

The Competitive Effects of Information Sharing*

John Asker, Chaim Fershtman, Jihye Jeon, and Ariel Pakes
(UCLA, Tel-Aviv, NYU Stern, and Harvard)

April 1, 2017.

Abstract

We investigate the impact of information sharing between rivals in a dynamic auction with asymmetric information. Firms bid in sequential auctions to obtain inputs. Their inventory of inputs, determined by the results of past auctions, are privately known state variables that determine bidding incentives. The model is analyzed numerically under different information sharing rules. The analysis uses the restricted experience based equilibrium concept of Fershtman and Pakes (2012) which we refine to mitigate multiplicity issues. We find that increased information about competitors' states increases participation and inventories, as the firms are more able to avoid the intense competition in low inventory states. While average bids are lower, social welfare is unchanged and output is increased. Implications for the posture of antitrust regulation toward information sharing agreements are discussed.

Keywords: *Information sharing, Procurement Auctions, Antitrust; Dynamic Games.*

*JEL Codes: D43, K21, L41, C63, C73. We would like to thank numerous seminar audiences for their comments and questions. El Hadi Caoui 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@stern.nyu.edu; and Ariel Pakes: apakes@fas.harvard.edu. The usual caveat applies.

1 Introduction

This paper examines the competitive impact of information sharing between rivals in a dynamic auction with asymmetric information. Firms bid in sequential auctions to obtain inputs. Their inventory of inputs, determined by the results of past auctions, are privately known state variables that determine bidding incentives. The model is analyzed numerically under different rules for sharing information about a firm's competitors' inventories.

Our goal is to shed light on the extent to which dynamic considerations can color the way antitrust regulators, procurement agencies, and other policy actors approach the regulation of information sharing. The antitrust laws of both the U.S. and the E.U., for instance, allow cases to be brought alleging that information sharing depresses competition (as distinct from being a facilitating device for a cartel). Recent enforcement activity suggests that the E.U. takes the less permissive approach, with a stance that strongly discourages firms from sharing strategically sensitive information.¹ Much of the academic economic analysis of information sharing has occurred in static settings. This paper presents a numerically analyzed example of a sequence of auctions that illustrates important competitive effects that static analyses omit.

Information sharing increases the precision of the firm's beliefs about its competitors' states. This helps the firm predict the outcomes from its bids and hence both the states the firm is likely to transition to and the continuation values from those states. An implication is that the increased information makes it easier to ascertain when both firms are likely to be in a low inventory state; a state in which there is increased competitive pressure. Firms respond by increasing their participation rates to increase their levels of inventory and move to portions of the state space in which competition is less intense. That is, information sharing changes bidding behavior due to the desire to transition to more profitable parts of the state space. The net effect is to increase firm profits, while decreasing the auctioneer's surplus. However output, auction participation, and firm inventories also increase. Importantly, social surplus is unaffected by information sharing, reflecting a balance between incurring participation costs and greater output (via higher inventories facilitating more downstream sales).

An important point to bear in mind is that, conditional on the information that they have, firms compete unilaterally. There is no sense in which the information sharing facilitates a scheme that enables firms to explicitly co-ordinate their bids at any point in time; i.e. we are not modeling a cartel. Our contribution is in providing a framework with which to evaluate the impact of information sharing, in and of itself, on the pattern of competition and on welfare.²

The model that is investigated is of an infinite sequence of auctions. The model is loosely based on the description of timber auctions in Baldwin, Marshall and Richard (1997), although, to keep the model simple, many departures are made from the precise institutional features described therein. 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 (inventories). This stock is private information, and its evolution is the source of dynamics. To compete in the auction,

¹See, for instance, *Dole Food Company et al. v. Commission*, Case T-588/08 and Case C-286/13 P.

²The strategy space, at least theoretically, could admit some limited forms of super-game strategies (e.g. tit-for-tat). However, these are not imposed or, as far as we can tell, observed.

firms must pay a participation fee and simultaneously submit a bid. 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. Harvest then occurs, which depletes the stock of timber each firm has.³

In our benchmark model, once every T periods, there is a full revelation of the state variable. That is, during this revelation period each firm observes the stock of unharvested timber of its competitor.⁴ Information sharing is modeled as shrinking the time interval between full revelation periods. 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 every period for the next T periods, or not. For voluntary information sharing to occur over the next T periods, all firms must want to share information.

In our benchmark model firms are more profitable, in the long run, when T is small, or when there is more information on the competitors' state. However either being in a low inventory state, or facing a competitor that is likely in a low inventory state, creates an incentive to withhold information to soften competition. This retards information sharing despite it being in the firms' collective long term interest. That is competitors find it difficult to commit to sharing information. So when the decision to share information is endogenous (rather than being dictated by institutions), a successful information sharing agreement may require the ability to commit for an extended period of time.

Analysis is conducted computationally, applying the Restricted Experience Based Equilibrium (REBE) concept developed in Fershtman and Pakes (2012). This equilibrium concept requires that, for all points in a recurrent class of states (i.e., those states that are visited repeatedly), the actions taken are optimal when evaluated using value functions that are consistent with the net present value generated by equilibrium play. The REBE equilibrium is further refined using a novel, additional, requirement we refer to as *boundary consistency*. Boundary consistency puts further structure on the values attached to actions that carry a firm outside the recurrent class. These values are required to be consistent with the net present value of play that eventuates should a firm play such a strategy. This puts discipline on strategies at the boundary reducing the possible multiplicity that may arise due to the specification of values at points outside the recurrent class.

Several things about the model are important to note. First, in contrast to the vast majority of the literature exploring dynamical oligopoly models, the model computed here is not a capital accumulation game.⁵ Second, it is the first computed model of a dynamic auction with persistent, action-dependent, asymmetric information.⁶ Third,

³The closest model to ours is that 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. In our model, the opposite is true. In particular, this means that information sharing has persistent value across periods and is integral to understanding competition and welfare.

⁴This is also one way to stop the state space from becoming unbounded.

⁵In a capital accumulation game the (stochastic) evolution of an agent's state depends on its own actions. For a survey, see Doraszelski and Pakes (2007).

⁶Athey and Bagwell (2008) analyze a repeated procurement auctions in which cost shocks are correlated

and most importantly, the firms do not collude in the sense of forming any explicit agreement. Rather, contingent on the information that they have, they act solely in their unilateral interest. This is important for the policy application, as it separates the information sharing investigated here from standard economic models of a cartel in which rivals coordinate according to some mutually agreed upon mechanism.

The application to which the model is directed is the antitrust treatment of information sharing. While explicit agreements to fix prices are *per se* violations of the antitrust laws, the legal treatment of information sharing among competitors is less clear.⁷ A useful distinction in organizing court decisions is whether the exchange involves price or non-price information.

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. However, it is also obvious that such an announcement is consistent with normal market competition and so will not attract antitrust sanction. More problematic is the communication of price information between competitors in a way that consumers do not have access to.⁸ U.S. courts apply a rule of reason test to decide whether the exchange of price information constitutes an unreasonable restraint of trade.⁹ 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.¹⁰

U.S. courts take a slightly more 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...”¹¹

over time. However, in their setting, the actions of the bidders do not affect the evolution of costs (types). In our model, actions do affect the evolution of types.

⁷The 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).

⁸In *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 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 stabilization (a form of price fixing).

⁹In 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.

¹⁰A 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 appeal, in *Todd v. Exxon Corp* 275 F.3d 191 (2001).

¹¹see *Maple Flooring Manufacturers' Assn. v. United States* 268 U.S. 563 (1925)

Contemporary guidance from the FTC and DoJ states that “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.”¹² This suggests a somewhat more nuanced view in modern times.

The E.U., 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.).¹³ This may include non-price strategic information.

In this paper, the information exchanged is about the firms inventory of harvestable lumber. Hence, this is a non-price information exchange. However, this information is strategically important, and primary input into the formulation of future prices. As such, it falls in the area in which both U.S. and E.U. regulation is both vague and, likely, divergent.

The numerical results suggest that a harsh approach to the sharing of information is likely misguided. In our setting the amount of timber supplied to the downstream market by the bidders in the auction increases with information sharing. This means more timber is sold in the auction. A classical rule-of-thumb test of whether competition is potentially harmed in a market is whether output is reduced. This test is not satisfied in this setting. That said, the net social impact of the information sharing is close to zero. This is because, to increase output, bidders participate more than is socially optimal, resulting in a cost from increased participation costs. Thus, while the auctioneer is worse off from the information sharing, and prices are depressed, the net economic impact does not suggest a competition problem exists, at least if the test is based on either social welfare or realized output. That is it seems that when dynamic considerations are taken into account the criteria for harm to competition matters crucially for deciding the appropriate policy treatment of information sharing. At least according to standard welfare measures, the numerical results investigated here do not give rise to a case for intervention. By contrast, if the test were solely based on the price effect, a case for intervention may exist.

This paper is organized as follows. Subsection 1.1, which follows, discusses the related literature. Following that, in section 2 the baseline model is described in detail and then the information sharing and the voluntary information sharing variants of the model are articulated. Then, in section 3, computational details are described. A reader not concerned with computational details can skip this section and proceed directly to section 4. Section 4 discusses the numerical analysis, focusing on the competitive impact of information sharing. Section 5 concludes.

1.1 Related Literature

Our paper is closely related to the literature on the numerical analysis of dynamic oligopolistic games. For a survey of this literature see Doraszelski and Pakes (2007).¹⁴ Within this literature, the closest papers to ours are Saini (2013) and Jeziorski and

¹²See FTC/DoJ’s April 2000 *Antitrust Guidelines for Collaborations Among Competitors* at page 15.

¹³See the E.U. 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*.

¹⁴Recent 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.

Krasnokutskaya (2016). Both these papers apply the Markov Perfect equilibrium concept to auction settings, exploring the optimal procurement policy given capacity constrained suppliers and subcontracting, respectively.¹⁵ The framework applied in these papers dates back to Ericson and Pakes (1995). Our setting differs from the original Ericson and Pakes setting as our focus is on information asymmetry; hidden states in our case. Firms do not observe the state of their competitor when they need to make a decision (a bid). Our framework is therefore related to Fershtman and Pakes (2012) in which the numerical analysis is extended to games with incomplete information but the setup is restricted only to capital accumulation games. This restriction implies that the state of each firm is affected only by the firm's action and not by the actions of other firms. This assumption simplifies the analysis considerably. In our current setup the state of each firm is its timber inventory which is affected by the auction in each period and the outcome of the auction is determined by the bids of all the firms so the model cannot be characterized as a capital accumulation game. We use, however, the same equilibrium concept of restricted experience based equilibrium but an important contribution of this paper is a further refinement that we suggest for this equilibrium concept. We define a boundary consistency condition and demonstrate how to implement and test for it.

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

Our paper also relates to the empirical literature on bidding collusion. There are several approaches in the literature 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).¹⁶ Aoyagi (2003) considers collusion in repeated auction when bidders are allowed to communicate with each other before each auction. In our paper we do not have specific collusion in our setup but we examine collusive information exchange regarding the firms' inventories on the firms' participation and bidding behavior. The policy implications of our paper relates 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).

2 A Model of a Dynamic Auction

We consider a model in which there are n firms in the market and no entry and exit from the industry. Each of the firms can harvest and sell a portion of their stock of lumber each year at a fixed price. 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

¹⁵Both these papers build on Jofre-Bonet and Pesendorfer (2003).

¹⁶Baldwin, Marshall and Richard (1997) also focused on timber auctions.

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 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. There are a number of possible rationals for this and it keeps the information set finite.¹⁷

We begin with the timing of the events that occur within a period. Then we describe the overall structure of the game. Following that, 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 a first price sealed bid auction.
3. Firms observe the realization of their stochastic participation fee. We assume that $F_{it} \sim U[F_l, F_h]$. 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 $\{\underline{b}, 2\underline{b}, 3\underline{b}, \dots, \bar{b}\}$.
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 $p^w(b_i, b_{-i})$. The winning bid, the identity of the winner, and the participants in the auction become public information.
7. If there is information exchange it occurs 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 p_t , who won the auction at period t , denoted by i_t^w , and the winning bid b_t^* . We denote the new public information as $\xi_t^n \equiv [i_t^w, b_t^*, p_t]$. In a period of information exchange the new public information is $[i_t^w, \omega_t]$, the identity of the firm that won the auction and the observed state $\omega_t \equiv \{\omega_{i,t}\}$.¹⁸

¹⁷See Fershtman and Pakes, 2012, for a list of ways to keep the information set finite. Information revelation every T periods is convenient for us as it allows us to directly compare equilibria based on a sequence of larger information sets. We can justify our structure by; assuming that a regulator imposes mandatory periodic information revelation, or by the existence of a trade group that facilitates the sharing of information every T periods.

¹⁸Note that at a period of information revelation the winning bid and the participation decision of that 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, as these ω 's are revealed at that periods.

8. The winner discovers the amount of timber on the plot it won. This 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 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 agents and not observed by competitors.
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 $\min\{\omega_i + \mathbb{I}_{\{i=win\}}(\theta + \eta), e + \epsilon_i\}$, where $\mathbb{I}_{\{i=win\}}$ is an indicator function which takes the value of one if i wins the auction and zero elsewhere.¹⁹ The quantity harvested by firm i is not observable by other firms.²⁰ Note that if $b_i = \emptyset$ signifies no participation, a firm's expected profit, given $(b_i, b_{-i}, F_i, \omega_i)$, are

$$\pi^e(b_i, b_{-i}, F_i, \omega_i) =$$

$$\sum_{\eta, \epsilon_i} \left[p^w(b_i, b_{-i}) [\min\{\omega_i + (\theta + \eta), e + \epsilon_i\} - b_i] + [1 - p^w(b_i, b_{-i})] \min\{\omega_i, e + \epsilon_i\} \right] p(\epsilon) p(\eta) - \{b_i \neq \emptyset\} F_i.$$

10. All the firms updates their private ω_i .

Agents' Strategy Spaces.

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 which 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 even if no agents' strategy depended upon the variable some player can improve its expected discounted value of net cash flows by conditioning on it; for more details see Fershtman and Pakes (2012). That paper also shows that in models with periodic revelation of information there exists an equilibrium which 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.²¹

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 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 $\xi_t = \{\tau_t, \omega_{t-\tau_t}, \xi_{t-1}^n, \dots, \xi_{t-\tau_t}^n\}$.²² 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

¹⁹Here, and in what follows, we drop time subscripts, except where they add clarity.

²⁰Otherwise the observable harvested quantity may serve as a signal regarding ω_i .

²¹Equilibrium is defined formally in section 2.2.

²²Note 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 participants identity as these variables are not informationally relevant.

are made includes $\omega_{i,t}$ and $F_{i,t}$. However since $F_{i,t}$ is i.i.d., it enters the value function linearly, and does not have an independent effect on future values whereas the other state variables do. As a result it will be useful to have notation for $J_{i,t} = (\omega_{i,t}, \xi_t)$ separately from $I_{i,t} = (J_{i,t}, F_{i,t})$.

Strategies. There are two elements of a firm's strategy; the participation strategy and bidding strategy. We denote firm i strategy as $b(J_i, F_i) \rightarrow \{\mathcal{B} \cup \emptyset\}$ where $b = \emptyset$ signifies no participation.

2.1 The Dynamic System

We let $V(I_i)$ be the value of the game for a player i given his information set I_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 and we let i_w be the winning firm. Letting β be the discount factor, the firm's expectation of current period revenue (which excludes F_i) is

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

It follows that, for $b \in \mathcal{B}$,

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

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); i.e. $\omega'(\omega, \eta, \epsilon_i) = \max\{0, \omega_i - (e + \epsilon_i) + \theta + \eta\}$. 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 , i.e. $\omega'(\omega, \epsilon_i) = \max\{0, \omega_i - (e + \epsilon_i)\}$. $p(\xi'|\xi, \omega_i, b, i = i_w)$ is the probability distribution of future public information given the current public information ξ , the firm's private information ω_i and the identity of the firm winning the auction with bid b . Similarly, $p(\xi'|\xi, \omega_i, b, i \neq i_w)$ is the probability distribution of future public information given that the firm loses the auction.

Lastly the continuation value when the firm chooses not to participate in the auction, our $W(\emptyset|J_i)$, is

$$W(\emptyset|J_i) = \pi^e(\emptyset|J_i) + \beta \sum_{\epsilon_i, \xi', F'_i} V\left(\omega'(\omega, \epsilon_i), \xi', F'_i\right) p(\xi'|\xi, \omega_i, b = \emptyset) p(F'_i) p(\epsilon_i)$$

where $p(\xi'|\xi, \omega_i, b = \emptyset)$ is the probability distribution of future public information given the current public information, the firm's private information, and the choice of not participating in the auction.

2.2 The Restricted Experience Based Equilibrium

We now derive the conditions of a restricted experience based equilibrium for this game (see Fershtman and Pakes 2012). We let s be the set consisting of the payoff and informationally relevant states of all the firms, that is $s = (J_1, \dots, J_n)$ when all the J_i have the same public component ξ . We will say that $J_i = (\omega_i, \xi)$ is a component of s if it contains the information set of one of the firms whose information is combined in s . Note that we can also write $s = (\omega_1, \dots, \omega_n, \xi)$ and define the set of possible states $\mathcal{S} = \{s : (\omega_1, \dots, \omega_n) \in \Omega^n(\omega), \xi \in \Omega(\xi)\}$.

Definition of a REBE: A restricted experience based equilibria consists of the following three objects.

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 each firm and for every J_i which is a component of any $s \in \mathcal{S}$ and $F_i \in [F_l, F_h]$.
3. A set of numbers $\mathcal{W} \equiv \{W^*(b|J_i)_{b \in \mathcal{B} \cup \emptyset}\}$ that, for every J_i that is a component of any $s \in \mathcal{S}$, 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 $s_0 \in \mathcal{R}$, and the transition kernel generated by $\{b^*(J_i, F_i)\}_{J_i \in s \in \mathcal{S}, F_i \in [F_l, F_h]}^{i=1, \dots, n}$ has \mathcal{R} as a recurrent class; that is, with probability one, any subgame starting from an $s_0 \in \mathcal{R}$ will generate sample paths that are within \mathcal{R} forever.

C2: Optimality of strategies. Conditional on $\mathcal{W} \equiv \{W^*(b|J_i)_{b \in \mathcal{B} \cup \emptyset}\}_{J_i \in s, s \in \mathcal{S}}$, 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 $s \in \mathcal{R}$ equals the expected discounted value of returns generated by that choice from that J_i ; where expectations are taken using the empirical distribution of outcomes from that J_i (empirical distributions are denoted by a superscript E). Formally for every $b \in \mathcal{B} \cup \emptyset$, $W^*(b|J_i)$, the equilibrium evaluations satisfy

$$W^*(b|J_i) = \pi^E(b|J_i) + \beta \sum_{\epsilon, \eta, J_{-i}, F_i} V(\omega', \xi', F_i) \mu^E(\xi'|\xi, J_i, w, b, J_{-i}) \mu^E(J_{-i}|J_i) p(\eta) p(\epsilon) p(F_i) \quad (4)$$

where

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

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

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

and

$$\omega' = \left(\{i = i_w\} \max[0, \omega_i - (e + \epsilon_i) + \theta + \eta] + \{i \neq i_w\} \max[0, \omega_i - (e + \epsilon_i)] \right)$$

while

$$\begin{aligned} \mu^E(J_{-i}|J_i) &\equiv \frac{\mu^E(J_{-i}, J_i)}{\mu^E(J_i)}, \\ \mu^E(\xi'|\xi, J_i, b, J_{-i}) &\equiv \frac{\mu^E(\xi', \xi, J_i, b, J_{-i})}{\mu^E(\xi, J_i, b, J_{-i})} \quad \spadesuit. \end{aligned}$$

As noted in Fershtman and Pakes (2012) any Markov Perfect Bayes equilibrium will satisfy the conditions of a REBE. In fact a REBE admits more equilibria than does Markov Perfect Bayes. To understand the main reason why, it is helpful to distinguish between two types of points in the recurrent class; interior points and boundary points.

At an interior point an agent will stay in the recurrent class with probability one regardless of which of the feasible policies is chosen. At a boundary point the agent will stay in the recurrent class with probability one if the equilibrium policy is chosen. The agent may move outside of R if a feasible but non-equilibrium policy is chosen. In a restricted experienced based equilibrium the perceived discounted value of all feasible policies from an interior point equals the actual expected discounted value that would arise from all agents playing their equilibrium policies. However at boundary points only the perception of returns from the policies that lead to points in \mathcal{R} with probability one are required to equal the actual discounted values were all agents to play their optimal strategies. Policies that lead to points outside of the recurrent class are determined solely by perceptions and different perceptions on boundary points can support different equilibria.

There are situations where it might be reasonable to impose restrictions on off the equilibrium path behavior at boundary points. This would restrict the set of equilibria further. We consider one such restriction in the next section. The reader who is not interested in this refinement should be able to go directly to section 2.4.

2.3 Strengthening REBE: Boundary Consistency

If agents 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 perceived value of off-equilibrium-path play from a boundary point equals the expected discounted value of profits from that point when all agents use their equilibrium policies

(note that those policies are defined on all of \mathcal{S}). We call this a boundary consistency condition as it, together with condition C2, ensures that off the equilibrium path play at boundary points would lead to discounted values that are less than those of optimal play. Note that to impose this condition we need 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 which could be taken from points in the recurrent class that would generate outcomes which are not in the recurrent class. To this end let $\text{supp}[p_{s'}(\cdot|b_i, b_{-i}^*, s)]$ be the support of the probability distribution over next period states, generated by actions (b_i, b_{-i}^*) and initial state $s = (J_i, J_{-i})$. The boundary set of couples (b, s) , which we denote by B , are the set of action-state combinations such that if $s = (J_i, J_{-i}) \in \mathcal{R}$, action b is taken by i and equilibrium actions are taken by the other agents, then a probability distribution for s' is generated which has a point in its support which is not in the recurrent class, or

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

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

The additional condition that needs to be satisfied for the one-period deviation to actually yield an outcome which is less than the value of optimal play is C4 below. In this condition we use γ to index periods since the off-equilibrium-path policy is played. Let $F = (F_i, F_{-i})$. The probability distribution $p(s_\gamma|b, s, \{F_\tau\}_{\tau=1}^\gamma)$ is derived recursively, with $p(s_1|b, b^*, s) = \sum_{F_{-i}} p(s_1|b_i = b, b_{-i} = b^*(J_{-i}, F_{-i}), s)p(F_{-i})$, and for $\gamma > 1$, $p(s_\gamma|b, s_{\gamma-1}) = \sum_F p(s_\gamma|s_{\gamma-1}, b^*, F)p(F)$.

C4:Boundary Consistency. Let $\pi_i(b^*, s, F) \equiv \pi(b_i^*(J_i, F_i), b_{-i}^*(F_{-i}, J_{-i}), F_i, J_i)$ and $\pi_i(b, b_{-i}^*, s, F) \equiv \pi(b, b_{-i}^*(F_{-i}, J_{-i}), F_i, J_i)$. Then our condition is $\forall (b, J_i)$ component of $(b, s) \in B$ and for every F_i ,

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

where $p(s_\gamma|s_{\gamma-1}, b^*, F_\gamma)$ is the probability of reaching state s_γ at time γ given that at time $\gamma - 1$ the state is $s_{\gamma-1}$, participation fees are F_γ and the players play the equilibrium strategies b^* .

Definition. We call an equilibrium which satisfies C1 to C4 a ‘‘Boundary Consistent’’ *REBE*.

Notice that if for any sample path (i.e. any $\{s_\gamma\}_{\gamma=1}^\infty$), we define $\gamma_R = \min_\gamma \{\gamma : (s_\gamma) \in \mathcal{R}\}$, we can replace

$$\sum_{\gamma=\gamma_R}^{\infty} \beta^\gamma \sum_{s_\gamma, F_i} \pi_i(b^*, F_i, s_\gamma) p(s_\gamma|s_{\gamma-1}, b^*, F) p(F_i)$$

in C4 with $\beta^{\gamma R} \sum_{F_i} V(s_{\gamma R}, F_i)p(F_i)$. 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 $\beta^{\gamma R} \sum_{F_i} V(s_{\gamma R}, F_i)p(F_i)$ eases the burden of computing of the test.

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 to two models that allow 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 which facilitates coordination among firms. We also assume that when information is exchanged firms reveal their true state.²³

2.4.1 Information Exchange (IE)

In the first information sharing model there is mandatory information exchange every period.²⁴ 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 in which there is a forced information exchange firms also make a decision on whether to share information in every period for the next $T - 1$ periods hence. If one of the firms does not wish to share information, there is no voluntary information sharing over the next $T - 1$ periods, and in the T^{th} period firms' chose whether they wish to share information in the subsequent T periods. We call this model the VIE model and describe it in more detail now.

We have a period index $\tau = 0, 1, \dots, T$, which designates the time from the period of mandatory information exchange. 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 decides whether to participate in the auction. Participation is costly, it requires an expenditure of $F_{i,t}$ (a draw from the uniform distribution). If they decide to participate they simultaneously submit their bids and decide whether

²³Truthful revelation may require careful design of the incentives surrounding the agreement. For an exploration of this 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).

²⁴Formally we compute the model already described with the constraint that $T = 1$.

to reveal information over the next T periods. If both firms agree to reveal information, there is information exchange over the next T periods and the voluntary 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 (5) with

- “Information exchange occurs at this point. If $R = 1$, $\omega_{i,t}$ of all the firms is revealed. This is 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 $\xi_t^n = [i_t^w, b_t^*, p_t]$.”

In the *VIE* game the agents’ information set is different than in the *B* game in that the public information also includes the most recent information sharing indicator, or $R \in \{0, 1\}$.

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 ($\tilde{R} = 0$) and define

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

where $W(b, \tilde{R} = 1|J_i)$ ($W(b, \tilde{R} = 0|J_i)$) is 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 its information set.

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 chose 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, i.e. $R \in \{0, 1\}$. At $\tau = 0$ the optimal policies are given by

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

Finally since the agent needs a perception of the probability that $R = 1$ when he evaluates the returns from $\tilde{R} = 1$, there is a an additional consistency requirement that

the perception of this probability is, in equilibrium, equal to its empirical probability, or

$$\mu^E(R = 1 | J_i, \tilde{R} = 1) \equiv \frac{\mu^E(R = 1, J_i, \tilde{R} = 1)}{\mu^E(J_i, \tilde{R} = 1)}.$$

Before going to our results we explain the computational algorithm we used to obtain them. A reader who is not interested in the computational algorithm should be able to go straight to section 4.

3 Computation, relationship to learning, 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 is 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 which maximizes the perception of the value from that state. This process continues iteratively.

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

We could perturb the algorithm to maintain the analogy with learning by forcing agents to experiment with different policies at each state (as in Fudenberg and Levine (1998)). This would, however, increase the complexity of the algorithm. A less computationally burdensome way of proceeding to compute a REBE is to use knowledge that the computer has in its memory but the agent does not have to update the values associated with all policies (even those the agent does not take). Indeed from a computational point of view the fact that we can compute an equilibrium for a non-capital

accumulation game without explicitly calculating the impact of one firm's policies on the evolution of its competitors' states is an advantage of our algorithm relative to algorithms which require explicit computation of all continuation values (see, for e.g., Besanko et al. (2014)).

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 is consistent with the stronger notion of equilibrium that 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) $L^k = [J_1^k, \dots, J_n^k]$, and has objects in memory, $M^k = \{M^k(J) : J \in s \in S\}$. 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. A more detailed discussion of increasing the efficiency of the algorithm is provided in the results section.

Memory: The elements of $M^k(J)$ 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 with different ways having different advantages. Here we focus on a simple structure that will always be available (though not necessarily always efficient; see Fershtman and Pakes, 2012).

$M^k(J)$ contains a counter, $h^k(J)$, which keeps track of the number of times we have visited J prior to iteration k . If $h^k(J) > 0$ it also contains

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

If $h^k(J) = 0$ there is nothing in memory at location J . When we need to evaluate policies at a J at which $h^k(J) = 0$ we have an initiation procedure which sets

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

The choice of initial values will be discussed below.

Updating L^k : We find the values in memory associated with different b for each agent at location L^k (or use the initiation procedure if needed), 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 $b^k \neq \emptyset$ there is an auction. We assume that if there is an auction and more than one firm bids b^k there is a lottery that determines the winning bid.

The b^k , the identity of the winner (i_w^k), and the participation decisions of all agents (the vector p^k) enable us to update the public information sets as

$$\xi^{k+1} = \{\tau_k = 0\} \left(\omega^k, \tau_{k+1} = 1, i_w^k \right) + \{\tau_k \neq 0\} \left(\xi^k(\tau^k + 1), p^k, i_w^k, b_w^k \right),$$

where $\xi^k(\tau^k + 1)$ is 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 (p^k, i_w^k, b_w^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 agents sell what they can sell to the market. The random draws from the harvest (η) and from the market sale (ϵ_i for each i) are realized and each agent's stock of timber is augmented as

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

Thus the information prior to the next auction is given by

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

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 agents' perceptions of the values associated with the different policies. We only update objects in memory associated with the location L^k , but we update each component of $\{W^k(b|J_i^k)\}_{b \in \mathcal{B} \cup \emptyset}$ for all i . That is 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 agent 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.

The update of the expected value from pursuing strategy b at state J_i^k , i.e. $W^k(b|J_i^k)$, is obtained by assuming 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 formula. This 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

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

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 agent taken that action every time in the past that it had reached J_k^i . Though this averaging procedure does satisfy the Robbins and Monroe (1951) criteria for convergence of a stochastic integral, it is unlikely to be efficient. This because 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.

3.2 Testing Procedures.

Appendix A provide 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 which differ from the test used 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.

We stop the algorithm at a particular iteration and fix the values for that iteration.

$$\left(\{W(b|J_i)\}_{b \in B}, W(\emptyset|J_i) \right) \forall J_i \text{ at } \left(\{W^*(b|J_i)\}_{b \in B}, W^*(\emptyset|J_i) \right),$$

The test is designed to check whether these values, together with the policies and the recurrent class that they generate, satisfy conditions C1 to C3 above.

The test is based on simulating a sample path with the optimal policies generated by $\left(\{W^*(b|J_i)\}_{b \in B}, W^*(\emptyset|J_i) \right)$. Since 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 separate memory for each point visited in the test's simulation run.

The first time a particular point is visited we record the simulated continuation value resulting from taking every possible action at that point. I.e. the profits plus the discounted continuation value (evaluated by $\{W^*(\cdot|\cdot)\}$) generated by; their action, the policy chosen by their competitors and simulated random draws on the primitives.²⁵ We also record the square of this continuation value and initiate a counter for the amount 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 continuation value for playing policy b at J_i is the mean square error of our estimate of $W^*(b|J_i)$. It can

²⁵Since the stage game is simultaneous move, we can evaluate a counterfactual choice of a given agent's policy by substituting it, and the optimal policies of competitors, into this calculation.

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 estimate squared. So by differencing the mean square 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.

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 recurrent class 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 the Appendix.

3.2.2 Testing for Boundary Consistency or for C4.

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

We then start new simulation runs from every point visited in this simulation run for every possible policy from that point. This 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 this is added to 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 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 looking to see if any of the simulated runs starting at J_i with policy b had a simulation run which did 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 we run after determining our estimate of the recurrent class. At each point, say J_i , chose every $b \in B \cup \emptyset$ and, using the policies generated by $\{W(\cdot|\cdot)\}$, start R simulation runs. Index the runs from each (J_i, b) couple by r and let the sequence of states visited during the r^{th} simulation run be $\{J_{i,\gamma_r}\}_{\gamma_r=1}^{\gamma_r^*}$, 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

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

$$- \sum_{\gamma_r=1}^{\gamma_r^*-1} \beta^{\gamma_r} \{b(J_{i,\gamma_r}, F_{i,\gamma_r}) \neq \emptyset\} F_{i,\gamma_r} + \beta^{\gamma_r^*} W(b^*|J_{i,\gamma_r^*}),$$

where it is understood that $b(J_{i,1}, F_{i,1}) = b$ or the policy we are evaluating. We keep in memory the average of the $\hat{W}_r(b|J_i)$, 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 it is understood that $\gamma_r^*(b, J_i)$ 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 $\overline{W}^R(b|J_i)$ which is an unbiased estimate of the population mean from R sample paths (in our case $R = 20$), and 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. Since we are testing inequalities, i.e. that the boundary point policies lead to discounted values of future net cash flows which are less than the optimal policy at the J_i associated with the boundary point, we will have to use a test statistic which is not pivotal, i.e. 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 test 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 $(b, J_i) \in \mathcal{B}$, thereby insuring the size of the test.²⁶ For each $(b, J_i) \in$

²⁶The 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) or Andrews and Pakes (2016).

\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

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

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

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

if and only if

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

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 resulting computational burden. An equilibrium is computed in 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. In each model, there are two firms (bidders) and four possible bids. This structure is adopted to limit the size of the state space, such that computation is feasible (the computational burden will be discussed in the next subsection). Similarly, the time between forced revelation periods in the baseline model is 4 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. Participation costs are assumed to be 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. Thus, participation costs, in this setting, are economically meaningful.²⁷

4.2 Computational burden and updating procedure

A restricted EBE is computed using the algorithm provided in section 3.1. Recall that there may be many equilibria that satisfy our equilibrium conditions. The choice of

²⁷Note that harvesting and production costs are normalized to zero.

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

initial conditions for continuation values (our $\{W^0(\cdot)\}$) is one determinant of which equilibria the algorithm will compute. If the initial conditions are higher than possible equilibrium values then all policies are likely to be explored, and, as a result, any equilibrium the algorithm converges to is likely to be boundary consistent. The cost of choosing high initial conditions is that they are likely to cause the algorithm to require many iterations before it converges to equilibrium values.

We incurred that cost and used as initial conditions

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

for all $(b, J_i) \in (\mathcal{B}, \mathcal{J})$. To see why we chose these initial values, 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 that the firm would have to win the auction in order to have the timber needed to sell e units in every period. So $(F + .5)e/[(\theta + 1)(1 - \beta)]$ 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 $\omega(F + .5)/(\theta + 1)$ 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 the recurrent classes were (less than) an order of magnitude smaller than this (less than 330,000), there was a significant computational burden in finding them. 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 2089 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 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.

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 one above is run for 10 million iterations and it took only one round of our steps before convergence. The boundary consistency test was run, as described in section 3.2.2, after we accepted the test for the Restricted EBE.²⁸ All the equilibria we describe here were boundary consistent, though we did find one that was not which we do not report on. A summary of compute times is provided in the bottom half of table 2 and the footnote to the table describes the program and computer used for the runs.²⁹

²⁸The number of simulation runs used to determine whether a point in the recurrent class was a boundary point was fifty, and the number of repetitions to form the averages used in the test of the boundary points was twenty.

²⁹The total computation times, including testing, for each of the models, were (in hours): B - 110, IE - 4.5, VIE - 185. There are many, likely quite helpful, ways one might improve on this, but optimizing the algorithm is beyond the scope of this paper.

Table 2: Computational details

Size of recurrent class:		
<i>B</i>	<i>IE</i>	<i>VIE</i>
325,843	2,089	328,688
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 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.³⁰ Were it not for the existence of a non-zero minimum bid, which distorts participation somewhat, the planner’s allocation problem would be equivalent to that of the ideal, perfectly coordinated, cartel; the planner maximizes the discounted value of the sum of future net cash flows.

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 both phenomena generate lower bids.

On the other hand static auction theory implies that increased participation signifies more competition which should, in turn, lead to higher bids; and there is more participation in the *IE* than in the *B* equilibrium. Part of the participation difference

³⁰Specifically, 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 assume that the planner has access to the F_i and ω_i realizations of both firms.

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 conditional on ≥ 1 firm participating	1.16	0.98	1.12	-
Avg. winning bid conditional on 1 firm participating	1.06	0.67	0.99	-
Avg. winning bid conditional on 2 firms participating	1.23	1.16	1.20	-
Avg. # of participants	1.52	1.63	1.52	1
Avg. # of participants, conditional on ≥ 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 in which a firm with the lowest omega wins conditional on ≥ 1 firm participating	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 $\pi(\omega_i) - \mathbb{I}_{\{i=win\}}b_i - \{b_i \neq \emptyset\}F_i = \min\{\omega_i + \mathbb{I}_{\{i=win\}}(\theta + \eta), e + \epsilon_i\} - \mathbb{I}_{\{i=win\}}b_i - \{b_i \neq \emptyset\}F_i$. Total revenue is defined as $\sum_i \pi(\omega_i) = \sum_i \min\{\omega_i + \mathbb{I}_{\{i=win\}}(\theta + \eta), e + \epsilon_i\}$. 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.

might be attributed to the more detailed information structure in the IE equilibrium facilitating more coordinated bids, as there are less periods in the IE equilibrium when neither firm bids (.015 vs .04 percent). However, the statement that more information leads to softer competition seems to be clearly at odds with the relationship between bids and participation in the periods with at least one bidder, as even in those periods there is more participation in the IE than the B equilibrium (1.63 vs 1.59).

Of course what might be confusing differences in behavior in a model of a static (or a repeated) game, might not be confusing in the context of a dynamic game. In particular differing dynamic incentives will generate differences in the propensity to hold different stocks of lumber. We expect participation and bidding to differ with differences in those stocks, and the table's comparisons between the IE and B outcomes are comparing different weighted averages of the stock combinations. The probable role of dynamics in explaining differences in the implications of the information environment also comes out clearly when we compare Table 3 to Table 4. Table 3 indicates that more information (the IE equilibrium) generates a higher discounted cash flow and therefore higher average profits, but table 4 makes it clear that once we condition on the stock of timber the B equilibrium generates higher profits almost always.³¹

Before leaving table 3 we note that all three models deliver (essentially) the same social surplus (albeit with *IE* being 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 -

³¹The only exception are states which are visited only .15% (1.12%) of the periods in the B (IE) equilibrium.

tuple). In the *IE* equilibrium the firms generate almost the same revenue (equivalently, output) per period as does the planner, but require much greater participation to do so, thus generating a lower social surplus. By contrast, firms in *B* are less effective at revenue generation (their stocks are not always high enough to satisfy the demand that faces them), but generate less wasteful participation.³²

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

(ω_i, ω_{-i})	Prob. Dist. (%)			Profit	
	<i>B</i>	<i>IE</i>	<i>SP</i>	<i>B</i>	<i>IE</i>
$(\leq 4, \leq 4)$	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, and 7, 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 4.3.1). 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.

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 $\omega \leq 4$ is 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*.

³²All the effects described in the preceding paragraphs become much more muted when the models are computed with $\beta = 0.8$. For instance, 'Avg. bid' across the three models (*B*, *IE* and *VIE* in order) is 0.82, 0.82 and 0.80; 'Avg. # of participants' is 1.45, 1.46 and 1.46 and social surplus is 2.77, 2.82 and 2.77 (as compared to 3.07 in *SP*). All of which suggests that continuation values matter for the observed conduct.

In contrast the social planner spends more time in the $(\leq 4, \leq 4)$ -tuples than either 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 bids by intervals of ω -tuples for B and IE . \emptyset indicates non-participation.

Table 5 contains the probability distributions over bids that underlie the distribution over the ω -tuples examined in table 4 together with average profits in those states. 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; there is both more participation in IE and a higher fraction of bids are higher than the minimal bid in these states. This reinforces the impression that the increased information created when moving from B to IE is not allowing the firms in IE to better coordinate; 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 and this softens competition.

The opposite seems to be true when at least one of the firm's has an ω greater than eight, or both firms have an ω between five and seven. In these states participation in IE is sometimes greater than in B but, 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 when their F_i draw is sufficiently low, and then bidding the minimal amount. 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 only 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. In the $(\geq 8, 5 - 7)$ -tuple and the $(\geq 8, \leq 4)$ 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. In contrast when the B equilibrium is at the tuple $(\geq 8, 5 - 7)$ the high ω firm bids in 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).

So when at least one of the firms has an ω greater than eight, or both firms have an ω between five and seven, it seems that more information enables better co-ordination of bids. The one couple of states in table 5 that we have not discussed is when one firm has an ω less than or equal to four and the other has an ω between five and seven. There is a sense in which this couple 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 this state 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, as it underlies the fact that the IE model typically generates disproportionate number of states where at least one firms has a high ω stock.

Tables 6 and 7 examine the differences in bids between the B and IE model in more detail. Table 6 looks at bids in the low ω states and shows the rather dramatic 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 chose the maximal bid (versus 28% in B). The differences between the bids in IE and B are similar in state $(1,1)$. Even when there is some asymmetry in the states, as long 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. For example at $(2,0)$, the firm with $\omega = 2$ participates 95% of the time in IE (versus 72% of the time in B) and the IE firm bids 1.5 or more 91% of the time (versus 64% of the time in B).

Table 6: Competition in low ω -tuples

(ω_i, ω_{-i})	Prob. Dist. (%)		Bids										Profit	
	<i>B</i>	<i>IE</i>	<i>B</i>					<i>IE</i>					<i>B</i>	<i>IE</i>
	\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*.

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 bidding in states in which 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 increases as their state moves from 0 to 1 to 2. By contrast, at least for the couples (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. So the low ω firm is likely to win more often, and if it does win, it wins with the minimal bid. This insures that both firms profits increase as the low ω firm's state 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, so that by the time the state (7,7) is reached each firm either does not participate or bids the minimal amount (in about equal proportions). 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) less profitable in the B than in the IE equilibrium.

Table 7: Bidding and participation in asymmetric ω -tuples

(ω_i, ω_{-i})	Prob. Dist. (%)		Bids										Profit	
	<i>B</i>	<i>IE</i>	<i>B</i>					<i>IE</i>					<i>B</i>	<i>IE</i>
			\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*.

Tables 6 and 7 are central to understanding how increasing a firm's information about its competitor changes the path of play. Providing more information about a competitor increases competition at low ω states which reduces profits in those states. In a static game a fall in profits that accompanies the increase in information would decrease participation. However here, despite the fact that profits are lower, participation is higher in the game with more information. This because firms respond to the possibility of higher future profits if they increase their stock of timber, and an increase in information at low ω states intensifies the competition over those future profits. Consider two firms initially at low ω states. A firm that does win the initial auctions proceeds to a higher ω state, and then participates less often in subsequent auctions. Compared to the B equilibrium, firms in the IE equilibrium are better able to asses when their competitor has a large stock. So a firm that loses the initial auctions is more certain of the extent to which the winning firm's stock increases and knows that when the increase is large enough its competitor is less likely to participate in

the auction. Thus the firm with a low ω knows that it is likely to win subsequent auctions with a minimal bid and bids accordingly. Though this 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.

To summarize, in a static model the intensified competition at low ω states caused by being more certain that your competitor is at a low ω state would induce firms to stay out of the market. Here, participation increases because the static incentive is dominated by; (i) the dynamic incentive to move to a more profitable, i.e. a higher, ω state, and (ii) the impact of the increased information on bidding thereafter, and through those bids, on the profits of the firm that does not initially pull ahead.

For a more formal look at what underlies the dynamic incentives consider the value function for a firm which choses $b \in \mathcal{B}$, (that is, equation 3), reproduced here as

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

Note first that if $\omega > 2$ the only motivation to bid comes via the continuation value since when $\omega > 2$ the timber won in the auction will not be harvested in the current period. Then the value of participating in the auction comes solely through the anticipation of profit in the future. Second, note that dynamic incentives enter through both the expected profits associated with the increase in ω_i (in the event of winning) and the evolution of ξ , the public information. The latter differs between the two models and, for all combinations of timber stocks, the firms in the IE model always have more precise information on the competitors' state than in the B model.

With this in mind one way of measuring the difference in incentives between the two models at a state and realization of F_i can be obtained by comparing the continuation value of the optimal bid in the two models to the continuation of the “statically optimal” bid, where the latter is defined as the optimal bid at the state ignoring the future (given that the rival is playing the optimal strategy of the dynamic model). More formally the statically optimal bid is $b_s^*(J_i, F_i) = \max_{b \in \mathcal{B} \cup \emptyset} \{\pi^e(b|J_i) - \{b \in \mathcal{B}\}F_i\}$. The optimal bid is $b^*(J_i, F_i) = \max_{b \in \mathcal{B} \cup \emptyset} \{W(b|J_i) - \{b \in \mathcal{B}\}F_i\}$ and the measure of the extent of dynamic incentives is $D(b^*, J_i, F_i)$ where

$$\begin{aligned}
D(b^*, J_i, F_i) &= [W(b^*(J_i, F_i)|J_i) - \pi^e(b^*(J_i, F_i)|J_i, p^w)] \\
&\quad - [W(b_s^*(J_i, F_i)|J_i) - \pi^e(b_s^*(J_i, F_i)|J_i, p^w)].
\end{aligned}$$

Figure 1 shows, for *IE* and *B*, the expectation of $D(b^*, J_i, F_i)$ by ω_i and the rival's ω , that is $E_{J_i, F_i} [D(b^*, J_i, F_i)|\omega_i, \omega_{-i}]$. In all three panels of figure 1 the solid dark line (*IE*) lies above the dashed line (*B*), but for i) $\omega = 0$, where the two lines coincide, and ii) at high ω 's when facing an intermediate- ω opponent (panel (b)). This suggests that, for most states, increasing information increases the extent to which incentives driving firm conduct are derived from dynamic considerations. The largest difference between dynamic and static incentives are for bidders with $\omega \in \{3, 4, 5\}$ facing rivals

with $\omega \geq 5$ (panels (b) and (c)). If the bidder wins at these states it is more likely to transition to that part of the state space where play softens and firms (roughly) divide their actions between not participating or bidding the minimum amount (see the prior tables). Thus, for bidders with $\omega \in \{3, 4, 5\}$, whose competitor's ω is five or more, the future with a winning bid is extremely attractive. The increased information in the IE equilibrium allows firms to have more precise information on when this situation is likely to occur, and hence to take better advantage of it.

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; and so colors competition throughout the recurrent class.³³ 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. Somewhat perversely this occurs precisely through the intensification of competition caused by a reduction in asymmetric information in those states (low- ω -tuples) in which competition was most vigorous to start with.

4.3.1 Voluntary information exchange (VIE)

In the VIE model firms can elect, every 4 periods, to share information. If both firms elect to share information then the model switches, for the next four periods, from the B to the IE setting. After the four periods they chose between B and IE again. If one or both firms choose not to share, then firms spend the next four periods in the B setting.

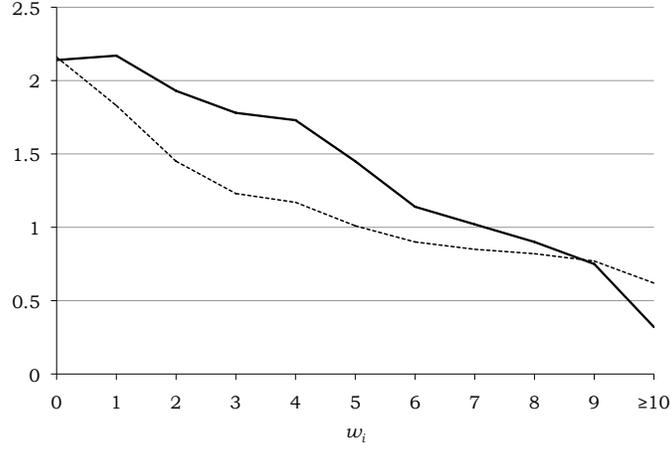
In the discussion above, the VIE model was largely ignored, other than noting that it did not generate outcomes meaningfully different from the B model. Despite the fact that average profits in IE are larger than average profits in B , firms in VIE only choose to share information in 5% of the states where that choice is made (i.e. firms choose to share information in 24% of those states but both firms choose to share information in only 5%). As a result, replacing IE with VIE in the analogs of tables 4, 5, 6 and 7 show little difference between B and VIE . The question then is why firms in VIE cannot reliably coordinate on sharing information, given that it appears to be in their long term interest.

Table 8 shows the propensity to elect to share information by ω -tuple. It shows that the propensity to share information is substantial only when both ω 's are greater than 4, and the highest is greater than 8. In VIE these states occur relatively rarely, due to the default being to not share information (and hence play according to the pattern in B). Hence, the states in which electing to share is highest are rarely visited, and the frequency of choosing to share information stays low at 5%.

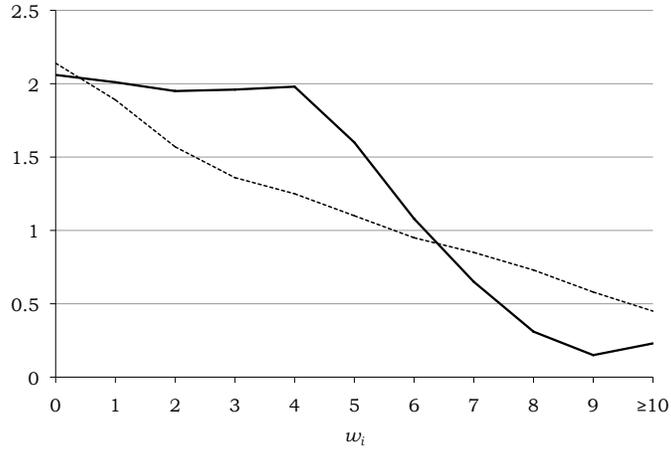
Though, as Table 3 shows, average profit are higher in IE , for a given ω pair profits are higher in the B equilibrium. The reason for this difference is that the path of play generated by the IE equilibrium places more weight on the high ω states, and these are the more profitable states in both equilibria. Notice, however, that since for a given state profits are higher in the B equilibrium, there is a cost of switching from the B to the IE regime in that the switch generates lower current and near future profits.

This tradeoff comes out clearly in the comparison presented in Table 9. 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

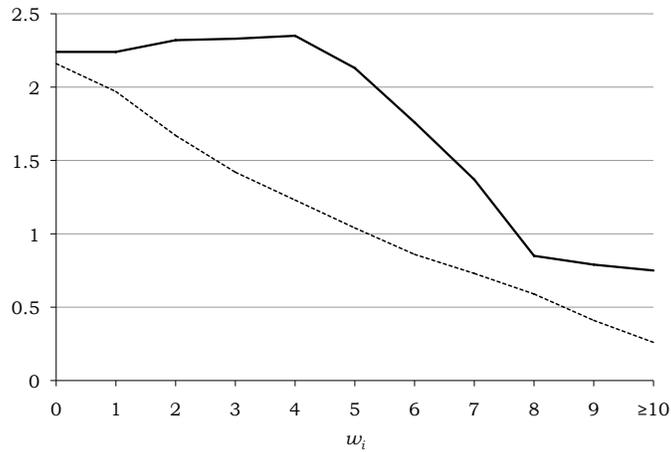
³³Recall that the recurrent class are those states visited repeatedly in the course of equilibrium play.



(a) $E_{J_i, F_i} [D(b^*, J_i, F_i) | \omega_i, \omega_{-i} \leq 4]$



(b) $E_{J_i, F_i} [D(b^*, J_i, F_i) | \omega_i, \omega_{-i} \in \{5, 6, 7\}]$



(c) $E_{J_i, F_i} [D(b^*, J_i, F_i) | \omega_i, \omega_{-i} \geq 8]$

Figure 1: $E_{J_i, F_i} [D(b^*, J_i, F_i) | \omega_i, \omega_{-i}]$ for different ω_{-i} sets. The solid line is IE . The dashed line is B . Values are a probability weighted average where the probability is the frequency with which each state is hit during a 5 million iteration simulation

Table 8: Individual firm’s choices to reveal by ω -tuple

(ω_i, ω_{-i})	Prob. Dist. (%)	Pr(Choose to reveal _i = YES) (%)	Pr(Both Reveal= YES) (%)	Profit	
	<i>VIE</i>	<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: This table shows the probability of an individual firm’s choices to reveal by ω -tuple in the *VIE* model. Only periods in which firms decide on information sharing (or periods with $\tau = 0$) are used in the calculation.

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 four periods and then shifting to information sharing. The last column reports the frequency, in the simulated data, with which the value for *IE* was larger than the calculation with four periods of waiting; i.e. 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.

The results in table 9 indicate that, for low- ω states, on average the expected value of waiting to switch to information sharing is greater than switching immediately. This changes once $\omega_i \geq 4$; then the immediate switch is preferred. This reflects the fact that information sharing induces fierce competition in low- ω states, with the returns to *IE* occurring as firms invest to avoid these states. When in a low- ω state a firm would rather wait and see if their ω transits to a higher level than elect to share information immediately and incur the lower profits associated with being in the *IE* regime. This changes once ω ’s are high enough, as then the returns from information sharing are more immediate. There are times when we start from the *B* regime and get to a decision period in which it is in the participants’ interest to switch to the *IE* regime. However even in the *IE* regime the firms’ periodically find themselves in low ω states, and if they do so in a decision period they tend to chose the *B* regime, and thereafter will be disproportionately in low ω states where typically the *B* regime is preferred.

The difficulty that firms, collectively, have in maintaining information sharing, despite it’s long term benefits, suggests the importance of commitment devices in establishing an effective information sharing arrangement. In *IE* perfect commitment is externally imposed. In *VIE*, firms are unable to commit for only four periods at a time

Table 9: $E_{F_i} [V (J_i, F_i) | \tau = 1]$ 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 % of times with (A) \geq (B) during a 1 million iteration simulation of the B model.

and this is sufficient to break down information sharing. The range of mechanisms that may be effective for committing to (truthful) information sharing is beyond the scope of the present model to unravel, but the VIE results presented here are sufficient to highlight the central role of commitment in determining the actual market impact of any informational sharing arrangement.

5 Conclusion

This paper presents a model that, when analyzed numerically, illustrates the difficulties in judging the competitive impact of information sharing. Information tends to intensify competition, at least in low- ω states, but the dynamic incentives in the model tend to push play toward higher ω states in which competition is less fierce, or (in the case of very high- ω states) outright reduced. As a consequence the auctioneer captures less surplus and the bidders do better. Importantly for market outcomes, the quantity produced (timber harvested) increases, indicating an increase in gains from trade (albeit one that is offset by the participation costs).

The example considered here is policy relevant, in that the information shared is clearly strategically relevant (prior ω is informative as to current ω , which in turn is highly informative as to the submitted bid), although not being directly about current bidding plans. In addition, the firms in the model, conditional on information, bid in their unilateral self interest and not according to some pre-arranged collusive mechanism. Thus, this model sits squarely in the thick grey line of ambiguity surrounding the proper treatment of information sharing arrangements. It highlights the sense in which a static view (prices in the auction dropped) can be in conflict with a dynamic view (inventory holdings, ω , increased moving competition to a different part of the state space). Taking into account the dynamics of the problem suggests that information sharing, at least in this setting, is less of a competition issue, and may well have significant pro-competitive aspects (note that output increased). This suggests that treating information sharing, even of strategically important data, as a *per se* offense (in the case of U.S.) or as a restriction of competition by object (in the case of the E.U.), needs to be carefully weighed against the serious possibility of type 1 error, falsely rejecting the hypothesis that conduct is pro-competitive, and incurring a reduction in welfare or output as a result.

Further, at a methodological level, this paper illustrates how the Experience Based Equilibrium concept facilitates investigation of the dynamics of complex auction environments with persistent, imperfectly correlated, asymmetric information. The dynamics that are accommodated are economically important. As demonstrated, states emerge in equilibrium that can only be reached when firms are responding to dynamic incentives. Further, the equilibrium concept is extended through the Boundary Consistency requirement, which mitigates the propensity for these types of dynamic games to exhibit multiplicity.

6 Appendix A

6.1 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 case.

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 storage we have particular values of $\left(\{W(b|J_i)\}_{b \in B}, W(0|J_i)\right)$, say $\left(\{W^*(b|J_i)\}_{b \in B}, W^*(0|J_i)\right)$, 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

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

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 together with 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, i.e. 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, labelled Υ , converges to an $L^2(P_{n_s}|W^*)$ norm in the percentage bias of the our estimates of W^* , where P_{n_s} 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

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

so that

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

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

6.2 Appendix B: NOT FOR PUBLICATION

Table 10: Probability Distribution and Actions by ω -tuple for B and IE

(ω_i, ω_{-i})	Prob. Dist. (%)		Bids										Profit		Joint Profit	
	B	IE	B						IE				B	IE	B	IE
	0	0.5	1	1.5	2	0	0.5	1	1.5	2	B	IE	B	IE		
(0, 0)	3.17	0.50	0.12	0.07	0.12	0.41	0.28	0.01	0.00	0.09	0.12	0.78	-0.22	-0.48	-0.43	-0.97
(0, 1)	3.70	0.88	0.12	0.08	0.13	0.46	0.20	0.04	0.00	0.09	0.44	0.43	-0.17	-0.44	0.24	-0.53
(0, 2)	4.91	1.48	0.11	0.09	0.17	0.49	0.15	0.05	0.08	0.05	0.60	0.23	-0.09	-0.31	0.92	0.35
(0, 3)	4.83	1.94	0.10	0.10	0.25	0.47	0.08	0.03	0.17	0.02	0.71	0.07	-0.02	-0.26	1.32	0.52
(0, 4)	3.83	2.27	0.08	0.13	0.39	0.38	0.03	0.02	0.19	0.17	0.62	0.00	0.04	-0.24	1.40	0.49
(0, 5)	3.02	2.47	0.07	0.15	0.53	0.24	0.02	0.02	0.20	0.38	0.40	0.00	0.14	-0.09	1.63	0.94
(0, 6)	2.19	2.48	0.06	0.18	0.61	0.13	0.02	0.02	0.19	0.61	0.18	0.00	0.19	-0.00	1.72	1.15
(0, 7)	1.49	2.36	0.05	0.23	0.61	0.09	0.03	0.01	0.33	0.62	0.03	0.00	0.22	0.02	1.76	1.19
(0, 8)	0.97	2.14	0.04	0.36	0.53	0.05	0.02	0.00	0.58	0.42	0.00	0.00	0.24	-0.01	1.78	1.13
(0, 9)	0.64	1.80	0.04	0.51	0.41	0.04	0.01	0.00	0.82	0.18	0.00	0.00	0.31	0.17	1.86	1.47
(0, 10)	0.41	1.29	0.04	0.64	0.27	0.04	0.01	0.00	1.00	0.00	0.00	0.00	0.38	0.69	1.99	2.44
(0, 11+)	0.53	1.21	0.03	0.78	0.12	0.05	0.02	0.00	1.00	0.00	0.00	0.00	0.50	0.81	2.22	2.67
(1, 0)	3.70	0.88	0.18	0.06	0.13	0.49	0.15	0.01	0.04	0.00	0.29	0.66	0.41	-0.08	0.24	-0.53
(1, 1)	2.36	0.80	0.18	0.12	0.23	0.40	0.07	0.03	0.09	0.00	0.74	0.15	0.46	0.20	0.93	0.39
(1, 2)	2.54	1.07	0.17	0.14	0.32	0.32	0.05	0.03	0.10	0.07	0.81	0.00	0.49	0.32	1.52	0.96
(1, 3)	2.09	1.16	0.15	0.16	0.43	0.23	0.02	0.02	0.13	0.33	0.53	0.00	0.55	0.32	1.89	1.19
(1, 4)	1.42	1.13	0.13	0.22	0.52	0.13	0.01	0.01	0.16	0.59	0.24	0.00	0.59	0.40	1.97	1.45
(1, 5)	0.98	1.08	0.11	0.29	0.51	0.08	0.01	0.00	0.28	0.72	0.00	0.00	0.62	0.52	2.10	1.65
(1, 6)	0.64	0.97	0.08	0.43	0.43	0.05	0.01	0.00	0.52	0.48	0.00	0.00	0.66	0.55	2.20	1.74
(1, 7)	0.40	0.83	0.08	0.50	0.38	0.03	0.01	0.00	0.79	0.21	0.00	0.00	0.69	0.64	2.26	1.85
(1, 8)	0.24	0.63	0.07	0.62	0.27	0.03	0.01	0.00	1.00	0.00	0.00	0.00	0.74	0.82	2.37	2.34
(1, 9)	0.14	0.40	0.08	0.68	0.19	0.03	0.02	0.00	1.00	0.00	0.00	0.00	0.75	0.89	2.40	2.63
(1, 10)	0.08	0.19	0.08	0.70	0.16	0.03	0.03	0.00	1.00	0.00	0.00	0.00	0.76	0.95	2.47	2.84
(1, 11+)	0.09	0.10	0.10	0.69	0.15	0.04	0.03	0.00	1.00	0.00	0.00	0.00	0.79	0.97	2.56	2.92
(2, 0)	4.91	1.48	0.28	0.07	0.19	0.41	0.05	0.05	0.10	0.00	0.86	0.00	1.01	0.66	0.92	0.35
(2, 1)	2.54	1.07	0.28	0.14	0.27	0.29	0.02	0.06	0.09	0.00	0.85	0.00	1.03	0.64	1.52	0.96
(2, 2)	2.57	1.32	0.26	0.17	0.34	0.22	0.01	0.04	0.18	0.11	0.66	0.00	1.01	0.62	2.02	1.24
(2, 3)	2.02	1.36	0.24	0.23	0.39	0.13	0.01	0.03	0.21	0.39	0.37	0.00	1.02	0.72	2.36	1.66
(2, 4)	1.33	1.26	0.21	0.32	0.40	0.06	0.01	0.01	0.29	0.65	0.05	0.00	1.04	0.86	2.40	1.94
(2, 5)	0.91	1.20	0.18	0.45	0.32	0.04	0.01	0.01	0.43	0.56	0.00	0.00	1.06	0.92	2.51	2.02
(2, 6)	0.58	1.06	0.15	0.57	0.24	0.03	0.01	0.01	0.73	0.27	0.00	0.00	1.06	0.97	2.59	2.27
(2, 7)	0.35	0.89	0.14	0.64	0.18	0.02	0.01	0.00	1.00	0.00	0.00	0.00	1.06	1.07	2.63	2.46
(2, 8)	0.22	0.62	0.14	0.69	0.14	0.02	0.01	0.00	1.00	0.00	0.00	0.00	1.08	1.03	2.71	2.75
(2, 9)	0.13	0.37	0.14	0.70	0.12	0.03	0.01	0.00	1.00	0.00	0.00	0.00	1.08	1.01	2.73	2.87
(2, 10)	0.07	0.17	0.13	0.72	0.11	0.03	0.01	0.00	1.00	0.00	0.00	0.00	1.06	1.01	2.78	2.92
(2, 11+)	0.07	0.09	0.17	0.68	0.11	0.03	0.01	0.00	1.00	0.00	0.00	0.00	1.08	1.00	2.87	2.97
(3, 0)	4.83	1.94	0.35	0.07	0.26	0.30	0.02	0.06	0.03	0.07	0.83	0.00	1.34	0.78	1.32	0.52
(3, 1)	2.09	1.16	0.34	0.16	0.31	0.16	0.02	0.12	0.10	0.23	0.55	0.00	1.34	0.87	1.89	1.19
(3, 2)	2.02	1.36	0.33	0.22	0.33	0.11	0.02	0.12	0.13	0.39	0.35	0.00	1.33	0.93	2.36	1.66
(3, 3)	1.54	1.34	0.31	0.30	0.32	0.06	0.01	0.11	0.20	0.57	0.12	0.00	1.34	1.03	2.68	2.06
(3, 4)	0.97	1.22	0.28	0.40	0.28	0.04	0.00	0.07	0.43	0.50	0.00	0.00	1.35	1.20	2.72	2.20
(3, 5)	0.65	1.17	0.25	0.48	0.23	0.03	0.00	0.04	0.67	0.29	0.00	0.00	1.33	1.19	2.79	2.35
(3, 6)	0.41	1.03	0.22	0.57	0.19	0.03	0.00	0.01	0.93	0.05	0.00	0.00	1.30	1.16	2.84	2.65
(3, 7)	0.25	0.80	0.20	0.61	0.16	0.03	0.01	0.00	1.00	0.00	0.00	0.00	1.27	1.13	2.88	2.78
(3, 8)	0.15	0.51	0.21	0.67	0.09	0.02	0.01	0.00	1.00	0.00	0.00	0.00	1.31	1.06	2.94	2.91
(3, 9)	0.08	0.27	0.21	0.67	0.09	0.02	0.01	0.00	1.00	0.00	0.00	0.00	1.28	1.05	2.97	2.95
(3, 10)	0.05	0.11	0.22	0.65	0.11	0.01	0.01	0.00	1.00	0.00	0.00	0.00	1.28	1.02	3.00	2.98
(3, 11+)	0.05	0.06	0.27	0.63	0.09	0.01	0.00	0.00	1.00	0.00	0.00	0.00	1.33	1.02	3.10	2.98
(4, 0)	3.83	2.27	0.37	0.06	0.34	0.22	0.01	0.09	0.00	0.13	0.78	0.00	1.36	0.73	1.40	0.49
(4, 1)	1.42	1.13	0.35	0.14	0.40	0.10	0.01	0.17	0.00	0.61	0.23	0.00	1.38	1.06	1.97	1.45
(4, 2)	1.33	1.26	0.35	0.19	0.38	0.07	0.01	0.17	0.00	0.78	0.05	0.00	1.36	1.08	2.40	1.94
(4, 3)	0.97	1.22	0.34	0.29	0.33	0.03	0.01	0.16	0.07	0.77	0.00	0.00	1.36	1.00	2.72	2.20
(4, 4)	0.58	1.11	0.31	0.44	0.22	0.02	0.01	0.13	0.32	0.56	0.00	0.00	1.40	1.13	2.80	2.25
(4, 5)	0.38	1.11	0.29	0.51	0.17	0.02	0.01	0.10	0.61	0.30	0.00	0.00	1.38	1.13	2.87	2.47
(4, 6)	0.23	0.95	0.26	0.57	0.13	0.02	0.02	0.08	0.85	0.08	0.00	0.00	1.35	1.13	2.88	2.81
(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	2.95	2.94
(4, 8)	0.07	0.38	0.27	0.62	0.08	0.02	0.01	0.02	0.98	0.00	0.00	0.00	1.35	1.04	2.97	2.98
(4, 9)	0.04	0.17	0.25	0.63	0.09	0.03	0.01	0.00	1.00	0.00	0.00	0.00	1.33	1.02	2.99	3.00
(4, 10)	0.02	0.06	0.32	0.56	0.09	0.02	0.01	0.00	1.00	0.00	0.00	0.00	1.38	0.98	3.08	2.98
(4, 11+)	0.02	0.02	0.40	0.52	0.07	0.01	0.00	0.05	0.95	0.00	0.00	0.00	1.46	1.06	3.24	2.96
(5, 0)	3.02	2.47	0.41	0.09	0.42	0.08	0.01	0.16	0.00	0.51	0.34	0.00	1.49	1.02	1.63	0.94
(5, 1)	0.98	1.08	0.41	0.19	0.35	0.05	0.01	0.24	0.00	0.71	0.05	0.00	1.48	1.13	2.10	1.65
(5, 2)	0.91	1.20	0.40	0.25	0.30	0.04	0.00	0.25	0.00	0.75	0.00	0.00	1.45	1.10	2.51	2.02
(5, 3)	0.65	1.17	0.39	0.34	0.24	0.03	0.00	0.28	0.16	0.56	0.00	0.00	1.46	1.16	2.79	2.35
(5, 4)	0.38	1.11	0.37	0.44	0.16	0.02	0.01	0.29	0.41	0.30	0.00	0.00	1.49	1.33	2.87	2.47
(5, 5)	0.24	1.07	0.35	0.49	0.12	0.03	0.01	0.28	0.65	0.07	0.00	0.00	1.46	1.39	2.93	2.77
(5, 6)	0.15	0.86	0.34	0.52	0.10	0.03	0.01	0.21	0.79	0.00	0.00	0.00	1.44	1.29	2.94	2.93
(5, 7)	0.08	0.59	0.33	0.55	0.07	0.03	0.01	0.12	0.88	0.00	0.00	0.00	1.42	1.16	3.03	2.98
(5, 8)	0.04	0.32	0.33	0.57	0.06	0.01	0.02	0.05	0.95	0.00	0.00	0.00	1.42	1.07	3.06	3.00
(5, 9)	0.02	0.14	0.35	0.58	0.05	0.01	0.01	0.00	1.00	0.00	0.00	0.00	1.44	1.01	3.10	2.98
(5, 10)	0.01	0.05	0.38	0.56	0.05	0.01	0.00	0.01	0.99	0.00	0.00	0.00	1.46	1.01	3.13	2.98
(5, 11+)	0.01	0.02	0.47	0.49	0.03	0.01	0.00	0.10	0.90	0.00	0.00	0.00	1.57	1.17	3.36	2.98
(6, 0)	2.19	2.48	0.44	0.10	0.43	0.03	0.01	0.20	0.00	0.74	0.06	0.00	1.54	1.15	1.72	1.15
(6, 1)	0.64	0.97	0.46	0.20	0.31	0.02	0.01	0.34	0.00	0.66	0.00	0.00	1.54	1.18	2.20	1.74
(6, 2)	0.58	1.06	0.46	0.26	0.26	0.02	0.01	0.40	0.12	0.48	0.00	0.00	1.53	1.30	2.59	2.27
(6, 3)	0.41	1.03	0.44	0.34												

(6, 4)	0.23	0.95	0.42	0.42	0.12	0.02	0.02	0.54	0.42	0.04	0.00	0.00	1.53	1.67	2.88	2.81
(6, 5)	0.15	0.86	0.40	0.47	0.09	0.02	0.02	0.54	0.46	0.00	0.00	0.00	1.51	1.64	2.94	2.93
(6, 6)	0.08	0.65	0.40	0.50	0.08	0.02	0.00	0.44	0.56	0.00	0.00	0.00	1.53	1.49	3.06	2.97
(6, 7)	0.04	0.42	0.40	0.52	0.07	0.01	0.00	0.31	0.69	0.00	0.00	0.00	1.52	1.33	3.07	2.98
(6, 8)	0.02	0.22	0.39	0.51	0.08	0.01	0.00	0.15	0.85	0.00	0.00	0.00	1.48	1.17	3.10	2.99
(6, 9)	0.01	0.10	0.45	0.48	0.05	0.02	0.00	0.06	0.94	0.00	0.00	0.00	1.57	1.10	3.25	3.00
(6, 10)	0.00	0.03	0.45	0.47	0.07	0.02	0.00	0.07	0.93	0.00	0.00	0.00	1.57	1.10	3.17	3.00
(6, 11+)	0.00	0.02	0.59	0.34	0.03	0.03	0.00	0.23	0.77	0.00	0.00	0.00	1.71	1.31	3.46	3.05
(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	1.76	1.19
(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	2.26	1.85
(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	2.63	2.46
(7, 3)	0.25	0.80	0.47	0.36	0.14	0.02	0.00	0.65	0.17	0.18	0.00	0.00	1.61	1.65	2.88	2.78
(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	2.95	2.94
(7, 5)	0.08	0.59	0.43	0.47	0.09	0.01	0.00	0.76	0.24	0.00	0.00	0.00	1.61	1.82	3.03	2.98
(7, 6)	0.04	0.42	0.42	0.49	0.08	0.02	0.01	0.63	0.37	0.00	0.00	0.00	1.55	1.65	3.07	2.98
(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	3.22	2.99
(7, 8)	0.01	0.13	0.49	0.46	0.05	0.01	0.00	0.25	0.75	0.00	0.00	0.00	1.58	1.28	3.16	2.98
(7, 9)	0.00	0.06	0.52	0.40	0.06	0.01	0.00	0.19	0.81	0.00	0.00	0.00	1.57	1.24	3.37	3.02
(7, 10)	0.00	0.02	0.56	0.34	0.09	0.01	0.00	0.24	0.76	0.00	0.00	0.00	1.57	1.35	3.28	3.18
(7, 11+)	0.00	0.01	0.69	0.27	0.03	0.00	0.00	0.34	0.66	0.00	0.00	0.00	1.80	1.44	3.50	3.16
(8, 0)	0.97	2.14	0.47	0.12	0.38	0.03	0.00	0.32	0.02	0.66	0.00	0.00	1.54	1.14	1.78	1.13
(8, 1)	0.24	0.63	0.48	0.32	0.18	0.01	0.00	0.56	0.17	0.27	0.00	0.00	1.63	1.52	2.37	2.34
(8, 2)	0.22	0.62	0.49	0.38	0.12	0.01	0.00	0.67	0.21	0.12	0.00	0.00	1.63	1.71	2.71	2.75
(8, 3)	0.15	0.51	0.48	0.42	0.10	0.01	0.00	0.78	0.20	0.03	0.00	0.00	1.64	1.85	2.94	2.91
(8, 4)	0.07	0.38	0.48	0.43	0.08	0.01	0.00	0.90	0.10	0.00	0.00	0.00	1.62	1.94	2.97	2.98
(8, 5)	0.04	0.32	0.49	0.43	0.07	0.01	0.00	0.91	0.09	0.00	0.00	0.00	1.64	1.93	3.06	3.00
(8, 6)	0.02	0.22	0.49	0.44	0.05	0.02	0.00	0.80	0.20	0.00	0.00	0.00	1.63	1.82	3.10	2.99
(8, 7)	0.01	0.13	0.49	0.44	0.04	0.02	0.00	0.66	0.34	0.00	0.00	0.00	1.58	1.69	3.16	2.98
(8, 8)	0.00	0.07	0.53	0.42	0.04	0.01	0.00	0.45	0.55	0.00	0.00	0.00	1.69	1.52	3.38	3.04
(8, 9)	0.00	0.03	0.56	0.43	0.01	0.00	0.00	0.50	0.50	0.00	0.00	0.00	1.66	1.60	3.45	3.18
(8, 10)	0.00	0.01	0.66	0.28	0.07	0.00	0.00	0.42	0.58	0.00	0.00	0.00	1.72	1.56	3.52	3.19
(8, 11+)	0.00	0.00	0.83	0.17	0.00	0.00	0.00	0.44	0.56	0.00	0.00	0.00	1.89	1.49	3.51	3.08
(9, 0)	0.64	1.80	0.50	0.16	0.32	0.02	0.00	0.47	0.03	0.51	0.00	0.00	1.56	1.29	1.86	1.47
(9, 1)	0.14	0.40	0.51	0.34	0.12	0.03	0.01	0.69	0.21	0.11	0.00	0.00	1.65	1.74	2.40	2.63
(9, 2)	0.13	0.37	0.51	0.39	0.08	0.02	0.00	0.76	0.24	0.01	0.00	0.00	1.65	1.86	2.73	2.87
(9, 3)	0.08	0.27	0.52	0.38	0.08	0.01	0.00	0.84	0.16	0.00	0.00	0.00	1.68	1.90	2.97	2.95
(9, 4)	0.04	0.17	0.54	0.37	0.08	0.01	0.01	0.97	0.03	0.00	0.00	0.00	1.65	1.98	2.99	3.00
(9, 5)	0.02	0.14	0.55	0.37	0.07	0.01	0.00	0.97	0.03	0.00	0.00	0.00	1.66	1.97	3.10	2.98
(9, 6)	0.01	0.10	0.54	0.39	0.06	0.01	0.00	0.88	0.12	0.00	0.00	0.00	1.68	1.90	3.25	3.00
(9, 7)	0.00	0.06	0.61	0.35	0.04	0.00	0.00	0.73	0.27	0.00	0.00	0.00	1.80	1.78	3.37	3.02
(9, 8)	0.00	0.03	0.64	0.33	0.02	0.00	0.00	0.50	0.50	0.00	0.00	0.00	1.79	1.58	3.45	3.18
(9, 9)	0.00	0.01	0.67	0.31	0.02	0.00	0.00	0.49	0.51	0.00	0.00	0.00	1.77	1.57	3.54	3.14
(9, 10)	0.00	0.00	0.58	0.42	0.00	0.00	0.00	0.42	0.58	0.00	0.00	0.00	1.48	1.56	3.24	3.13
(9, 11+)	0.00	0.00	0.83	0.17	0.00	0.00	0.00	0.38	0.62	0.00	0.00	0.00	2.04	1.52	3.73	3.24
(10, 0)	0.41	1.29	0.51	0.23	0.24	0.01	0.01	0.63	0.32	0.05	0.00	0.00	1.61	1.75	1.99	2.44
(10, 1)	0.08	0.19	0.54	0.35	0.09	0.02	0.00	0.79	0.21	0.00	0.00	0.00	1.70	1.88	2.47	2.84
(10, 2)	0.07	0.17	0.54	0.37	0.08	0.01	0.00	0.85	0.15	0.00	0.00	0.00	1.72	1.91	2.78	2.92
(10, 3)	0.05	0.11	0.56	0.37	0.06	0.01	0.00	0.90	0.10	0.00	0.00	0.00	1.72	1.96	3.00	2.98
(10, 4)	0.02	0.06	0.58	0.33	0.08	0.01	0.00	0.99	0.01	0.00	0.00	0.00	1.70	2.00	3.08	2.98
(10, 5)	0.01	0.05	0.62	0.30	0.07	0.01	0.00	0.96	0.04	0.00	0.00	0.00	1.66	1.97	3.13	2.98
(10, 6)	0.00	0.03	0.58	0.32	0.09	0.02	0.00	0.86	0.14	0.00	0.00	0.00	1.60	1.89	3.17	3.00
(10, 7)	0.00	0.02	0.66	0.31	0.02	0.01	0.00	0.75	0.25	0.00	0.00	0.00	1.71	1.83	3.28	3.18
(10, 8)	0.00	0.01	0.74	0.26	0.00	0.00	0.00	0.51	0.49	0.00	0.00	0.00	1.81	1.63	3.52	3.19
(10, 9)	0.00	0.00	0.79	0.21	0.00	0.00	0.00	0.47	0.53	0.00	0.00	0.00	1.75	1.56	3.24	3.13
(10, 10)	0.00	0.00	0.83	0.17	0.00	0.00	0.00	0.30	0.70	0.00	0.00	0.00	1.57	1.49	3.13	2.98
(10, 11+)	0.00	0.00	NaN	NaN	NaN	NaN	NaN	0.36	0.64	0.00	0.00	0.00	NaN	1.48	NaN	3.12
(11+, 0)	0.53	1.21	0.57	0.30	0.12	0.01	0.00	0.75	0.25	0.00	0.00	0.00	1.72	1.86	2.22	2.67
(11+, 1)	0.09	0.10	0.62	0.32	0.06	0.01	0.00	0.91	0.09	0.00	0.00	0.00	1.78	1.94	2.56	2.92
(11+, 2)	0.07	0.09	0.63	0.32	0.05	0.00	0.00	0.94	0.06	0.00	0.00	0.00	1.79	1.97	2.87	2.97
(11+, 3)	0.05	0.06	0.64	0.31	0.05	0.00	0.00	0.94	0.06	0.00	0.00	0.00	1.77	1.96	3.10	2.98
(11+, 4)	0.02	0.02	0.67	0.29	0.04	0.00	0.00	0.89	0.11	0.00	0.00	0.00	1.78	1.91	3.24	2.96
(11+, 5)	0.01	0.02	0.65	0.31	0.04	0.00	0.00	0.80	0.20	0.00	0.00	0.00	1.79	1.81	3.36	2.98
(11+, 6)	0.00	0.02	0.66	0.29	0.05	0.00	0.00	0.65	0.35	0.00	0.00	0.00	1.76	1.74	3.46	3.05
(11+, 7)	0.00	0.01	0.69	0.27	0.03	0.00	0.00	0.57	0.43	0.00	0.00	0.00	1.70	1.73	3.50	3.16
(11+, 8)	0.00	0.00	0.75	0.25	0.00	0.00	0.00	0.50	0.50	0.00	0.00	0.00	1.62	1.59	3.51	3.08
(11+, 9)	0.00	0.00	0.83	0.17	0.00	0.00	0.00	0.47	0.53	0.00	0.00	0.00	1.69	1.71	3.73	3.24
(11+, 10)	0.00	0.00	NaN	NaN	NaN	NaN	NaN	0.58	0.42	0.00	0.00	0.00	NaN	1.63	NaN	3.12
(11+, 11+)	0.00	0.00	NaN	NaN	NaN	NaN	NaN	0.62	0.38	0.00	0.00	0.00	NaN	1.73	NaN	3.46

References

- Andrews, Isaiah, and Ariel Pakes (2016), Linear Moment Inequalities, mimeo, Harvard University.
- 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 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.
- 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 (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.

- Graham, Daniel A., and Robert C. Marshall (1987), Collusive Bidder Behavior at Single-Object Second-Price 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, working paper, Johns Hopkins University.
- 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.
- 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 (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.
- 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.