determine which, if any, player wins the auction. Let $b^k \equiv \operatorname{Max}_i \{b^*(J_i^k, F_i)\}$ be the highest bid at iteration k . If $b^k \neq \emptyset$, an auction occurs. We assume that if an auction occurs, and more than one firm bids b^k , a lottery determines the winning bid.

The winning bid (b^k), the identity of the winner (i_*^k), and the participation decisions of all firms (the vector \mathbf{p}^k) enable us to update the public information sets as

$$\xi^{k+1} = \{\tau_k = 0\}(\omega^k, \tau_{k+1} = 1, i_*^k) + \{\tau_k \neq 0\}(\xi^k(\tau^k + 1), \mathbf{p}^k, i_*^k, b^k),$$

where $\xi^k(\tau^k + 1)$ is the notation for ξ^k with τ changed from τ_k to $\tau_k + 1$. That is, if we are in a full-information-exchange period (if $\tau_k = 0$), we reveal all information about ω , delete the variables in ξ^k (as the revelation of ω makes them irrelevant), and add the identity of the winning bidder. If $\tau_k \neq 0$, we simply add the newly generated information (i_*^k, b^k, \mathbf{p}^k) to the old information set and increase its τ by one.

After bids are submitted and information is revealed, but before the next auction occurs, the firm that wins the auction gathers its new timber and all firms sell what they can to the market. The random draws from the auctioned lot (η) and from the harvests (ϵ_i for each i) are realized and each firm's stock of timber is augmented as

$$\omega_i^{k+1} = \max\{0, \omega_i^k - (e + \epsilon_i) + \{i = i_*\}(\theta + \eta)\}.$$

Thus, the information prior to the next auction is given by

$$J_i^{k+1} = \{\xi^{k+1}, \omega_i^{k+1}\}, \quad \text{and} \quad \mathbf{s}^{k+1} = \{J_1^{k+1}, \dots, J_n^{k+1}\},$$

where it is understood that ω_i^k is omitted from firm's J_i^{k+1} .

Updating the Values in Memory: The algorithm uses the information generated by the random draws that lead to the new location to update firms' perceptions of the values associated with the different policies. The updating is "asynchronous" in that we only update objects in memory associated with the location \mathbf{s}^k , but we update each component of $\{W^k(b|J_i^k)\}_{b \in \mathcal{B} \cup \emptyset}$ for all i at that \mathbf{s}^k . In other words, we update the continuation values for the policies not taken as well as for those taken.

The update for each $W^k(b|J_i^k)$ assumes that the profits and the continuation state that would have accrued to the firm had it chosen that b are those that would have been generated by the competitor's chosen policy, the current state, and random draws from the primitive processes. That is, we assume that the "realized" value that would have been obtained from playing that b was one draw from the expected value of choosing strategy b at J_i^k . The "realized" value is evaluated as the profits it would have earned had it played "b" plus its current perception of the discounted continuation value from the state that it would have moved to. More formally, let $J_i^{k+1}(b, b_{-i}^k, \cdot)$ be the updated information set were we to follow the updating procedure defined above after substituting b for b_i^k in those formulas. This procedure generates $\xi^{k+1}(b, b_{-i}^k, \cdot)$ and $\omega_i^{k+1}(b, b_{-i}^k, \cdot)$. Then, the perceptions of the value for taking action b at state J_i^k are updated as

$$W^{k+1}(b|J_i^k) = \frac{h^k(J_i^k)}{h^k(J_i^k) + 1} W^k(b|J_i^k) + \frac{1}{h^k(J_i^k) + 1} \left[\pi \left(\omega_i^k(b, b_{-i}^k, \cdot), \xi^k(b, b_{-i}^k, \cdot) \right) - b \{i = i_*^k\} + \beta \sum_{\mathbf{F}'} V \left(\omega_i^{k+1}(b, b_{-i}^k, \cdot), \xi^{k+1}(b, b_{-i}^k, \cdot), \mathbf{F}' \right) p(\mathbf{F}') \right].$$

This updating procedure sets the current perception of the value of taking action b at state J_i^k equal to a simple average of what the perception of taking action b would have been had the firm taken that action every time in the past that it had reached J_i^k . Notice that if $W^k(b|J_i^k) = W^*(b|J_i^k)$, then the expectation of $W^{k+1}(\cdot)$ is $W^k(\cdot)$; that is, if we get to an equilibrium set of valuations, the algorithm will tend to stay there.²⁸

Though this averaging procedure does satisfy the Robbins and Monroe (1951) criteria for convergence of a stochastic integral, it is unlikely to be efficient. One reason is that the earlier values are associated with less precise evaluations. We come back to discussing ways of increasing computational efficiency in the results section, and now turn to the testing procedure.

²⁸As in other methods of computing solutions to dynamic games, this algorithm is not converge. Though we have not had convergence problems in our calculations, formally, all we can do is test whether it has converged. We provide the test in the next subsection.

3.2 Testing Procedures

Appendix A provides a detailed explanation of how to test whether the output of the algorithm satisfies the conditions of a REBE. It is analogous to the test described in Fershtman and Pakes (2012), so in the text, we suffice with a brief overview of how to construct the test statistic. We then consider testing for boundary consistency. This concept is new to this paper, and the test has elements that differ from the test for REBE as it requires testing for the validity of moment inequalities. Accordingly, we go over the test for boundary consistency in more detail.

3.2.1 Testing for a REBE

The test is designed to check whether the computed values, together with the policies and the recurrent class that they generate, satisfy conditions C1 to C3 above. We stop the algorithm at a particular iteration, k , and conduct the test based on the values of $W^k(b|J_i^k)$ in memory at that point. In describing the test, we denote these values by $W^*(b|J_i)$.

The test is based on simulating a sample path with the optimal policies generated by $W^*(b|J_i)$. Because the state space is finite, the simulated path will wander into a recurrent class after a finite number of iterations, and stay within that class thereafter. Every point within that class will be visited repeatedly. We keep a record of each point visited in the test's simulation run in separate memory.

The first time a particular point is visited, we record the simulated continuation value resulting from taking every possible action at that point. The values are given by the profits plus the discounted continuation value (evaluated by $\{W^*(\cdot|\cdot)\}$) generated by each action, the policy chosen by competitors, and the simulated random draws on the primitives.²⁹ We also record the square of this continuation value and initiate a counter for the number of times this point was visited in the simulation run. Recall that we visit each point in the recurrent class repeatedly. At each subsequent time a given point is visited, we again calculate a simulated continuation value for each possible policy and then form an average of the simulated continuation values from each time the point was visited for all policies at the point. A similar averaging is used for the continuation value squared. When the simulation run is stopped, the memory for each point visited consists of the average of past simulated continuation values from that point, the average of the continuation values squared, and the number of times the point has been visited in the test run.

The squared difference between $W^*(b|J_i)$ and the estimated average continuation value for playing policy b at J_i is the mean squared error of our estimate of $W^*(b|J_i)$. It can be additively decomposed in the standard way into the bias squared of our estimate and the variance of our estimate. The variance is unbiasedly estimated by the average of the squared value minus the estimated average squared. So by differencing the mean squared error from the estimate of the variance, we are able to get an unbiased estimate of the bias in our estimator for $W^*(b|J_i)$. Our test statistic is a weighted average of the percentage bias (squared) in our estimates of $W^*(b|J_i)$. We weight the different b at a given J_i equally, and the sum over b at different J_i by the number of times that J_i was visited in the simulation run.

²⁹Since the stage game has simultaneous moves, we can evaluate a counterfactual choice of a given firm's policy by substituting the policy and the optimal policies of competitor into this calculation.

More formally, the test is an $L^2(P_{R(ns)})$ norm of the bias in the sum of simulated continuation values as estimates for W^* , where $P_{R(ns)}$ refers to the simulated estimate of the distribution of recurrent states generated by W^* . We accept the test when the test statistic is less than .001—heuristically when our R^2 is above .999. For more details, see appendix B.

3.2.2 Testing for Boundary Consistency (Condition C4).

We begin with a verbal explanation of the test for a given $\{W(b|J_i)\}_{b,J_i}$. Initially, we run a 5-million-iteration simulation from the last point visited in the algorithm. We consider the points visited during the simulation as the points in the recurrent class, and tabulate the fraction of times each of those points is visited during this simulation run, denoted by $\{h(J_i)\}_{J_i}$.

We then start new simulation runs from every point visited in this simulation run for every possible policy from that point. This procedure is analogous to the simulation procedure used in the test for a REBE, except that in the boundary-consistency test, we have to do it for every possible policy. We continue each of the simulation runs for every (b, J_i) until the run enters a point in our estimate of the recurrent class. We keep track of the discounted profits that the firm earns from the simulation run until the simulation enters the recurrent class, and add to this the discounted proposed equilibrium continuation value from the entry point to the recurrent class. Under the null of a boundary consistent REBE, the result is an unbiased estimate of the expected discounted value from taking the policy b at J_i . This value is tabulated and averaged with the other simulated discounted values obtained from the given (b, J_i) . We then determine which of the (b, J_i) are boundary couples by checking if there exists any simulation from the given (b, J_i) that does not enter the recurrent class immediately. Finally, we introduce a test of C4 and apply it to the boundary couples.

We now provide a more formal description of the testing procedure. For each point J_i and each $b \in B \cup \emptyset$, start R simulation runs using the policies generated by $\{W(\cdot|\cdot)\}$. Index the runs from each (b, J_i) couple by r and let the sequence of states visited during the r^{th} simulation run be $\{J_{i,\gamma_r}\}_{\gamma_r=0}^{\gamma_r^*}$, where γ_r^* is the period in the simulation run where the simulation enters the recurrent class (or some sufficiently large number, which we take as 100).

Our estimate of the discounted value of net cash flows from run r for the couple (b, J_i) is

$$\begin{aligned} \hat{W}_r(b|J_i) \equiv & \sum_{\gamma_r=0}^{\gamma_r^*-1} \beta^{\gamma_r} \left(\pi \left(b(J_{i,\gamma_r}, F_{i,\gamma_r}), b(J_{-i,\gamma_r}, F_{-i,\gamma_r}), \omega_{i,\gamma_r}, \epsilon_{i,\gamma_r}, \eta_{\gamma_r} \right) \right) \\ & - \sum_{\gamma_r=0}^{\gamma_r^*-1} \beta^{\gamma_r} \{b(J_{i,\gamma_r}, F_{i,\gamma_r}) \neq \emptyset\} F_{i,\gamma_r} + \beta^{\gamma_r^*} W(b^*|J_{i,\gamma_r^*}), \end{aligned}$$

where $b(J_{i,0}, F_{i,0}) = b$ is the policy we are evaluating. We keep in memory the average of $\hat{W}_r(b|J_i)$, the average of $\hat{W}_r(b|J_i)^2$, and the maximum of γ_r^* from the R simulation runs from each (b, J_i) .

Let $\chi(b, J_i) = 1$ whenever $\max_r \gamma_r^*(b, J_i) \neq 1$, where $\gamma_r^*(b, J_i)$ is the γ^* associated with a particular (b, J_i) . Then,

$$\hat{B} = \{(b, J_i) : \chi(b, J_i) = 1\}$$

is our estimate of the set of boundary couples. For each of these couples, we have a sample mean $\overline{W}^R(b|J_i)$ that is an unbiased estimate of the population mean from R sample paths (in our case $R = 20$). Then, we use the average of the sum of squares of $\hat{W}_r(b|J_i)$ and this sample mean to calculate an unbiased estimate of $Var[\overline{W}^R(b|J_i)]$.

We now use this information to form a test. We are testing inequalities, that is, whether the boundary point policies lead to lower discounted values of future net cash flows that are compared to the optimal policy at the J_i associated with the boundary point. Therefore, we have to use a test statistic that is not pivotal whose distribution does not have a standard form (like the chi-square or normal). We define the statistic below and then explain how we can construct its distribution under the null that our conditions are satisfied. We accept the the null of boundary consistency if the observed value of the test statistic is less than the 95th quantile of the distribution we construct.

The observed test statistic for boundary consistency for the points in \hat{B} . Let $\hat{B}(J_i) = \{b : (b, J_i) \in \hat{B} \subset \mathcal{B} \cup \emptyset\}$ and $\#\hat{B}(J_i)$ be the number of elements in $\hat{B}(J_i)$. Also let

$$T(J_i) = \frac{1}{\#\hat{B}(J_i)} \sum_{b \in \hat{B}(J_i)} \left(\frac{[\overline{W}^R(b|J_i) - W(b^*|J_i)]_+}{W(b^*(J_i))} \right),$$

where $[\overline{W}^R(b|J_i) - W(b^*|J_i)]_+ = \max[\overline{W}^R(b|J_i) - W(b^*|J_i), 0]$.

Let $\mathcal{J}_{\hat{B}}$ be the set of J_i for which there is an element in \hat{B} . Recall that $h(J_i)$ is the number of visits to the point J_i in the initial simulation run, and calculate for each $J_i \in \mathcal{J}_{\hat{B}}$

$$q(J_i) = \frac{h(J_i)}{\sum_{J_i \in \mathcal{J}_{\hat{B}}} h(J_i)}.$$

Our test statistic is

$$T(\hat{B}) = \sum_{J_i \in \mathcal{J}_{\hat{B}}} q(J_i)T(J_i).$$

The simulated distribution of the test statistic under a conservative null: We now simulate the distribution of, $T(\hat{B})$, under the null that $W(b|J_i) = W(b^*|J_i)$ for each $(b, J_i) \in B$, thereby insuring the size of the test.³⁰ For each $(b, J_i) \in \hat{B}$, take ns independent random draws from a normal with mean zero and variance $Var[\overline{W}^R(b|J_i)]$, and call them, $z(b, J_i)_1, \dots, z(b, J_i)_{ns}$ (we set $ns = 50$). For each draw, indexed by $r = 1, \dots, ns$, calculate

$$\tilde{T}(J_i)_r = \frac{1}{\#B(J_i)} \sum_{b \in \hat{B}(J_i)} \left(\frac{[z(b, J_i)_r]_+}{W(b^*(J_i))} \right),$$

and

$$\tilde{T}(\hat{B})_r = \sum_{J_i \in \mathcal{J}_{\hat{B}}} q(J_i)\tilde{T}(J_i)_r.$$

³⁰The test used here is often referred to as the least favorable test statistic in the econometric literature; see, for example, Romano, Shaikh, and Wolf (2014).

Let $\tilde{T}(\hat{B})_{ns}^{.95}$ be the 95th percentile of the distribution of $\tilde{T}(\hat{B})_r$. Then, we accept the null of

$$H_0 : \text{Boundary Consistency}$$

if and only if

$$\tilde{T}(B)_{ns}^{.95} > T(B). \spadesuit$$

3.3 Initial Conditions and Boundary Consistency

Recall that there may be many equilibria that satisfy our equilibrium conditions. In computational problems, the choice of initial conditions for continuation values (our $\{W^0(\cdot)\}$) is a determinant of which equilibria the algorithm will compute.³¹ If the initial conditions are higher than possible equilibrium values, then all policies are likely to be explored. As a result, an equilibrium obtained starting with high initial values is more likely to be boundary consistent (though a formal test for boundary consistency is still advised).³²

The cost of choosing high initial conditions is that they are likely to require many iterations before the test for a REBE in section 3.2.2 is satisfied. So starting with lower values is likely to lead to computed values that satisfy the conditions of a REBE with shorter computation times. Moreover, if a particular equilibrium is desired (boundary consistent or not), and there is prior information on the likely value functions for that equilibrium, for example, from a computed equilibrium from a related model, more carefully chosen initial conditions are likely to lead to the desired equilibrium values faster.

Because we were interested in equilibria that explored all possible policies, we have incurred the cost associated with exceptionally high initial conditions. Hence, all but one of our computed REBEs were boundary consistent.

4 Numerical Analysis

A parameterized version of each of the baseline (B), information-exchange (IE), and voluntary-information-exchange (VIE) models is computed, using the computational algorithm described above. The parameterization and the implementation of the algorithm are discussed below, together with a description of the computational burden. An equilibrium is computed for each of the three models. These equilibria are described in section 4.3, together with a discussion of the economic content of these numerical results.

4.1 Parameter Values

The parameter values that are used in the numerical analysis are given in Table 1 below. Each model has two firms and four possible bids. This structure enables us to

³¹In empirical problems one would use the data to help mitigate multiplicity issues.

³²High initial conditions are also attractive in that there is a concern that low initial conditions may pre-determine an equilibrium. For instance, the recurrent class may be forced to be small if initial conditions are chosen such that there is no incentive to explore a certain region. Boundary consistency also prevents this problem by evaluating whether perceptions at the boundary of a recurrent class are robust to deviating behavior.

Table 1: Parameter specifications

	<i>B</i>	<i>IE</i>	<i>VIE</i>	
Parameters:				
Periods between ω revelation	T	4	1	{1,4}
Common Parameters:				
Distribution of fixed cost of participation	F_i			U[0,1]
Discount factor	β			0.9
Mean timber in a lot	θ			3.5
Disturbance around θ	η			{-0.5,0.5}
Probability on η realizations				{0.5,0.5}
Mean harvest capacity	e			2
Disturbance around e	ϵ			{-1,0,1}
Probability on ϵ realizations				{0.33,0.33,0.33}
Bidding grid				{0.5,1,1.5,2}
Number of firms/bidders				2
Retail price of a unit of timber				1

compute an assortment of models in reasonable time (we discuss computational burdens in the next subsection).³³ Similarly, the time between forced-revelation periods in the baseline model is four periods, a choice arrived at through balancing the desire to have meaningful private information evolving over time with the need to keep the state space at a manageable scale. We assume participation costs are uniformly distributed $U[0, 1]$. To give some sense of scale, this means that the participation costs are between 0 and 50%, and, on average, 25% of the mean revenue generated by a harvested lot of timber.³⁴

4.2 Computational Burden and Updating Procedure

A REBE is computed using the algorithm provided in section 3.1 and the following initial conditions:

$$W^0(b|J_i) = e \left(1 - \frac{\bar{F} + 0.5}{\theta + 1} \right) \frac{1}{1 - \beta} + \omega_i \frac{\bar{F} + 0.5}{\theta + 1}$$

³³In an alternative specification, we restrict the set of information that firms condition on and accommodate an additional firm. The restricted information set includes the time since the last information exchange and the last revealed ω vectors. We find that: (i) qualitatively, restricting the information set has the same effect on bids and participation as revealing information less frequently (moving from *IE* to *B*); and (ii) including an additional firm intensifies competition. The online appendix (available on the authors' webpages) describes the details of the model, computation, and results.

³⁴We elaborate on this assumption here. First, recall that harvesting and production costs are normalized to zero. Second, we have also computed a model with participation cost that distribute $U[0, .5]$, and find that the qualitative results are unchanged.

for all (b, J_i) combinations, and where \bar{F} is the average value of F_i and is 0.5 under our parametrization.

As noted, the initial conditions are likely much higher than equilibrium values. To see why, note that $e/(1 - \beta)$ is the discounted value of being able to sell the mean harvest forever and $e/(\theta + 1)$ is smaller than the periodicity with which the firm would have to win the auction in order to have the timber needed to sell e units in every period. So $(\bar{F} + 0.5)e/[(\theta + 1)(1 - \beta)]$ is less than the cost of bidding in enough periods to be able to sell e units in every period if all the auctions that the firm bid on were won and the winning bid was the lowest bid possible. Finally, $\omega(\bar{F} + .5)/(\theta + 1)$ adds back in the cost of the timber the firm has already stored.

Table 2 provides statistics that summarize different aspects of the computational burden we incurred in computing the equilibria. Partly as a result of our choice of initial conditions, the number of states visited (and hence explored) in both the B and the VIE algorithms was large: 7.5 and 7.9 million respectively. Though recurrent classes were (less than) an order of magnitude smaller than the sets of states ever visited (less than 330,000), there was a significant computational burden in finding them.³⁵ Computation of the IE equilibrium was much less difficult; the number of states visited was only 2,724 and the cardinality of the recurrent class was 2,089, reflecting the fact that the IE model does not require the continuation values associated with every possible different four-period history after the period of revelation.

To lessen the computational burden for the B and VIE model, we used the following simple way of reducing the impact of the bias in the early iterations resulting from the high initial conditions.

1. First, the computational algorithm was run for 50 million iterations, resetting the counters for the states every 10,000 iterations as follows:

$$h(J_i | \text{iteration } 10,001) = \begin{cases} 10 & \text{if } h(J_i | \text{iteration } 10,000) \geq 10 \\ h(J_i | \text{iteration } 10,000) & \text{otherwise.} \end{cases}$$

2. Then, the algorithm was run for 5 million iterations without resetting the counter.
3. Next, a run of additional 5 million iterations was used to form the test for the REBE (recall that the test requires an R^2 statistic to be greater than .999).
4. If the test passed, we stopped the algorithm. Otherwise, we repeated steps 1 to 3.

Steps 1 through 3 were repeated six times for B before the test was satisfied and eight times for VIE . To obtain our results for the IE model, we used a similar procedure but with shorter runs; step 1 above was run for 10 million iterations, and we achieved convergence after only one round of our steps. The boundary-consistency test was run, as described in section 3.2.2, after we accepted the test for the restricted EBE.³⁶ All the equilibria we describe here are boundary consistent, though we have

³⁵The size of the problem illustrated here demonstrates the usefulness of having a computationally feasible approach to deriving equilibrium (that is, one that can be implemented algorithmically on a computer). By comparison, we are unaware of a similarly feasible computation approach that would succeed in computing all the elements of a Markov perfect Bayesian equilibrium in this setting.

³⁶The number of simulation runs used to determine whether a point in the recurrent class was a boundary point was 50, and the number of repetitions to form the averages used in the test of the boundary points was 20.

found one that is not, which we do not report. We provide a summary of computation times in the bottom half of Table 2, and the footnote to the table describes the program and computer used for the runs.

To ensure our estimate of the recurrent class was accurate, we extended the last 5 million run by an additional 5 million and asked what fraction of the incremental iterations visited points that had already been visited in the initial 5 million. For the baseline, information-exchange, and voluntary-information-exchange models, the fractions were 99.42%, 100%, and 98.9%, respectively.³⁷

Table 2: Computational details

Size of recurrent class:		
<i>B</i>	<i>IE</i>	<i>VIE</i>
325,843	2,081	328,692
Number of all states visited during computation:		
<i>B</i>	<i>IE</i>	<i>VIE</i>
7,495,307	2,724	7,908,122
Computation times per 5 million iterations (in hours):		
<i>B</i>	<i>IE</i>	<i>VIE</i>
1:38	1:06	1:56
Computation times for testing for a REBE (5 million iterations, in hours):		
<i>B</i>	<i>IE</i>	<i>VIE</i>
1:43	1:09	2:00
Computation times for testing for boundary consistency (100,000 iterations, in hours):		
<i>B</i>	<i>IE</i>	<i>VIE</i>
3:03	0:16	75:41

Notes: Computation was conducted in MATLAB version R2013a using (a Dell Precision T3610 desktop with) a 3.7 GHz Intel Xeon processor and 16GB RAM on Windows 7 Professional. A round of computation includes steps 1 and 2 of the computational procedure given above. It is 55 million iterations for *B* and *VIE* and 15 million iterations for *IE*.

4.3 Computational Results

Table 3 shows a summary of average per-period performance metrics for each of the *B*, *IE*, and *VIE* models and for a social planner (*SP*) version of the model. The social planner observes all private information of both firms and maximizes total revenues minus participation fees.³⁸ If not for the existence of a non-zero minimum bid, which

³⁷Note also that the incremental points in the *B* and *VIE* cases are likely to be points that satisfy the boundary-consistency conditions.

³⁸Specifically, the planner’s objective is to maximize revenues minus participation fees. That is, the planner views the bid payment as a transfer between players, while participation payments represent real costs to the society. As in the baseline case, each firm draws a stochastic i.i.d. participation cost from $F_i \sim U[0, 1]$ in each period. After observing the realization of the participation costs, the planner chooses which firm to assign the lot to or chooses not to assign the lot to any firm. In terms of the informational structure, we

distorts participation somewhat, the planner’s allocation problem would be equivalent to that of an ideal, perfectly coordinated, cartel: the planner maximizes the discounted value of the sum of future net cash flows.

Table 3: Summary statistics, in per-period terms, by model

	<i>B</i>	<i>IE</i>	<i>VIE</i>	<i>SP</i>
Avg. bid	1.09	0.94	1.04	-
Avg. winning bid (revenue for the auctioneer)	1.11	0.98	1.07	-
Avg. winning bid with ≥ 1 firm participating	1.16	0.98	1.12	-
Avg. winning bid with 1 firm participating	1.06	0.67	0.99	-
Avg. winning bid with 2 firms participating	1.23	1.16	1.20	-
Avg. # of participants	1.52	1.63	1.52	1
Avg. # of participants with ≥ 1 firm participating	1.59	1.63	1.59	1
Avg. participation rate	0.76	0.81	0.76	0.50
% of periods with no participation	4.39	0.15	3.85	0.004
Avg. total revenue	3.35	3.49	3.37	3.50
Avg. profit	0.81	0.87	0.84	-
% of periods; lowest omega wins	66.37	60.80	65.32	85.96
Average total social surplus	2.73	2.72	2.74	3.10

Notes: Here, and in Tables 4, 5, 6, and 7, the per-period profit is defined as $\pi(\omega_i) - \mathbb{I}_{\{i=i^*\}}b_i - \{b_i \neq \emptyset\}F_i = \min\{\omega_i + \mathbb{I}_{\{i=i^*\}}(\theta + \eta), e + \epsilon_i\} - \mathbb{I}_{\{i=i^*\}}b_i - \{b_i \neq \emptyset\}F_i$. Total revenue is defined as $\sum_i \pi(\omega_i) = \sum_i \min\{\omega_i + \mathbb{I}_{\{i=i^*\}}(\theta + \eta), e + \epsilon_i\}$. Total social surplus is defined as $\sum_i \{\pi(\omega_i) - \{b_i \neq \emptyset\}F_i\}$. Averages are taken over periods. The statistics are computed based on a 5-million-iteration simulation of each model.

The average bid for *B*, *IE*, and *VIE* is 1.09, 0.94, and 1.04, respectively. The ordering of bids across models is the same if we look at winning bids, or winning bids conditional on the number of bidders. So if lower prices correspond to weakened competition, the view that information sharing (of strategic data) is akin to collusion has some support, in that information exchange (*IE* and *VIE*) generates lower bids.

However, the statement that more information leads to softer competition is at odds with the observed participation pattern. Participation is higher in *IE* than in *B*. Even if we condition on at least one bidder participating, we see more participation in the *IE* than the *B* equilibrium (1.63 vs 1.59). A common intuition would be that increased participation indicates more competition, which should, in turn, lead to higher bids; however, in this instance, this intuition is contradicted.

Before leaving Table 3, we note that all three models deliver (essentially) the same social surplus (albeit with *IE* being the lowest by 0.01). However, the maximal social surplus from the market equilibria, 2.73, is much lower than the social surplus attained by the planner (3.10). The participation numbers indicate why the planner does so much better. The planner only ever lets one firm enter the auction, thus saving on the cost *F* (the planner also benefits from being able to better coordinate the path of the ω -tuple). In the *IE* equilibrium, the firms generate almost the same revenue (equivalently, output) per period as the planner, but requires much greater participation to do so, thus generating a lower social surplus. By contrast, firms in *B* are less effective at

assume that the planner has access to the F_i and ω_i realizations of both firms.

revenue generation (their stocks are not always high enough to satisfy the demand they face), but generate less wasteful participation.

We now turn to resolving the contrasting patterns of bids and participation between B and IE visible in Table 3: bids are higher in the B game, but participation is higher in IE .

Table 4: Probability Distribution by ω -tuple for B , IE , and SP

(ω_i, ω_{-i})	Prob. Dist. (%)			Profit	
	B	IE	SP	B	IE
$(\leq 4, \leq 4)$	65.51	32.59	90.12	0.68	0.52
$(\leq 4, 5 - 7)$	12.61	19.09	4.52	0.57	0.58
$(\leq 4, \geq 8)$	4.05	10.55	0.28	0.60	0.59
$(5 - 7, \leq 4)$	12.61	19.09	4.52	1.51	1.26
$(5 - 7, 5 - 7)$	0.88	5.72	0.22	1.49	1.46
$(5 - 7, \geq 8)$	0.14	1.12	0.02	1.49	1.13
$(\geq 8, \leq 4)$	4.05	10.55	0.28	1.62	1.58
$(\geq 8, 5 - 7)$	0.14	1.12	0.02	1.66	1.87
$(\geq 8, \geq 8)$	0.01	0.17	0.00	1.72	1.56

Notes: This table shows the probability of intervals of ω -tuples for B , IE , and SP . Here, and in Tables 5, 6, 7, and 9, the per-period profit is a probability-weighted average, over the states underlying each ω -tuple.

To explain these phenomena, we have to consider the relationship between the different information structures and dynamic incentives. We begin with the differences between the IE and B equilibria (the discussion of VIE is delayed until section 4.3.2). Table 4 divides the state space by ω -tuples, and shows the probability distribution over these ω -tuples for each of B and IE as well as the average per-period profits earned by the firms with ω 's in the tuple. The distribution for SP is also provided for comparison.

Table 4 shows that:

1. bidders in IE spend less time in $(\leq 4, \leq 4)$ states than is the case in B ; and,
2. conditional on the state, profits in IE tend to be lower than in B .

Both B and IE are dynamic games in which the control that the firm uses to change its stock of timber is its bid. Hence, to understand how differences in information sets shape the different paths taken through the state space, an examination of bidding is required. The salient feature of the data in Table 4 that the bids must explain is how the IE information structure generates bids that keep the firms in higher ω -tuples. The lower ω -tuples, the tuples in which both firms have $\omega \leq 4$, are the least profitable tuples in *either* equilibrium; indeed the maximal profits for a firm with $\omega \leq 4$ are less than half the minimal profits with $\omega \geq 4$. What is evident from Table 4 is that the additional information available to firms in the IE equilibrium enables them to stay away from states with $\omega \leq 4$ with greater propensity than the firms in the B equilibrium are able to. The fraction of periods with both firms with $\omega \leq 4$ is 65.5% in B compared to 32.6% in IE , while the fraction of states with at least one firm with $\omega \leq 4$ is just over 62% for IE compared to just over 82% for B .

By contrast, the social planner spends more time in the $(\leq 4, \leq 4)$ -tuples than firms in B or IE , thereby generating a smaller cost of holding the timber already procured. So IE firms maintain ω stocks that are greater, and in that sense even less efficient, than in the B equilibrium. Table 4 also reveals that firms in IE spend more time in states that are asymmetric, in the sense of having one firm with a high ω and one with a low ω .

Table 5: Bids by ω -tuple for B and IE

(ω_i, ω_{-i})	Bids										Profit	
	B					IE					B	IE
	\emptyset	0.5	1	1.5	2	\emptyset	0.5	1	1.5	2		
$(\leq 4, \leq 4)$	0.22	0.13	0.27	0.31	0.07	0.07	0.13	0.28	0.47	0.06	0.68	0.52
$(\leq 4, 5-7)$	0.11	0.32	0.45	0.11	0.02	0.02	0.53	0.37	0.08	0.00	0.57	0.58
$(\leq 4, \geq 8)$	0.08	0.58	0.29	0.04	0.02	0.00	0.88	0.12	0.00	0.00	0.60	0.59
$(5-7, \leq 4)$	0.43	0.18	0.34	0.04	0.01	0.33	0.10	0.52	0.05	0.00	1.51	1.26
$(5-7, 5-7)$	0.37	0.50	0.09	0.02	0.01	0.40	0.59	0.01	0.00	0.00	1.49	1.46
$(5-7, \geq 8)$	0.39	0.53	0.06	0.01	0.01	0.11	0.89	0.00	0.00	0.00	1.49	1.13
$(\geq 8, \leq 4)$	0.51	0.25	0.22	0.02	0.00	0.60	0.14	0.26	0.00	0.00	1.62	1.58
$(\geq 8, 5-7)$	0.53	0.39	0.06	0.01	0.00	0.84	0.16	0.00	0.00	0.00	1.66	1.87
$(\geq 8, \geq 8)$	0.61	0.36	0.03	0.00	0.00	0.47	0.53	0.00	0.00	0.00	1.72	1.56

Notes: This table shows the probability of bids by intervals of ω -tuples for B and IE . \emptyset indicates non-participation.

Table 5 contains the probability distributions over bids by ω -tuples, together with average profits in those states. These ω -tuples are the same ones examined in Table 4. Grey shaded cells indicate bids that are more frequent in IE than in B . Notice first that when both firms have $\omega \leq 4$, bidding is more “aggressive” in the IE than in the B equilibrium; we see both more participation in IE (the fraction on \emptyset is lower)

and a higher fraction of bids that are higher than the minimal bid. This observation reinforces the finding that the increased information created when moving from B to IE is not allowing the firms in IE to better coordinate; rather, more information actually intensifies competition when stocks of timber are low. Relative to IE , the firms in the B model are less certain about their competitor’s states, which softens competition.

The opposite is true when both of the firms have $\omega \geq 5$. In these states, conditional on bidding, the bids in IE are smaller. The result is that the winning bid in IE is the minimal bid much more frequently. For example, when both firms have an ω between five and seven, the IE bidding patterns are consistent with firms participating and bidding the minimal amount, when their F_i draw is sufficiently low. The result is that in virtually every case the winning bid is the minimal bid. This essentially reduces the auction to a lottery. When both firms have an ω between five and seven in the B equilibrium, participation is somewhat lower, but conditional on participating about a quarter of the bids are more than the minimal bid. A similar comparison holds when both firms have an ω greater than eight.

Similarly, in the $(\geq 8, 5-7)$ -tuple, the IE equilibrium has the high- ω firm typically sitting out the auction, deferring to the lower- ω rival who most often wins with the minimal bid. By contrast, when the B equilibrium is at the tuple $(\geq 8, 5-7)$ the high- ω firm bids 47 % of the time (compared to only 16% of the time in the IE equilibrium), and 15% of those bids are greater than the minimal bid (compared to 0% for the IE equilibrium).

Thus, when both of the firms have $\omega \geq 5$, more information enables better coordination of participation and, consequently, bids.

One set of states in Table 5 that we have not yet discussed is when one firm has an ω less than or equal to four and the other has an ω between five and seven. In a sense, this set of states lies “in-between” the low-stock states in which more information intensifies competition, and the high-stock states in which more information facilitates coordination. In these in-between states, the high- ω firm participates *more* in the IE equilibrium (67% vs. 57%), and 85% of the time that the high- ω firm participates in the IE equilibrium it bids more than the minimum bid (compared to 68% of the time in the B equilibrium). The low- ω firm in the $(\leq 4, 5-7)$ participates more in the IE equilibrium but bids less aggressively than it does in the B equilibrium. The fact that the high- ω firm bids more aggressively in the IE equilibrium but the low- ω firm does not, explains part of the difference between the probabilities of different states between the IE and B model provided in Table 4. That is, this pattern of aggressive bidding by the high- ω firm contributes to the IE model more frequently reaching states in which at least one firms has a high- ω stock. A similar pattern emerges in the states in which one firm has an ω less than or equal to four and the other has an ω of eight or more (with the only qualitative difference being a lower propensity to participate by the high- ω type in the IE equilibrium).

Thus, Table 5 shows that

1. bidders in IE bid more aggressively when, for both firms, $\omega \leq 4$;
2. in IE , when both firms have $\omega \geq 5$, the participation decisions move the auction closer to a form of lottery; and
3. in the remaining (asymmetric) states, in IE , the high- ω firm bids more aggressively than the low- ω firm, conditional on participating.

Tables 6 and 7 examine the differences in bids between the B and IE models in more detail. Table 6 looks at bids in the low- ω states and shows the increase in aggressiveness that results from providing firms with the increased information in the IE equilibrium. At state $(0, 0)$, firms in IE participate 99% of the time (compared to 88% in B), and when they participate, 78% of the time they choose the maximal bid (compared to 28% in B). The differences between the bids in IE and B are similar in state $(1,1)$. Even when some asymmetry exists in the states, for example in $(2,0)$, as long as both states are low, the increased information in IE causes the firm with a higher ω to bid more aggressively in IE than in B .

Table 6: Competition in low ω -tuples

(ω_i, ω_{-i})	Prob. Dist. (%)		Bids										Profit	
	B	IE	B					IE					B	IE
			\emptyset	.5	1	1.5	2	\emptyset	.5	1	1.5	2		
$(0, 0)$	3.17	.50	.12	.07	.12	.41	.28	.01	.00	.09	.12	.78	-.22	-.48
$(0, 1)$	3.70	.88	.12	.08	.13	.46	.20	.04	.00	.09	.44	.43	-.17	-.44
$(0, 2)$	4.91	1.48	.11	.09	.17	.49	.15	.05	.08	.05	.60	.23	-.09	-.31
$(1, 0)$	3.70	.88	.18	.06	.13	.49	.15	.01	.04	.00	.29	.66	.41	-.08
$(1, 1)$	2.36	.80	.18	.12	.23	.40	.07	.03	.09	.00	.74	.15	.46	.20
$(2, 0)$	4.91	1.48	.28	.07	.19	.41	.05	.05	.10	.00	.86	.00	1.01	.66

Notes: This table shows the probability of selected ω -tuples and bids by those ω -tuples for B and IE . \emptyset indicates non-participation.

Table 7 focuses on bidding behavior when states are asymmetric. The firm with the larger stock has an $\omega = 7$, but the pattern is representative of states with its $\omega \in \{5, 6, 7, 8, 9\}$. Relative to the B equilibrium, the low- ω firms in IE have a higher propensity to bid and, when bidding, to bid the minimum bid. Moreover, those propensities increase as their stock moves from 0 to 1 to 2. By contrast, at least in states $(7,0)$, $(7,1)$, and $(7,2)$, the high- ω rival either does not participate or tends to bid 1 (and so is likely to win if it does bid). As the low- ω firm's stock increases, the high- ω firm participates less, making the low- ω firm more likely to win and win with the minimal bid. This ensures that both firms' profits increase as the low- ω firm's stock increases.

In the IE equilibrium, this pattern of play shifts as the low- ω firm passes $\omega=4$. Then, the high- ω firm (if it bids) moves its bids toward the minimal bid. The behavior in the B equilibrium in these cases is quite different. Participation and bids conditional on participation are higher, making the relative profitability of those states (relative to the low- ω states) lower in the B than in the IE equilibrium.

Tables 6 and 7 show how increasing a firm's information about its competitor changes the path of play. Although providing more information about a competitor increases competition and lowers profits at low- ω states, participation is higher in the game with more information. The reason is that an increase in information at low- ω states intensifies the competition over future profits in high- ω states. Consider two firms initially at low- ω states. Compared to the B equilibrium, the losing firm in the IE equilibrium has more information about the competitor's stock after winning, and

Table 7: Bidding and participation in asymmetric ω -tuples

(ω_i, ω_{-i})	Prob. Dist. (%)		Bids										Profit	
	B	IE	B					IE					B	IE
			\emptyset	.5	1	1.5	2	\emptyset	.5	1	1.5	2		
(0, 7)	1.49	2.36	.05	.23	.61	.09	.03	.01	.33	.62	.03	.00	.22	.02
(1, 7)	.40	.83	.08	.50	.38	.03	.01	.00	.79	.21	.00	.00	.69	.64
(2, 7)	.35	.89	.14	.64	0.18	0.02	0.01	0.00	1.00	0.00	0.00	0.00	1.06	1.07
(4, 7)	0.13	0.69	0.26	0.61	0.10	0.02	0.02	0.04	0.96	0.00	0.00	0.00	1.36	1.09
(7, 0)	1.49	2.36	0.46	0.10	0.41	0.03	0.01	0.26	0.00	0.74	0.00	0.00	1.55	1.17
(7, 1)	0.40	0.83	0.48	0.23	0.26	0.02	0.00	0.40	0.03	0.57	0.00	0.00	1.57	1.21
(7, 2)	0.35	0.89	0.48	0.29	0.21	0.02	0.00	0.50	0.11	0.39	0.00	0.00	1.57	1.39
(7, 4)	0.13	0.69	0.46	0.43	0.09	0.02	0.01	0.76	0.24	0.00	0.00	0.00	1.59	1.84
(7, 7)	0.02	0.26	0.45	0.47	0.06	0.01	0.00	0.47	0.53	0.00	0.00	0.00	1.61	1.49

Notes: This table shows the probability of selected ω -tuples and bids by those ω -tuples for B and IE . \emptyset indicates non-participation.

thus the likelihood that the competitor will participate in the following auction. Based on this information, the firm can better assess whether it is likely to win subsequent auctions with a minimal bid, and bid accordingly. Though this fact certainly does not dull the incentive to bid aggressively when both stocks are low, it does ameliorate the consequences of initial losses and support an equilibrium where both firms are at high- ω (and hence highly profitable) states more often.³⁹

More generally, the reduction in asymmetric information, caused by moving from B to IE , intensifies competition in low- ω -tuples (causing a reduction in profits in those states), but mitigates competition in high- ω states (in the sense of making coordinating participation easier). The result is an environment in which firms invest in maintaining higher ω stocks, and thus spend more time in parts of the state space in which competition is less intense.

4.3.1 The Model with Static Incentives (i.e. $\beta = 0$)

Note that when we set $\beta = 0$, the firms still use the prior history as signals on the likely current stock of timber held by their competitors. However, they now bid to maximize current profits with no interest in investing for future use. The striking implication of the computational results in Table 8 is that in the absence of an incentive to invest in the future, whether or not firms share information has little impact on their behavior. That is, when $\beta = 0$, the outcomes generated by B and IE are not different in any economically meaningful way. This finding confirms that the primary impact of the additional information in the IE equilibrium is to enable the firms to plan for the

³⁹This finding resembles the intuition provided in Jofre-Bonet and Pesendorfer (2003) in which low-inventory (unconstrained) firms may bid less aggressively in the presence of a competitor with medium inventory, allowing the competitor with medium inventory to win. This allows the low-inventory firm to face a high-inventory competitor in the future, which effectively gives monopoly power to the low-inventory firm. It is easier to settle on such a pattern of play when inventory levels are relatively transparent to bidders. Hence, this pattern is more apparent in the IE case than in the baseline.

future, and this, in turn, changes the equilibrium distribution of states⁴⁰.

Table 8: $\beta = .9$ versus $\beta = 0$.

	$\beta = 0.9$		$\beta = 0$	
	<i>B</i>	<i>IE</i>	<i>B</i>	<i>IE</i>
Avg. bid	1.09	0.94	0.61	0.59
Avg. winning bid (revenue for the auctioneer)	1.11	0.98	0.54	0.53
Avg. winning bid conditional on ≥ 1 firm participating	1.16	0.98	0.62	0.60
Avg. winning bid conditional on 1 firm participating	1.06	0.67	0.55	0.53
Avg. winning bid conditional on 2 firms participating	1.23	1.16	0.82	0.82
Avg. # of participants	1.52	1.63	1.10	1.10
Avg. # of participants conditional on \geq one firm participating	1.59	1.63	1.25	1.25
Avg. participation rate	0.76	0.81	0.55	0.55
% of periods with no participation	4.39	0.15	11.98	11.65
Avg. total revenue	3.35	3.49	3.08	3.09
Avg. profit	0.81	0.87	1.03	1.04
% of periods; lowest omega wins				
Average total social surplus	2.73	2.72	2.60	2.61

4.3.2 Voluntary Information Exchange (*VIE*)

In the *VIE* model, firms can elect to share information every 4 periods. If both firms elect to share information, then the model switches, for the next four periods, from the *B* to the *IE* setting.⁴¹

Table 9 indicates that despite the fact that average profits in *IE* are larger than average profits in *B*, firms in *VIE* only choose to share information in 5% of the states where that choice is made (though one of the two firms chooses to share in 24% of those states). As a result, when we calculate the prior tables, we find little difference between *B* and *VIE*. This finding raises the question of why firms in *VIE* cannot reliably coordinate on sharing information; after all, doing so appears to be in their long-term interest.

Table 9 shows that the propensity to share information is substantial only when both ω 's are greater than four, and the highest is greater than eight. Since the default is *B*, in *VIE* these states occur relatively rarely, hence the low frequency of choosing to share information. Recall that profits are higher in the *B* equilibrium. As a result, to enjoy the benefit of switching to the *IE* equilibrium, the firm has to forsake profits in an intermediate period.

This tradeoff comes out clearly in the comparison presented in Table 10. It reports, for *IE*, the average of $E_{F_i} [V(J_i, F_i) | \tau = 1]$ by the underlying state's ω_i , weighted by the relative frequency with which a state is visited. It also reports the same expectation for an alternate scenario in which optimal policies from the *B* model are followed (from the same initial state) for four periods, and then, for all subsequent states, *IE*-optimal policies are followed. Comparing the two expected valuations indicates the value of switching from no-information sharing directly to information sharing versus waiting

⁴⁰We have also computed for $\beta \in [.25, .5, .8]$. As we increase β the difference between the *IE* and *B* equilibria in the rows of tables analogous to Table 8 grows.

⁴¹If one or both firms choose not to share, then firms spend the next four periods in the *B* setting. For more detail on *VIE*, refer back to section 2.4.2.

Table 9: Individual firm’s choices to reveal by ω -tuple

(ω_i, ω_{-i})	Prob. Dist. (%)	$\Pr(\cup_i \psi_i \geq 1)$	$\Pr(\prod_i \psi_i = 1)$	Profit	
	<i>VIE</i>	<i>VIE</i>		<i>B</i>	<i>IE</i>
$(\leq 4, \leq 4)$	62.98	24.75	4.76	0.68	0.52
$(\leq 4, 5 - 7)$	13.17	24.57	4.47	0.57	0.58
$(\leq 4, \geq 8)$	4.58	28.06	6.09	0.60	0.59
$(5 - 7, \leq 4)$	13.17	21.38	4.47	1.51	1.26
$(5 - 7, 5 - 7)$	1.13	18.94	4.59	1.49	1.46
$(5 - 7, \geq 8)$	0.19	24.38	9.73	1.49	1.13
$(\geq 8, \leq 4)$	4.58	23.39	6.09	1.62	1.58
$(\geq 8, 5 - 7)$	0.19	24.60	9.73	1.66	1.87
$(\geq 8, \geq 8)$	0.02	38.14	20.34	1.72	1.56

Notes: $\psi_i \in \{0, 1\}$, $\psi_i = 1$ indicates that firm i chose to reveal, so $\cup_i \psi_i \geq 1$ indicates that at least one firm chose to reveal, and $\prod_i \psi_i = 1$ indicates both firms chose to reveal. Only periods in which firms decide on information sharing (or periods with $\tau = 0$) are used in the calculation.

four periods and then shifting to information sharing. The last column reports the probability of the value for *IE* being larger than the calculation with four periods of waiting in the simulated data. This probability indicates the fraction of times when any losses in the interim four periods of information exchange are worth less than any gains from information sharing in subsequent periods.

Tables 9 and 10 show the difficulty that the collective of firms have in maintaining information sharing, despite its long-term benefits. This finding suggests the importance of commitment devices in establishing an effective information-sharing arrangement. In *IE*, perfect commitment is externally imposed. In *VIE*, firms are able to commit for only four periods at a time, and this is sufficient to break down information sharing.

5 Conclusion

This paper illustrates how the experience-based equilibrium concept facilitates investigation of the dynamics of complex auction environments. It also extends this equilibrium concept through a boundary-consistency requirement which mitigates the the problem of multiplicity that can be generated by the conditions of experience-based equilibrium.

Our application shows that allowing for the dynamics implicit in many auction environments is important in that it enables the emergence of equilibrium states that can only be reached when firms are responding to dynamic incentives. It also shows that the impact of information sharing can depend crucially on the extent of dynamics, and suggests that treating information sharing, even of strategically important data, as a *per se* offense (in the case of the U.S.) or as a restriction of competition by object (in the case of the EU), needs to be weighed against the possibility of type 1 error, falsely rejecting the hypothesis that conduct is welfare neutral (or even welfare enhancing).

Table 10: $E_{F_i} [V (J_i, F_i) | \tau = 1]$ by ω_i

ω_i	Number of states	IE (A)	B for 4 periods, then IE (B)	Probability of (A) \geq (B)
0	146	6.22	6.34	22.92
1	120	6.89	7.01	32.57
2	131	7.72	7.79	36.47
3	136	8.54	8.58	29.87
4	127	9.35	9.30	63.57
5	120	10.10	10.02	44.79
6	113	10.87	10.70	75.12
7	94	11.60	11.37	87.34
8	87	12.27	11.98	90.58
9	75	12.86	12.52	94.66
10	63	13.40	13.02	99.93
11+	186	14.25	13.88	99.53


Notes: This table shows, for IE , the average of $E_{F_i} [V (J_i, F_i) | \tau = 1]$ by the underlying state's ω_i , weighted by the relative frequency with which a state is visited during a 1-million-iteration simulation of the B model. It then replaces the first four periods of IE by B (and the IE continuation from the resulting end state) to form the same computation for “ B for 4 periods, then IE .” States are selected by taking all $\tau = 1$ states visited during a 1 million iteration simulation of the B model. The number of states is the count of distinct states. The probability of (A) \geq (B) is the percent of times with (A) \geq (B) during a 1-million-iteration simulation of the B model.

A Appendix: The mechanics of REBE in a simple example

We begin with an informal statement of a REBE (see section 2.2 for a more formal statement). A REBE consists of

- C1 A set \mathcal{R} that is a subset of the state space: \mathcal{R} contains all states that can possibly be reached given equilibrium play.
- C2 A set of strategies, one for each player, such that at every state in which a player can take an action, the action that the player takes is optimal, conditional on the W 's attached to each possible action at that state.
- C3 A number, denoted by W , corresponding to each possible action at each state. For states in \mathcal{R} , W 's are equivalent to the sum of the period payoff and the discounted, probability-weighted, continuation value of play from the states that may be reached in the next period.⁴²

Figure 1: The stage game Example 1



A.1 Example 1: Comparing MPBE and REBE

The extensive form of a stage game is shown in Figure 1. We consider an infinitely repeated version of this game (although we consider the perfect Bayesian equilibrium of the stage game as a first step). The game in Figure 1 differs conceptually from the

⁴²The continuation value achieved from a state is the highest W attached to the actions that can be taken at that state. That is, for state J_i , the continuation value from J_i , denoted $V(J_i)$, is $V(J_i) = \max_{a \in A_{J_i}} W(a|J_i)$, where A_{J_i} is the set of actions that can be taken at state J_i . See equations 1 and 4.

game considered in this paper and in Fershtman and Pakes (2012) in that it has an alternating move structure in the stage game. (This example is adapted from Gibbons (1992)).⁴³

In the stage game in Figure 1, player 1 begins by choosing an action from $\{L, M, R\}$. If player 1 chooses either L or M , player 2 finds itself at information set $2a$ and must choose an action from $\{L', R'\}$ without knowing the nature of player 1's proceeding choice. If player 1 chooses R , then player 2 finds itself at information set $2b$ (a singleton) and must choose an action from $\{L'', R''\}$.⁴⁴

A perfect Bayesian equilibrium in this game is for player 1 to play L , for player 2 to play L' at information set $2a$ supported by a belief that player 1 played L conditional on being at information set $2a$, denoted by $\Pr(L|2a)$ such that $\Pr(L|2a) = 1$, and for player 2 to play R'' at information set $2b$. This is summarized as $\langle L, L', \Pr(L|2a) = 1, R'' \rangle$.

For ease of exposition, we restrict the discussion to pure-strategy perfect Bayesian equilibria. This game has two: $\langle L, L', \Pr(L|2a) = 1, R'' \rangle$, and $\langle R, R', \Pr(L|2a) \leq 0.5, R'' \rangle$.

Now consider an infinitely repeated version of this game, such that Figure 1, is the stage game. The repeated game works by transitioning back to player 1's move following any terminal node in the stage game, and assumes that β is the common discount rate. In this repeated game, consider the pure-strategy Markov perfect Bayesian equilibria (MPBE). The Markov restriction limits the state space to pay-off, and informationally, relevant states, which in this simple setting is just the information set in the stage game that each player finds himself at.⁴⁵ Given this simple state space, the pure-strategy MPBE mirror those in the one shot game: $\langle L, L', \Pr(L|2a) = 1, R'' \rangle$, and $\langle R, R', \Pr(L|2a) \leq 0.5, R'' \rangle$.

In this repeated game, it is possible to consider what a REBE would constitute. The infinite repetition allows the empirical distributions, central to the definition of a REBE, to be well defined.

Example 1.1: The following is a REBE:

1. A subset of the state space $\mathcal{R} \subset \mathcal{S} \equiv \{1, 2a, 2b\}$ such that $\mathcal{R} \equiv \{1, 2a\}$.
2. A set of strategies such that player 1 plays L and player 2 plays L' at $2a$ and R'' at $2b$.
3. A set of numbers having an interpretation as the firm's perceptions of the expected discounted values of current and future cash flows conditional on its information set (equivalently, the state) and action. That is,

$$\mathcal{W} \equiv \{W(L|1), W(M|1), W(R|1), W(L'|2a), W(R'|2a), W(L''|2b), W(R''|2b)\}.$$

Since each terminal node in the stage game transfers to the same point in the stage game for the next period, the future cash flows are the same for every element of \mathcal{W} . Denote the discounted value of these cash flows for each player by βV_i . Thus, consider the following values:

$$(a) \quad W(L|1) = 2 + \beta V_1$$

⁴³The adaptation is from p.176 of Gibbons.

⁴⁴Node $2b$ is added to allow a state to lie outside the recurrent class in a REBE.

⁴⁵Note that this very simple structure means that a full-revelation period is not needed in this (and the next) example - the Markov restriction is sufficient to keep the state space finite. This is true for any game that returns to a particular point in finite time with probability 1.

$$(b) \ W(M|1) = 0 + \beta V_1$$

⋮

$$\text{such that } \mathcal{W} = \{2, 0, 1, 1, 0, 0, 3\} + \beta \mathbf{V}.^{46}$$

One can easily check that given the above values \mathcal{W} the above strategies are optimal, R is indeed the recurrent class, and that the consistency condition C3. is satisfied. For illustration, $W(R|1) = 1 + \beta V_1$ because, if player 1 plays R , then player 2 plays R'' (see the definition of the equilibrium strategies). This leads to a period payoff of 1 for player 1. The game then transitions to a new period, starting at node 1, and player 1 gets the discounted continuation value from this state (equaling βV_1).

In this example, $W(L''|2b) = 0 + \beta V_2$ and $W(R''|2b) = 3 + \beta V_2$. These numbers are actually arbitrary. Since node $2b$ is not in the recurrent class, the only constraint on $W(L''|2b)$ and $W(R''|2b)$ comes from needing to provide a pattern of play in the current period such that if player 1 plays R , the period payoff is 1 (we need that $W(R|1) = 1 + \beta V_1$). This minor restriction, which is solely an artifact of the stage game having an alternating move structure, is all that constrains $W(L''|2b)$ and $W(R''|2b)$. Indeed, any numbers such that $W(L''|2b) \leq W(R''|2b)$ would work perfectly well.⁴⁷ Note that not for the alternating move structure, even this minimal restriction would be absent.

Unsurprisingly, as in an MPBE, other patterns of play can be supported by a REBE. For instance, the following is also a REBE

Example 1.2:

1. $\mathcal{R} \equiv \{1, 2b\}$.
2. Player 1 plays R and player 2 plays R' at $2a$ and R'' at $2b$.
3. $\mathcal{W} = \{0, 0, 1, 0, 1, 0, 3\} + \beta \mathbf{V}$.

Here, $W(L'|2a) = 0 + V_2$ and $W(R'|2a) = 1 + V_2$. This makes playing R' optimal for player 2 at $2a$, but since $2a$ is outside the recurrent class, there is no discipline on these numbers. By contrast, $W(L|1)$ and $W(M|1)$ are consistent with what would happen were player 1 to play L or M , and player 2 played according to its equilibrium strategy. This satisfies consistency (C3).⁴⁸

A.2 Example 2

The purpose of this second example is to provide yet another example of REBE and also to demonstrate the mechanics of boundary consistency. Consider the following extensive form game, which comprises two possible subgames.

The following is a REBE:

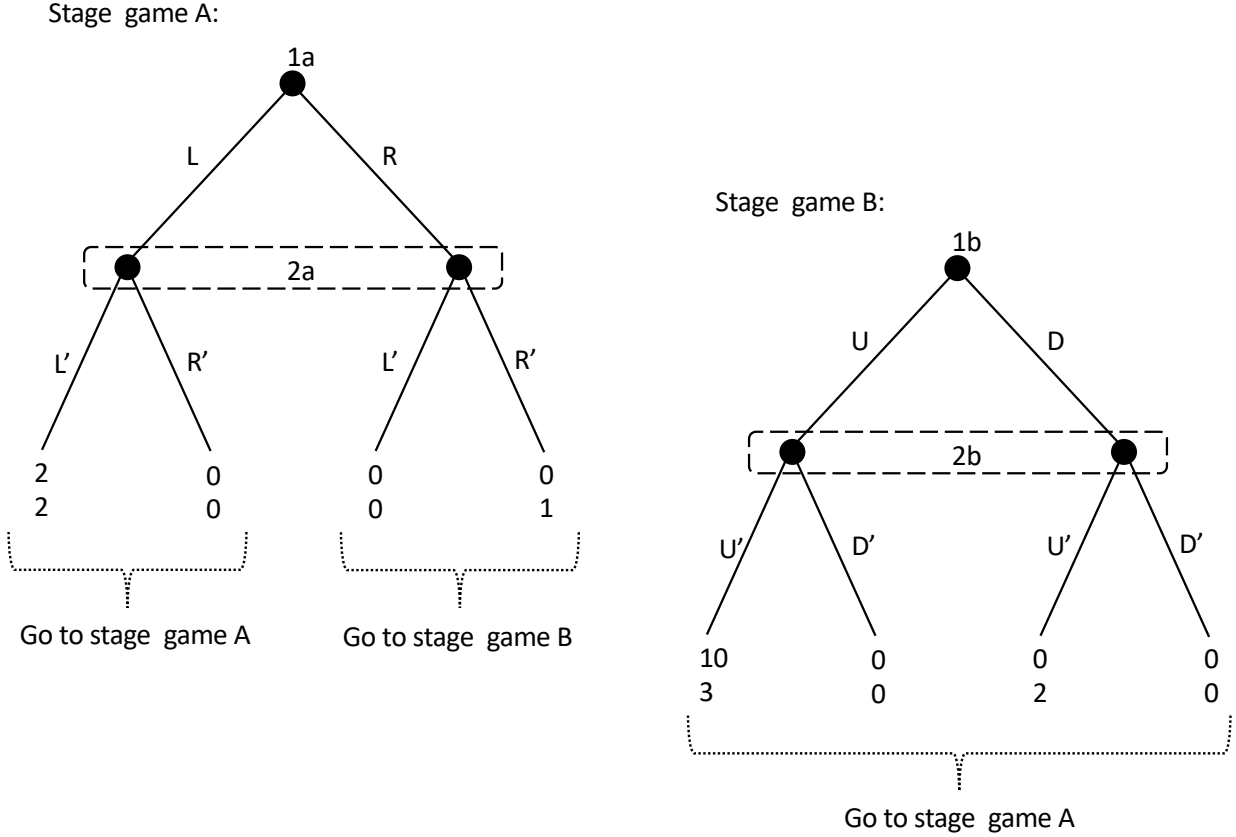
Example 2.1:

⁴⁶ $\mathbf{V} = \{V_1, V_1, V_1, V_2, V_2, V_2, V_2\}$

⁴⁷Formally, this would give rise to multiplicity of equilibria, but no change in the pattern of play. Similarly if $W(L''|2b) > W(R''|2b)$ then all that would change is that the strategies would have to be redefined such that player 2 plays L'' at $2b$ and $W(R|1)$ would have to be changed so that $W(R|1) = 0$ (or else condition C3 would be violated). This would be a distinct REBE, but again the path of equilibrium play would be unchanged.

⁴⁸Note that \mathbf{V} in examples 1.1 and 1.2 are different since the pattern of play is different.

Figure 2: The stage games of Example 2



1. A subset of the state space $\mathcal{R} \subset \mathcal{S} \equiv \{1a, 2a, 2b\}$ such that $\mathcal{R} \equiv \{1a, 2a\}$.
2. A set of strategies such that player 1 plays L at node $1a$, and player 2 plays L' at $2a$. Player 1 plays U at node $1b$ and player 2 plays U' at $2b$.
3. A set of numbers having an interpretation as the firm's perceptions of the expected discounted values of current and future cash flows conditional on its information set and action. That is

$$\mathcal{W} \equiv \{W(L|1a), W(R|1a), W(L'|2a), W(R'|2a); W(U|1b), W(D|1b), W(U'|2b), W(D'|2b)\}.$$

Since each of the terminal nodes in each stage game transitions to node $1a$ eventually, it is useful to denote the continuation value for each player at $1a$ by βV_i . Let \mathcal{W} be such that (with commentary added explaining the values for player 1):

- (a) $W(L|1a) = 2 + \beta V_1$: In stage game A, player 2 plays L' , so if player 1 plays L , then player 1 gets a payoff of 2 in the stage game. The game then transitions to the new period, and stage game A is played again. The discounted continuation value of the subsequent play is βV_1 . Hence, $W(L|1a) = 2 + \beta V_1$ and condition C3 is satisfied.
- (b) $W(R|1a) = \beta$: In stage game A, player 2 plays L' , so if player 1 plays R then player 1 gets a payoff of 0 that period. The game then transitions to the new

period, and stage game B is played. The continuation value from this point is the maximum of $W(U|1b) = 1$ and $W(D|1b) = 0$ (see below, for where these values come from). Hence, the continuation value is equal to 1, which when discounted yields β . Hence, $W(R|1a) = 0 + \beta = \beta$ and condition C3 is satisfied.

- (c) $W(L'|2a) = 2 + \beta V_2$.
- (d) $W(R'|2a) = \beta V_2$.
- (e) $W(U|1b) = 1$: Since node $1b$ is outside the recurrent class, this number does not need to satisfy condition C3. This leaves it unconstrained, provided that it remains the case that $W(R|1a) < W(L|1a)$. (Recall that, to satisfy condition C3, $W(R|1a)$ must be consistent with a period profit of 0 and the discounted continuation value generated by transitioning to stage game B. This continuation value is determined by $\max\{W(U|1b), W(D|1b)\} = W(U|1b)$). Any numbers such that $\max\{W(U|1b), W(D|1b)\} = W(U|1b) \leq (2 + \beta V_1)/\beta$ would suffice.
- (f) $W(D|1b) = 0$: This is arbitrary. See commentary on $W(U|1b)$ above.
- (g) $W(U'|2b) = 2 + \beta V_2$.
- (h) $W(D'|2b) = \beta V_2$.

The pattern of play in this game is to stay in subgame 1 indefinitely, with player 1 playing L and player 2 playing L' . The key to supporting this REBE is that $W(R|1a) = \beta < W(L|1a)$, which means that L is the optimal action for player 1 at node $1a$. $W(R|1a) = \beta$ satisfies C3 (the consistency condition) because $W(U|1b) = 1 > W(D|1b)$. Hence, both the period profit from playing R at $1b$ is zero, and the continuation value from reaching $1b$ (given by $W(U|1b)$) is 1. So, C3 is satisfied. $W(U|1b)$, which is outside the recurrent class, does not have to satisfy C3, which is why $W(U|1b) = 1$ is possible. Note that $W(D|1b)$, $W(U'|2b)$, and $W(D'|2b)$ can also be arbitrary, since the nodes are outside the recurrent class. The commentary in the statement of the equilibrium above goes through this more slowly.

This pattern of play (always staying in stage game A) is counterintuitive for high enough β . Stage game A can be solved by iteration of dominated strategies - U, U' is a compelling pattern of play - and one would think that (for high β) player 1 would want to cycle through stage game B as often as possible to collect the period payoff of 10.

This REBE fails boundary consistency, for a high enough β . The reason is that, if player 1 were to play R at $1a$, the discounted cash flows that would result can be higher than $W(L|1a)$. To see this, note that playing R at $1a$ would result in a period payoff of 0, and then player 1 would next choose an action at $1b$. At $1b$, the policy played by player 1 is U , and player 2 plays U' , resulting in a payoff of 10 for player 1. Following that, the game transitions back to stage game A. This means that the discounted payoff from playing R at $1a$ is $0 + 10\beta + \beta^2 V_1$, which is greater than $W(L|1a) = 2 + \beta V_1$ if β is high enough (close to 1). Thus, boundary consistency can be violated by this REBE.

Boundary consistency eliminates example 2.1, but does not eliminate all REBE supporting this unattractive pattern of play. For instance,

Example 2.2:

1. A subset of the state space $\mathcal{R} \subset \mathcal{S} \equiv \{1a, 1b, 2a, 2b\}$ such that $\mathcal{R} \equiv \{1a, 2a\}$.

2. A set of strategies such that player 1 plays L at node $1a$, and player 2 plays L' at $2a$. Player 1 plays D at node $1b$ and player 2 plays U' at $2b$.
3. $W \equiv \{W(L|1a), W(R|1a), W(L'|2a), W(R'|2a); W(U|1b), W(D|1b), W(U'|2b), W(D'|2b)\}$
 $\equiv \{2 + \beta V_1, \beta^2 V_1, 2 + \beta V_2, 2\beta + \beta^2 V_2; -1, 0, 2 + \beta V_2, \beta V_2\}$.

This example of a REBE satisfies boundary consistency, but still supports the path of play L, L' . The reason is that $W(D|1b)$ is now greater than $W(U|1b)$, resulting in D being the action played at $1b$ (recall that boundary consistency restricts $W(R|1a)$ but not the W 's attached to points outside the recurrent class, like $W(D|1b)$). This means that $W(R|1a)$ remains less than $W(L|1a)$ when the discounted cash flows resulting from playing R at $1a$ are computed.

To eliminate this equilibrium, the boundary-consistency condition would have to be extended to actions at points one step away from the boundary.

B Appendix: Testing for REBE

In this appendix, we discuss the testing for REBE and the boundary consistency for the baseline case. Analogous procedures are used for the *IE* and *VIE* cases.

Notation and Memory. Iterations of the test will be denoted by l (in contrast to the k notations for iterations of the algorithm for computing policies). At each iteration there will be two information sets, one for each firm, so $s_l \equiv (J_{1,l}, J_{2,l})$. In memory, we have particular values of $\left(\{W(b|J_i)\}_{b \in B}, W(0|J_i)\right)$, or $\left(\{W^*(b|J_i)\}_{b \in B}, W^*(0|J_i)\right)$, for all J_i with positive counters ($h^*(J_i) > 0$), and our goal is to determine whether these values satisfy the conditions of a REBE.

At each point visited during the simulation run, we draw an F_i for each firm and calculate

$$V(J_i, F_i) = \max\{\max_b(W^*(b|J_i) - F_i), W^*(0|J_i)\}.$$

The argmax of this equation for each firm will be denoted with a star. Together with the random draws that determine the quantity of timber in the newly acquired lot and those determining the harvest, these policies generate the next state. However, since we are calculating a REBE, we need to simulate the continuation values for all possible policies, that is, for $b \in B \cup \emptyset$.

That is, at iteration l , we calculate the simulated continuation values for firm i and policy b as

$$SCV^l(b|J_i^l) = \pi_i(b_i, b_{-i}^{*,l}, \omega_i^l, \epsilon_i^l, \eta_i^l) + \beta V^*\left(J_i^{l+1}\left(J_i^l, b_i, b_{-i}^{*,l}, \omega_i^l, \epsilon_i^l, \eta_i^l\right), F_i^{l+1}\right).$$

We also calculate $SCV^l(b|J_i^l)^2$. We then update our memory for that point, which consists of; an average of the simulated continuation values, an average of the square of the simulated values, and the counter for the number of times we have visited that point.

Say we stop the simulation routine at a particular $l = \bar{l}$. At that point, we have in memory an average of the estimated simulation value for each possible policy at each

point visited more than once,

$$\mu_{\bar{l}}(b|J_i) \equiv \frac{\sum_{l=1}^{\bar{l}} SCV^l(b|J_i)\{J_i = J_i^l\}}{h_{\bar{l}}(J_i)},$$

and can calculate an unbiased estimate of the variance of the simulated continuation values for each policy at every point,

$$\hat{\sigma}_{\bar{l}}^2(b|J_i) \equiv \frac{\sum_{l=1}^{\bar{l}} SCV^l(b|J_i)^2\{J_i = J_i^l\}}{h_{\bar{l}}(J_i) - 1} - \frac{\mu_{\bar{l}}(b|J_i)^2 h_{\bar{l}}(J_i)}{h_{\bar{l}}(J_i) - 1}.$$

Omitting the index \bar{l} for notational convenience and letting $\#B$ be the cardinality of the set B plus one (for choosing not to enter), we note that the percentage means square error of our estimates at $W^*(J_i)$ or

$$MSE\left(\frac{\mu(J_i)}{W^*(J_i)}\right) \equiv \frac{1}{\#B} \sum_{b \in B \cup \emptyset} \left(\frac{\mu(b|J_i) - W^*(b|J_i)}{W^*(b|J_i)} \right)^2 = Bias^2\left(\frac{\mu(J_i|W^*)}{W^*}\right) + Var\left(\frac{\mu(J_i)}{W^*}\right)$$

where if $E(\cdot)$ takes expectations over the simulated draws,

$$Bias^2\left(\mu(J_i)|W^*\right) \equiv \frac{1}{\#B} \sum_{b \in B \cup \emptyset} \left(E[\mu(b|J_i)] - W^*(b|J_i) \right)^2$$

and

$$Var\left(\mu(J_i)\right) \equiv \frac{1}{\#B} \sum_{b \in B \cup \emptyset} \sigma^2(b|J_i) = \frac{1}{\#B} \sum_{b \in B \cup \emptyset} \left(E[\mu(b|J_i)] - \mu(b|J_i) \right)^2.$$

Our test statistic, labeled Υ , converges to an $L^2(P_{ns}|W^*)$ norm in the percentage bias of the our estimates of W^* , where P_{ns} is the empirical measure of the number of times each J_i is visited in the simulation run (this will converge to $L^2(P_{\mathcal{R}}|W^*)$, the invariant measure of a recurrent class generated by W^*). To obtain a consistent estimate of Υ , we note that

$$\sum_{J_i} \left(\frac{1}{\#B} \sum_{b \in B \cup \emptyset} \hat{\sigma}_{\bar{l}}^2(b|J_i) - Var(\mu(J_i)) \right) p_{ns}(J_i) \rightarrow_{a.s.} 0,$$

so that

$$\Upsilon \equiv \sum_{J_i} \left(MSE\left(\mu(J_i)\right) - \left(\frac{1}{\#B} \sum_{b \in B \cup \emptyset} \hat{\sigma}_{\bar{l}}^2(b|J_i) \right) \right) p_{ns}(J_i) \rightarrow_{a.s.} \sum_{J_i} Bias^2\left(\mu(J_i)|W^*\right) p_{ns}(J_i). \spadesuit$$

We accept the test when $\Upsilon \leq .001$.

References

- Aoyagi, Masaki (2003), Bid Rotation and Collusion in Repeated Auctions, *Journal of Economic Theory*, 112, 79-105.
- Asker, John (2010), A Study of the Internal Organization of a Bidding Cartel, *American Economic Review*, 100(3), 724-762.
- Athey, Susan (2001), Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information, *Econometrica* 69(4), 861-889.
- Athey, Susan and Kyle Bagwell (2008), Collusion with Persistent Cost Shocks, *Econometrica*, 76(3), 493-540.
- Athey, Susan, Jonathan Levin and Enrique Seira (2011), Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions, *Quarterly Journal of Economics*, 126(1), 207-257.
- Bajari, Patrick and Lixin Ye (2003) Deciding Between Competition and Collusion, *The Review of Economics and Statistics*, 85(4), 971-989.
- Baldwin, Laura, Robert Marshall and Jean-Francois Richard (1997), Bidder Collusion at Forest Service Timber Sales, *Journal of Political Economy*, 105(4), 657-699.
- Balat, Jorge (2015), Highway Procurement and the Stimulus Package: Identification and Estimation of Dynamic Auctions with Unobserved Heterogeneity, Working paper, Johns Hopkins University.
- Besanko, David, Ulrich Doraszelski and Yaroslav Kryukov (2014), The Economics of Predation: What Drives Pricing When There Is Learning-by-Doing?, *American Economic Review*, 104(3), 868-897.
- Blundell, Wesley, Gautam Gowrisankaran and Ashley Langer (2018), Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations, NBER Working Paper No. 24810.
- Board, Simon and Moritz Meyer-ter-Vehn (2018), A Reputational Theory of Firm Dynamics, Working paper.
- Bonatti, Alessandro, Gonzalo Cisternas and Juuso Toikka (2017), Dynamic Oligopoly with Incomplete Information *Review of Economic Studies*, 84(2), 503-546.
- Clarke, Richard (1983), Collusion and the Incentive for Information Sharing, *Bell Journal of Economics*, 14(2), 383-394.
- Doraszelski, Ulrich, and Ariel Pakes (2007), A Framework for Applied Dynamic Analysis in IO. In Armstrong, Mark and Robert Porter (eds.) *The Handbook of Industrial Organization Vol. 3*. Elsevier, New York.
- Doyle, Maura, and Christopher Snyder (1999), Information Sharing and Competition in the Motor Vehicle Industry, *Journal of Political Economy*, 107, 1326-1364.
- Ericson, Richard, and Ariel Pakes (1995), Markov Perfect Industry Dynamics: A Framework for Empirical Work, *Review of Economic Studies*, 62(1), 53-82.
- Fershtman, Chaim, and Ariel Pakes (2012), Dynamic Games with Asymmetric Information: A Framework for Empirical Work, *Quarterly Journal of Economics*, 127(4), 1611-1661.

- Fudenberg, Drew, and David K. Levine (1993), Self-Confirming Equilibrium, *Econometrica*, 61(3), 523-545.
- Fudenberg, Drew, and David K. Levine (1998), *The Theory of Learning in Games*, MIT Press, Cambridge.
- Gal-Or, Ester (1985), Information Sharing in Oligopoly, *Econometrica*, 53(2), 329-343.
- Gal-Or, Ester (1986), Information Transmission – Cournot and Bertrand Equilibria, *Review of Economic Studies*, 53(1), 85-92.
- Gibbons, Robert (1992), *Game Theory for Applied Economists*, Princeton University Press, Princeton.
- Graham, Daniel A., and Robert C. Marshall (1987), Collusive Bidder Behavior at Single-Object SecondPrice and English Auctions, *Journal of Political Economy*, 95(6), 1217-39.
- Groeger, Joachim (2014), A Study of Participation in Dynamic Auctions, *International Economic Review*, 55(4), 1129-1154.
- Jeziorski, Przemyslaw, and Elena Krasnokutskaya (2016), Dynamic Auction Environment with Subcontracting, *RAND Journal of Economics*, 47(4), 751-791.
- Jofre-Bonet, Mireia and Martin Pesendorfer (2003), Estimation of a Dynamic Auction Game, *Econometrica*, 71(5), 1443-1489.
- Kirby, Alison (1988), Trade Associations as Information Exchange Mechanisms, *RAND Journal of Economics* 19(1), 138-146.
- Kubitz, Greg, and Kyle Woodwar (2019), Sharing Cost Information in Oligopoly, Working paper.
- Kuhn, Kai-Uwe and Xavier Vives (1995), *Information Exchanges Among Firms and their Impact on Competition*, Office for Official Publications of the European Community, Luxemburg.
- Luco, Fernando (2017), Who Benefits from Information Disclosure? The Case of Retail Gasoline, mimeo, Texas A&M University.
- Mailath, George J., and Peter Zemsky (1991), Collusion in Second Price Auctions with Heterogeneous Bidders, *Games and Economic Behavior*, 3(4), 467-86.
- Maskin, Eric and John Riley (1985), Auction Theory with Private Values, *American Economic Review Papers and Proceedings*, 75(2), 150-155.
- Maskin, Eric and John Riley (2000), Asymmetric Auctions, *Review of Economic Studies* 67, 413-438.
- Maskin, Eric and Jean Tirole (2001), Markov Perfect Equilibrium: Observable Actions, *Journal of Economic Theory*, 100, 191-219.
- Mermelstein, Ben, Volker Nocke, Mark Satterthwaite and Michael D. Whinston (2014), Internal vs. External Growth in Industries with Scale Economies: A Computational Model of Optimal Merger Policy, NBER Working Paper w20051, National Bureau of Economic Research.
- McAfee, Preston and John McMillan (1992), Bidding Rings, *American Economic Review*, 82(3), 579-599.

- Overgaard, Per Baltzer, and H. Peter Møllgaard (2008), Information Exchange, Market Transparency and Dynamic Oligopoly, in *Issues in Competition Law and Policy*, ed. W.D. Collins, 1241-1268, Chicago, IL: American Bar Association.
- Pakes A. Methodological Issues in Analyzing Market Dynamics. In: *Advances in Dynamic and Evolutionary Games: Theory, Applications, and Numerical Methods*, Switzerland: Springer International Publishing (F Thujisman and F Wagener editors); 2016. pp. 43-75
- Pesendorfer Martin (2000), A Study of Collusion in First-Price Auctions, *Review of Economic Studies*, 67(3), 381-411.
- Porter, Robert and Douglas Zona (1993), Detection of Bid Rigging in Procurement Auctions, *Journal of Political Economy*, 101(3), 518-538.
- Porter, Robert and Douglas Zona (1999), Ohio School Milk Markets: An Analysis of Bidding, *RAND Journal of Economics*, 30(2), 263-288
- Robbins, Herbert, and Sutton Monro (1951), A Stochastic Approximation Technique, *Annals of Mathematics and Statistics*, 22, 400-07.
- Romano, Joseph, Azeem Shaikh and Michael Wolf (2014), A Practical Two-Step Method for Testing Moment Inequalities, *Econometrica*, 82(5), 1979-2002
- Saini, Viplav (2013), Endogenous Asymmetry in a Dynamic Procurement Auction, *Rand Journal of Economics*, 69(3), 25-41.
- Shapiro, Carl (1986), Exchange of Information in Oligopoly, *Review of Economic Studies*, 53(3), 433-446.

A Online Appendix: NOT FOR PUBLICATION

A.1 Results from models with a restricted state space

In this section, we consider a model in which firms condition their strategies on a subset of payoff or informationally relevant variables. In this version, the information set includes the time since the last information exchange, $\tau_t \in [1, \dots, T]$, and the last revealed ω vectors, $\omega_{t-\tau_t}$. We compute a baseline model with the restricted state space (B^r) in which information is shared every $T = 4$ periods and an information-exchange model with the restricted state space (IE^r) with $T = 1$. In addition, we compute each of the models with the restricted state space (B^r and IE^r) with three firms.

The computation uses the same following procedure used in computing the B and VIE models.

1. First, the computational algorithm is run for 50 million iterations, resetting the counters for the states every 10,000 iterations as follows:

$$h(J_i | \text{iteration } 10,001) = \begin{cases} 10 & \text{if } h(J_i | \text{iteration } 10,000) \geq 10 \\ h(J_i | \text{iteration } 10,000) & \text{otherwise.} \end{cases}$$

2. Then, the algorithm is run for 5 million iterations without resetting the counter.
3. Next, a run of 5 million iterations is used to form the test for the REBE (recall that the test requires an R^2 statistic to be greater than .999).
4. If the test is passed, we stop the algorithm. Otherwise we repeat steps 1 to 3.

All four specifications with the restricted state space (B^r , IE^r , B^r with three firms, IE^r with three firms) converge after the first round of steps 1 through 3. A summary of computation times is available in table 1.

Table 2 reports a summary of outcomes from B , B^r , B^r with three firms, IE , IE^r , and IE^r with three firms. Restricting the information set leads to increased competition. The average bid increases from 1.09 to 1.29 for the baseline case (going from B to B^r), and from 0.94 to 1.16 for the information-exchange case (going from IE to IE^r). The winning bids and winning bids conditional on the number of bidders also increase. By contrast, the participation rate falls from 0.76 to 0.72 for the baseline model and from 0.81 to 0.71 for the information-exchange model. The results show that restricting the information set by conditioning on fewer variables (moving from B to B^r or from IE to IE^r) has qualitatively the same effects on bids and participation as restricting the information set by revealing information less frequently (moving from IE to B). The profit, total revenue, and social surplus are also lower in B^r than in B (and in IE^r than in IE).

Bids are lower in IE^r than in B^r . Participation is almost the same in B^r and IE^r in contrast to the unrestricted case where participation is higher in IE than in B . As in the unrestricted case, the profit is higher and social surplus is almost unchanged (slightly smaller) when there are more frequent information revelations. However, the total revenue is smaller in IE^r than in B^r (3.28 vs 3.32) unlike in the unrestricted case.

Competition intensifies with an additional firm. The average bid increases from 1.29 to 1.98 for the baseline case with the restricted state space, and from 1.16 to 1.97 for the information-exchange case with the restricted state space. The average number of participants increases from 1.44 for B^r with two firms to 1.71 in B^r with three firms,

Table 1: Computational details

Size of recurrent class:			
B^r	$B^r, N = 3$	IE^r	$IE^r, N = 3$
5776	7576	1467	1642
Number of all states visited during computation:			
B^r	$B^r, N = 3$	IE^r	$IE^r, N = 3$
9514	17282	1784	3405
Computation times per 5 million iterations (in hours):			
B^r	$B^r, N = 3$	IE^r	$IE^r, N = 3$
0:22	0:37	0:21	0:37
Computation times for testing for a REBE (5 million iterations, in hours):			
B^r	$B^r, N = 3$	IE^r	$IE^r, N = 3$
0:21	0:37	0:20	0:36

Notes: Computation was conducted in MATLAB version R2017b using (an iMac Pro desktop with) a 3.2 GHz Intel Xeon processor and 32GB Ram on macOS Mojave.

although the participation rate decreases from 0.72 to 0.57. There is a large decreases in the firm profit from 0.66 to 0.18. The revenue and social surplus decrease as well.

Table 2: Summary statistics, in per-period terms, by model

	B	B^r	$B^r, N = 3$	IE	IE^r	$IE^r, N = 3$
Avg. bid	1.09	1.29	1.98	0.94	1.16	1.97
Avg. winning bid (revenue for the auctioneer)	1.11	1.27	1.85	0.98	1.14	1.85
Avg. winning bid with ≥ 1 firm participating	1.16	1.34	1.99	0.98	1.22	1.99
Avg. winning bid with 1 firm participating	1.06	1.28	1.98	0.67	1.11	1.98
Avg. winning bid with 2 firms participating	1.23	1.40	2.00	1.16	1.31	2.00
Avg. # of participants	1.52	1.44	1.71	1.63	1.43	1.71
Avg. # of participants with ≥ 1 firm participating	1.59	1.52	1.84	1.63	1.52	1.84
Avg. participation rate	0.76	0.72	0.57	0.81	0.71	0.57
% of periods with no participation	4.39	5.06	7.27	0.15	6.33	7.23
Avg. total revenue	3.35	3.32	3.25	3.49	3.28	3.25
Avg. profit	0.81	0.66	0.18	0.87	0.71	0.18
% of periods; lowest omega wins	66.37	76.51	76.00	60.80	67.39	76.56
Avg. total social surplus	2.73	2.60	2.39	2.72	2.57	2.39

Table 3: Probability Distribution by ω -tuple for B , B^r , IE and IE^r

(ω_i, ω_{-i})	Prob. Dist. (%)				Profit			
	B	B^r	IE	IE^r	B	B^r	IE	IE^r
$(\leq 4, \leq 4)$	65.51	87.38	32.59	73.76	0.68	0.68	0.52	0.68
$(\leq 4, 5 - 7)$	12.61	5.80	19.09	10.67	0.57	0.41	0.58	0.61
$(\leq 4, \geq 8)$	4.05	0.45	10.55	2.21	0.60	0.42	0.59	0.70
$(5 - 7, \leq 4)$	12.61	5.80	19.09	10.67	1.51	1.73	1.26	1.59
$(5 - 7, 5 - 7)$	0.88	0.10	5.72	0.40	1.49	1.58	1.46	1.60
$(5 - 7, \geq 8)$	0.14	0.01	1.12	0.05	1.49	1.47	1.13	1.59
$(\geq 8, \leq 4)$	4.05	0.45	10.55	2.21	1.62	1.85	1.58	1.77
$(\geq 8, 5 - 7)$	0.14	0.01	1.12	0.05	1.66	1.54	1.87	1.73
$(\geq 8, \geq 8)$	0.01	0.00	0.17	0.00	1.72	1.67	1.56	1.61

Table 4: Probability Distribution by ω -tuple for B^r and IE^r with three firms

$\omega_i, (\omega_{-i})$	Prob. Dist. (%)		Profit	
	$B^r, N = 3$	$IE^r, N = 3$	$B^r, N = 3$	$IE^r, N = 3$
$\leq 4, (\leq 4, \leq 4)$	96.32	96.67	0.28	0.29
$\leq 4, (\leq 4, 5 - 7)$	1.21	1.10	-0.13	-0.15
$\leq 4, (\leq 4, \geq 8)$	0.02	0.01	-0.16	-0.23
$\leq 4, (5 - 7, 5 - 7)$	0.00	0.00	-0.29	-0.26
$\leq 4, (5 - 7, \geq 8)$	0.00	0.00	-0.02	0.56
$5 - 7, (\leq 4, \leq 4)$	1.21	1.10	1.93	1.96
$5 - 7, (\leq 4, 5 - 7)$	0.00	0.00	1.50	1.34
$5 - 7, (\leq 4, \geq 8)$	0.00	0.00	1.75	0.76
$\geq 8, (\leq 4, \leq 4)$	0.02	0.01	1.62	1.87
$\geq 8, (\leq 4, 5 - 7)$	0.00	0.00	1.91	2.07

Table 5: Bids by ω -tuple for B and B^r

(ω_i, ω_{-i})	Bids										Profit	
	B					B^r					B	B^r
	\emptyset	0.5	1	1.5	2	\emptyset	0.5	1	1.5	2		
$(\leq 4, \leq 4)$	0.22	0.13	0.27	0.31	0.07	0.26	0.09	0.10	0.54	0.00	0.68	0.68
$(\leq 4, 5 - 7)$	0.11	0.32	0.45	0.11	0.02	0.10	0.16	0.22	0.52	0.00	0.57	0.41
$(\leq 4, \geq 8)$	0.08	0.58	0.29	0.04	0.02	0.06	0.15	0.42	0.37	0.00	0.60	0.42
$(5 - 7, \leq 4)$	0.43	0.18	0.34	0.04	0.01	0.67	0.05	0.15	0.13	0.00	1.51	1.73
$(5 - 7, 5 - 7)$	0.37	0.50	0.09	0.02	0.01	0.52	0.23	0.22	0.02	0.00	1.49	1.58
$(5 - 7, \geq 8)$	0.39	0.53	0.06	0.01	0.01	0.40	0.41	0.18	0.01	0.00	1.49	1.47
$(\geq 8, \leq 4)$	0.51	0.25	0.22	0.02	0.00	0.76	0.08	0.13	0.03	0.00	1.62	1.85
$(\geq 8, 5 - 7)$	0.53	0.39	0.06	0.01	0.00	0.47	0.31	0.21	0.00	0.00	1.66	1.54
$(\geq 8, \geq 8)$	0.61	0.36	0.03	0.00	0.00	0.50	0.35	0.15	0.00	0.00	1.72	1.67

Notes: This table shows the probability of bids by intervals of ω -tuples for B and B^r . \emptyset indicates non-participation.

Table 6: Bids by ω -tuple for IE and IE^r

(ω_i, ω_{-i})	Bids										Profit	
	IE					IE^r					IE	IE^r
	\emptyset	0.5	1	1.5	2	\emptyset	0.5	1	1.5	2		
$(\leq 4, \leq 4)$	0.07	0.13	0.28	0.47	0.06	0.27	0.12	0.16	0.41	0.04	0.52	0.68
$(\leq 4, 5-7)$	0.02	0.53	0.37	0.08	0.00	0.12	0.29	0.50	0.09	0.00	0.58	0.61
$(\leq 4, \geq 8)$	0.00	0.88	0.12	0.00	0.00	0.07	0.48	0.45	0.00	0.00	0.59	0.70
$(5-7, \leq 4)$	0.33	0.10	0.52	0.05	0.00	0.54	0.06	0.37	0.03	0.00	1.26	1.59
$(5-7, 5-7)$	0.40	0.59	0.01	0.00	0.00	0.41	0.58	0.02	0.00	0.00	1.46	1.60
$(5-7, \geq 8)$	0.11	0.89	0.00	0.00	0.00	0.43	0.57	0.00	0.00	0.00	1.13	1.59
$(\geq 8, \leq 4)$	0.60	0.14	0.26	0.00	0.00	0.65	0.17	0.18	0.00	0.00	1.58	1.77
$(\geq 8, 5-7)$	0.84	0.16	0.00	0.00	0.00	0.54	0.46	0.00	0.00	0.00	1.87	1.73
$(\geq 8, \geq 8)$	0.47	0.53	0.00	0.00	0.00	0.47	0.52	0.01	0.00	0.00	1.56	1.61

Notes: This table shows the probability of bids by intervals of ω -tuples for IE and IE^r . \emptyset indicates non-participation.

Table 7: Bids by ω -tuple for B^r and IE^r with three firms

(ω_i, ω_{-i})	Bids										Profit	
	B^r					IE^r					B^r	IE^r
	\emptyset	0.5	1	1.5	2	\emptyset	0.5	1	1.5	2		
$\leq 4, (\leq 4, \leq 4)$	0.43	0.01	0.00	0.00	0.57	0.43	0.01	0.01	0.00	0.56	0.28	0.29
$\leq 4, (\leq 4, 5-7)$	0.34	0.00	0.00	0.00	0.66	0.31	0.01	0.01	0.00	0.67	-0.13	-0.15
$\leq 4, (\leq 4, \geq 8)$	0.28	0.02	0.01	0.00	0.68	0.19	0.07	0.02	0.03	0.69	-0.16	-0.23
$\leq 4, (5-7, 5-7)$	0.23	0.05	0.00	0.02	0.70	0.09	0.14	0.09	0.18	0.50	-0.29	-0.26
$\leq 4, (5-7, \geq 8)$	0.00	0.25	0.00	0.00	0.75	0.00	0.00	0.67	0.33	0.00	-0.02	0.56
$5-7, (\leq 4, \leq 4)$	0.78	0.19	0.01	0.00	0.03	0.80	0.14	0.06	0.00	0.00	1.93	1.96
$5-7, (\leq 4, 5-7)$	0.52	0.04	0.04	0.02	0.38	0.50	0.04	0.04	0.17	0.26	1.50	1.34
$5-7, (\leq 4, \geq 8)$	0.40	0.00	0.00	0.00	0.60	0.50	0.00	0.00	0.00	0.50	1.75	0.76
$\geq 8, (\leq 4, \leq 4)$	0.61	0.05	0.03	0.01	0.31	0.75	0.01	0.12	0.04	0.09	1.62	1.87
$\geq 8, (\leq 4, 5-7)$	0.67	0.00	0.00	0.00	0.33	0.60	0.00	0.20	0.00	0.20	1.91	2.07

Notes: This table shows the probability of bids by intervals of ω -tuples for B^r and IE^r with three firms. \emptyset indicates non-participation.

Table with 15 columns of numerical data, grouped into rows by index (e.g., (6, 2), (7, 0), (8, 0), (9, 0), (10, 0), (11+, 0)). The values are organized in a grid format, with some cells containing 'NaN'.