Information Feedback in First Price Auctions[*]

Ignacio Esponda[†]
New York University

December, 2007

Abstract

I apply the notion of a self-confirming equilibrium (SCE) to study how information feedback in first price auctions (e.g. whether all bids or only the winning bid are revealed at the end of each auction) influences bidders' perceptions about their strategic environment, and consequently their bidding behavior. In a private values setting, revealing the two highest bids is sufficient for bidders to have correct beliefs (which justifies the standard assumption of Nash equilibrium). In contrast, in every symmetric SCE of a symmetric, affiliated, private values model, bidding strategies and revenue are (weakly) higher when only the highest bid is revealed compared to the case where at least the two highest bids are revealed. I also obtain results when valuations are interdependent and discuss the implications for the empirical auction literature.

---

[*] I am especially grateful to Susan Athey for her advice and support since the early stages of this project, and to Liran Einav, Jon Levin, and Ilya Segal for helpful discussions. I also thank the Editor, Philip A. Haile, and two anonymous referees for extensive feedback and very helpful suggestions, and David McAdams, Paul Milgrom, Pedro Miranda, Muriel Niederle, Ariel Pakes, Demian Pouzo, Andy Skrzypacz, Steve Tadelis, and Larry White for their comments. I gratefully acknowledge financial support from the Bradley Graduate and Post Graduate Fellowship Program, through a grant to the Stanford Institute for Economic Policy Research.

[†] Department of Economics, NYU Stern School of Business, 44 W 4th Street, KMC 7-76, New York, NY 10012; iesponda@stern.nyu.edu, http://www.stern.nyu.edu/~iesponda.

> “The analysis of past bids can provide a wealth of valuable information to the contractor in clarifying the relationships among many different competitive factors.” (William R. Park, *Construction Bidding for Profit*, p. 210)

1 Introduction

In practice, the information that bidders receive about auction outcomes is often a design choice of the auctioneer. In the private sector, it is common for the auctioneer to reveal only the winning bid (and many times not even this) at the end of each auction.[1] In contrast, in the public sector the law usually mandates all bids to be revealed. In both cases, it may be of interest to know whether a different information policy would result in different outcomes.

I study how different information feedback policies affect competitive bidding in first price auctions. In particular, I focus on a specific role of feedback: the amount of information regarding outcomes of past, similar auctions may presumably affect bidders’ beliefs about their strategic environment, including their beliefs about how competitors are bidding in these auctions as well as beliefs about their own valuation of the object or procurement cost (in a common value setting). This role of feedback is ignored in the literature when attention is restricted to the (Bayesian) Nash equilibrium of the first price auction, which requires bidders to have correct beliefs about the equilibrium strategies of other players and about the distribution of valuations, so that, e.g., bidders correctly anticipate their probability of winning with any possible bid. In contrast, in this paper players’ beliefs about their probability of winning the auction are influenced by the amount of information that they observe from previous auctions.

The proposed setup is inspired by the empirical literature on auctions, where the stage game is a standard, single-object, first price sealed-bid auction and data are observed from play of several identical and independent auctions.[2] In this context, the empirical literature asks whether the primitives of the stage game (such as the distribution of valuations) can be identified from the data available to the researcher under the assumption that bidders play a Nash equilibrium, and then proceeds to estimate the model (see Athey and Haile (2006, forthcoming) for a review of this literature). But just as identification depends on the data observed by a researcher, whatever players can learn about their strategic environment is also likely to depend on the data that are available to them. As suggested by the quote at the beginning of the paper, bidders do seem to behave as empirical researchers and use information about past bids to learn about their strategic environment.[3]

---

[1] A documented example where only the winning bid is revealed in a sealed-bid auction is the mussels auction in the Netherlands (Schaik and Kleijnen, 2001). More generally, the results of this paper are also relevant for Dutch auctions, since these are equivalent to first price auctions where only the winning bid is revealed.

[2] In contrast, Thomas (2006) studies a setting where valuations are *correlated* across a sequence of auctions and players are assumed to have correct beliefs about the primitives and each others’ equilibrium strategies. The prices revealed in past auctions then affect players’ inferences about each others’ valuations in subsequent auctions, and he finds that procurement costs are non-monotonic in the amount of information revealed.

[3] Friedman (1956) initiated an early operations research literature that studies how bidders should use data from previous auctions to devise optimal bidding strategies (see Laffont (1997) for a discussion). In this paper I not only

I model the effect of different feedback policies on bidders' beliefs about opponents' strategies and the primitives of the economy by assuming that bidders play a self-confirming equilibrium (SCE) (Battigalli (1987), Fudenberg and Levine (1993a), Dekel, Fudenberg, and Levine (2004)). A SCE is often interpreted as the outcome of a learning process, in which players revise their beliefs using observations of previous play.[4] A SCE is similar to a NE in that it is a steady state solution concept applied to a well-defined stage game where players choose strategies that are optimal given their beliefs about the consequences of choosing any feasible strategy. In particular, repeated game considerations (e.g. how feedback affects collusion) are ignored in this paper.[5] The difference between a NE and a SCE is that in a SCE beliefs are not required to be correct, but instead restricted to be "consistent" with feedback that players obtain about equilibrium outcomes. To illustrate the meaning of "consistent", suppose that a bidder always observes the winning bid in the context of playing a series of identical first price auctions. The definition of SCE then requires her to have correct beliefs about the equilibrium distribution of the winning bid but not necessarily correct beliefs about the equilibrium distribution, say, of the second highest bid. By allowing for incorrect beliefs, the set of SCE contains the set of NE and coincides with it only when players have correct beliefs in every SCE.

The main objective is to answer two related questions. First, how much information about opponents' bids needs to be revealed so that bidders have correct beliefs in a SCE (therefore implying that each SCE is also a NE)? This question is important to the extent that several theoretical and empirical results in the literature are obtained under the assumption that bidders play a NE. Second, what is the effect of different information policies on (self-confirming) equilibrium bidding, with a particular emphasis on the revenue-maximizing information policy? The answers to these questions depend on whether types are independent or correlated (affiliated), valuations are private or interdependent, and bidders are symmetric or asymmetric.[6]

In the private values model, beliefs are always correct in a SCE when the auctioneer reveals at least the two highest bids (implying that the sets of SCE and NE coincide in this case). The reason is that the only uncertainty relevant to bidders concerns their probability of winning as a function of their offers. This probability is determined by the distribution of the maximum of opponents' bids, an statistic that is always observed when the highest two bids are revealed.

In contrast, when only the winning bid is revealed a bidder will not observe the maximum of opponents' bids when she wins the auction. Hence, in equilibrium she is only required to have correct beliefs about the probability of winning with offers higher than her equilibrium bid (since she observes the maximum of opponents' bids only in that range of bids). As a result, in a SCE

---

view bidders as learning from past data, but I also follow the equilibrium implications of everyone behaving in such a manner.

[4] Explicit learning-theoretic foundations have been provided by Fudenberg and Levine (1993b).

[5] The effect of feedback on collusion has been emphasized at least since Stigler (1964). For recent studies of this complementary role of feedback see Athey and Bagwell (2001), Athey, Bagwell, and Sanchirico (2004), and Skrzypacz and Hopenhayn (2004).

[6] Valuations are private when they do not depend on opponents' private signals. Otherwise, they are said to be interdependent. These signals can be either independent or affiliated (a strong form of positive correlation). Symmetry refers both to the valuations and to the underlying distribution of signals.

a bidder may overbid relative to the best response of a bidder who has correct beliefs about the probability of winning for each possible bid.[7]

The previous overbidding result establishes a testable relationship between the (perceived) best response of a player in a SCE and the best response of a player who has correct beliefs, as in a NE. However, it does not necessarily follow that self-confirming equilibrium bidding is more aggressive. The reason is that it is not true, in general, that our previous comparison of best responses suffices to obtain an unambiguous comparison of the set of equilibria. I do show, however, that in the symmetric affiliated private values model of Milgrom and Weber (1982) any symmetric SCE is (weakly) more aggressive than the unique NE when only the winning bid is revealed. Furthermore, I establish existence of non-Nash SCE's that are strictly more aggressive and provide strictly higher revenue than the unique NE.

To focus on the new issues that arise when valuations are interdependent, consider the case where the auctioneer reveals at least the two highest bids. In this case bidders may only have incorrect beliefs about the expected surplus from winning the object. Under the often realistic assumption that bidders only receive feedback about the value of the object when they win it, bidders learn about the times they were too optimistic about value, but need not learn they are sometimes too pessimistic. Using a line of reasoning analogous to the one applied for the private values case, it follows that a bidder may now underbid in a SCE relative to the optimal strategy of a bidder with correct beliefs. I then apply the findings on comparative statics in symmetric first price auctions to conclude that when at least the two highest bids are revealed in a symmetric, affiliated model with interdependent valuations, revenues are lower in a symmetric SCE as compared to a NE. When the two highest bids are not revealed, the effect identified in the case of private values affects outcomes in the opposite direction, and equilibria ranking is ambiguous even in a symmetric model.

The results for the symmetric model suggest that in a private value setting an auctioneer might want to limit the amount of information revealed about submitted bids to either increase revenues or lower procurement costs.[8] When valuations are interdependent, an auctioneer might benefit from providing information about the ex post value of an object. The reason is that when bidders observe the realized value of the object (irrespective of whether they win or lose), in a SCE they must have correct beliefs about the expected value of the object. Thus, the sets of SCE and NE coincide (assuming at least the two highest bids are revealed, as explained above), and there is consequently no underbidding in equilibrium. This information policy could be implemented by requiring the winner to provide information about the revenue obtained from the use of the object. For example, in US offshore oil and gas lease sales the federal government publishes monthly data on the production of oil and gas (Porter, 1995).

Another implication of assuming that bidders play a SCE is that the Dutch (i.e. descending)

---

[7] I also show that when (bidders know that) private valuations are independently distributed, then the set of SCE also coincides with the set of NE since bidders of one type may use information that bidders of another type have learned.

[8] In this paper I focus on a particular effect of information feedback. In general, when devising a revenue-maximizing mechanism, an auctioneer should consider other possible effects of information on bidding behavior, such as the effects on collusion, corruption, entry, and on bidders' privacy concerns.

and first price auctions are not equivalent, unless only the highest bid is revealed in the latter. Hence, information feedback constitutes an alternative, non-psychological explanation to the non-equivalence of the Dutch and first price auctions.[9]

Several papers in the experimental literature emphasize that behavior depends on the amount of feedback that players receive as they learn to play the game.[10] In the context of first price auctions, Isaac and Walter (1985), Neugebauer and Selten (2006), and Ockenfels and Selten (2006) show that in experimental settings with independent private values overbidding occurs when subjects only receive feedback about the winning bid, compared to the case where they receive more feedback about their opponents' bids. The results in this paper shed light not only on the forces that may be behind their findings but also suggest new directions for further experimental work.[11]

In a recent paper, Jehiel (2007) also looks at feedback in auctions as a design issue. In contrast, he allows for more general auction formats and focuses on a different kind of feedback. In his framework, bidders obtain coarse feedback about past bids, so that, for example, only the aggregate distribution of bids across the various auction formats that the auctioneer has implemented in the past are revealed.[12]

In Section 2, I describe the model and define both NE and SCE. I then present conditions for beliefs to be correct in a SCE in Section 3, and characterize non-Nash SCE behavior in Section 4. In Section 5, I discuss implications for the empirical auction literature and show that SCE behavior imposes testable restrictions on observed bids. I briefly conclude in Section 6, and relegate all proofs to the Appendix.

2 Setup and definition of equilibrium

Consider the following stage game, which describes a standard first price auction. There are N risk-neutral bidders who simultaneously submit bids for a single object.[13] The bidder who submits the highest bid gets the object and pays her bid; more generally, the winner is chosen randomly from the set of bidders who submit the highest bid. Before submitting their bid, each bidder i receives a signal or type $s_i \in [\underline{s}_i, \overline{s}_i] \subset \mathbb{R}_+$. Player i's utility from winning the object is $v_i - b_i$,

---

[9] A few experimental papers have shown that the winning bid tends to be higher in IPV first price auctions compared to Dutch auctions (e.g., Cox, Robertson, and Smith 1982). However, these experiments report only the winning bid in first price auctions, so that the observed difference must be attributed to reasons other than information feedback. It would be interesting to see whether this result would be reversed when more information is revealed in first price auctions, as predicted by the analysis in this paper.

[10] Kagel (1995, p. 521) argues that "the price information may provide [bidders] with a better sense of what their rivals' bids are likely to be and to help correct for overly pessimistic, or overly optimistic, initial expectations." In a sense, this paper formalizes and extends this intuition.

[11] Dufwenberg and Gneezy (2002) and Neugebauer and Perote (forthcoming) also study feedback in auction experiments. Alternative explanations for overbidding (relative to NE) in private value auctions include risk aversion (e.g. Cox et al. 1982), a non-equilibrium model of level-k thinking (Crawford and Iriberri, 2007), and a model where bidders feel regret depending on the feedback they receive about others' bids (Engelbrecht-Wiggans and Katok (2007) and Filiz and Ozbay (2007)).

[12] He also considers cases where a bidder only gets feedback about the aggregate distribution of bids that were submitted in other, similar auctions where, in contrast to this paper, the bidder has not participated.

[13] The results can be easily adapted to the case where an object is being procured rather than sold.

where $v_i \in \mathbb{R}_+$ is her finite valuation and $b_i \in \mathbb{R}_+$ is her offer. Losing bidders pay nothing, and their utility is normalized to zero. Signals and valuations are drawn according to a joint probability distribution $F_{S,V}$ over the random variables $S = (S_1, ..., S_N)$ and $V = (V_1, ..., V_N)$. Additional regularity assumptions will be made on $F_{S,V}$, but only in the context of the symmetric model in Section 4. A (pure) strategy β_i is a function mapping types to bids, and a profile of strategies is denoted by $\beta = (\beta_1, ..., \beta_N)$. Player i's bid under strategy β_i is a random variable $B_i = \beta_i(S_i)$, and B^1 and B^2 are the highest and second highest bids among $B_1, ..., B_N$.

□ **Beliefs.** Let $\mathbb{G}_i$ denote the set of all joint distributions over (B_{-i}, V_i). Each player i of type s_i has a belief $G_{s_i} \in \mathbb{G}_i$ over the bids of other players and her own valuation of the object. The expected profit from choosing bid b_i that is perceived by player i of type s_i who has belief G_{s_i} is

$$\pi_i(b_i \mid G_{s_i}) = E_{B_{-i},V_i}(V_i - b_i)1\{b_i \geq \max_{j \neq i} B_j\},$$

where the expectation is taken according to the distribution G_{s_i}.

Let $G^0_{s_i,\beta_{-i}}$ denote the joint distribution over $(\beta_{-i}(S_{-i}), V_i)$ conditional on $S_i = s_i$ – in words, $G^0_{s_i,\beta_{-i}}$ represents the true distribution over (B_{-i}, V_i) being faced by type s_i when her opponents choose strategies β_{-i}. In this paper, type s_i's belief, G_{s_i}, need not necessarily coincide with the true distribution $G^0_{s_i,\beta_{-i}}$.

A special case of the model restricts bidders to have correct beliefs about their own valuation of the object (as a function of their type), so that bidders only face uncertainty about their probability of winning.[14]

Definition 1 (private values) *In the private values model, the expected profit from choosing bid* b_i *that is perceived by player* i *of type* s_i *who has belief* G_{s_i} *is* $\pi_i(b_i \mid G_{s_i}) = (s_i - b_i)\Pr(b_i \geq \max_{j \neq i} B_j)$*, where the probability is taken using the marginal distribution over* B_{-i} *implied by* G_{s_i}*.*

In the more general case where the private values assumption does not hold, valuations are said to be interdependent.

□ **Consistency of beliefs**. In a SCE, beliefs are required to be consistent with the information feedback that players would obtain from repeatedly playing their equilibrium strategies in a sequence of independent auctions. Let $\mathcal{I}_i$ denote the set of all functions of the random variables $(B_1, ...B_N, V_i)$, and let $I_i \subset \mathcal{I}_i$ denote the feedback observed by bidder i. For example, $I_i = \{B_i, B^1, V_i\}$ implies that bidder i observes her own bid, the highest of all bids, and her own valuation at the end of each auction.[15] Throughout, I make the natural assumption that $B_i \in I_i$, so that players always observe their own bids.

---

[14] The standard private values assumption requires that $V_i \mid s_i \overset{D}{=} V_i \mid s_i, s_{-i}$. Here I also make explicit that players know how their signals map into their expected valuations, so that requiring $s_i = E(V_i \mid s_i)$ is just a normalization.

[15] See Esponda (2007) for a more general setting where players may also get feedback about the realized signals and valuations of other players and use such information, together with additional information about, for example, the rationality of other players, to make inferences about what other players are doing in equilibrium.

From the point of view of type s_i, the distribution of observable outcomes depends not only on her own type s_i and on others' strategies (which determine the joint distribution over (B_{-i}, V_i)), but also on the bid b_i (which determines $B_i = b_i$) that is chosen in equilibrium by type s_i. The distribution over $(B_1, ..., B_N, V_i)$ then determines the joint distribution over the set of observable random variables in I_i. Consistency requires a player to have correct beliefs about this distribution of observable outcomes.

Definition 2 (consistency) *A belief $G_{s_i} \in \mathbb{G}_i$ of player i of type s_i is I_i-consistent for (b_i, β_{-i}) if the joint distribution of the random variables in I_i, conditional on $B_i = b_i$, is the same whether the joint distribution over (B_{-i}, V_i) is given by G_{s_i} or by the correct distribution $G^0_{s_i,\beta_{-i}}$.*

☐ **Definition of Equilibrium**. In a SCE, players choose strategies that are optimal given their beliefs, and these beliefs are required to be consistent with observables obtained through feedback.[16]

Definition 3 (self-confirming equilibrium) *The strategy profile β is a self-confirming equilibrium (SCE) with information feedback $\{I_i\}_{i \in N}$ if for every player i and for almost every type s_i there exists a belief G_{s_i} such that:*
(i) $\beta_i(s_i) \in \arg\max_{b_i} \pi_i(b_i \mid G_{s_i})$, and
(ii) G_{s_i} is I_i-consistent for $(\beta_i(s_i), \beta_{-i})$.

When beliefs are correct, the definition of a SCE coincides with the definition of a Bayesian Nash equilibrium (NE). Let SCE and NE denote the set of self-confirming and Nash equilibria, respectively. It follows from the definition that $NE \subset SCE$, since consistency does not rule out the case where players may somehow have correct beliefs. The objective of this paper is to further characterize and compare these equilibrium sets for different information policies that may naturally arise in first price auctions.

In applications, it may also be desirable to place additional restrictions on beliefs. One example is the private values restriction defined above. Other examples will be considered in the paper.

☐ **Interpretation.** The idea underlying the concept of SCE is that the stage game is being repeatedly played and that players use the information revealed after each stage game to update their beliefs. This interpretation is analogous to thinking of bidders as econometricians in a context that is standard in the empirical auction literature. Here, bidders participate in several identical but independent auctions, where for each auction bidders draw a different type and valuation from the primitive $F_{S,V}$, and draws are independent across auctions. Bidders then use observable outcomes

---

[16] The definition of SCE is less restrictive than that in Dekel, Fudenberg, and Levine (2004) since it does not require bidders to derive their beliefs about the joint distribution of their valuation and opponents' bids from beliefs about the distribution of signals and valuations and the signal-contingent strategies of the other players. The results in this paper hold without having to make the above additional restriction, and it is possible to show that non-Nash SCEs would still exist under this additional restriction.

from previous auctions to form beliefs about their strategic environment. A SCE simply captures the steady state of such a dynamic process, but the dynamic process itself is left unspecified (see Fudenberg and Levine (1993a) and Dekel, Fudenberg, and Levine (2004) for more on the interpretation of SCE).

3 Conditions for correct beliefs in a SCE

In this section I study conditions on the feedback policy under which beliefs are correct in a SCE, and therefore the set of SCE coincides with the set of NE. This question is related to the literature on identification in auctions (Guerre, Perrigne and Vuong (2000), Athey and Haile (2002)). While that literature studies the observables that are needed for an *econometrician* to identify the primitives of an auction model, this section is concerned with the observables that are needed for a *bidder* to have correct beliefs in equilibrium and consequently justify the Nash equilibrium assumption. These questions are related because in order for an econometrician to identify the primitives of the model, he must first identify the strategic environment (e.g. probability of winning) being faced by each bidder. This relation explains the close parallels between the bidder identification results in this section and the econometrician identification results in the literature. The proofs of these results differ slightly to the extent that the bidders are part of the game (therefore observing their own types and bids, and determining endogenously through their actions what they are going to observe in equilibrium).

□ **Private values model.** In the private values model, observing at least the two highest bids results in correct beliefs in a SCE. The reason is that by observing the two highest bids, a bidder always observes the maximum of opponents' bids, and therefore (by the consistency requirement) must have correct beliefs about the distribution of the maximum of opponents' bids. Such distribution determines the probability of winning for every possible bid, which is the only source of uncertainty in the private values model.[17]

Proposition 1 *$SCE = NE$ in the private values first price auction where the two highest bids are observed.*

In contrast, when only the winning bid is observed, a bidder no longer observes the maximum of opponents' bids when she is the winner of the auction. Section 4 shows that non-Nash SCE exist in this case, so that the corresponding bidders' SCE strategies are sustained by incorrect beliefs. However, by making the additional restriction that valuations are independently distributed and that bidders are aware of this fact (which must be assumed, since it does not follow from consistency

---

[17] The parallel result in the identification literature is given by Athey and Haile (2002, Lemma 1). The difference is that in Proposition 1 bidders' identities need not be observed.

of beliefs), we recover the result that beliefs are correct in a SCE even if only the winning bid is observed.[18]

Definition 4 (independent valuations) *The independent private values model is a private values model where in addition (i) $S_1, ..., S_N$ are mutually independent and (ii) for each player i: for every s_i, s_i', G_{s_i} and $G_{s_i'}$ both imply the same distribution over B_{-i}.*

Proposition 2 *$SCE = NE$ (restricted to the set of bids that win with strictly positive probability[19]) in the independent private values first price auction where only the winning bid is observed.*

The proof relies on two arguments. First, since the winning bid is observed, a bidder of a particular type observes the maximum of the opponents' bids whenever she loses the auction. Consistency then requires her to have correct beliefs about the distribution of the maximum of opponents' bids in the range of bids that are higher than the equilibrium bid for her particular type. Second, under the independence assumption a bidder knows that the distribution of other bidders' bids is independent of her own type. Hence, she can use what she knows about other players' bids when she is of a type that chooses low bids, since this information continues to be valid when she is of a type that chooses higher bids.[20]

☐ **Interdependent values model.** Without the private values assumption, equilibrium outcomes depend not only on beliefs about the probability of winning but also on beliefs about the expected surplus from winning the object with a certain bid. Beliefs are correct in a SCE when the two highest bids are observed and, in addition, players observe the realization of their own valuations after each auction. Essentially, observing one's own valuation is equivalent to having correct beliefs about one's own valuation, as in the private values model.[21]

Proposition 3 *$SCE = NE$ in a first price auction where players observe the two highest bids and their own valuations.*

In many situations, it may be more reasonable to assume that players observe the realization of their own valuation when they win the object, but not otherwise. In this case, non-Nash SCE exist and are characterized in the next section.

---

[18] The restriction that players must *know* that signals are independent and take this knowledge into account when learning about their opponents' bids might explain why in the private-value experimental settings of Isaac and Walter (1985) and Neugebauer and Selten (2006) described in the introduction bidders do not actually play a Nash equilibrium, but rather behave in a manner consistent with our results for a non-independent setting.

[19] This weaker condition is sufficient for it to be valid to focus on Nash equilibrium. The reason is that bids that win with probability zero are arbitrary to the extent that they can be replaced with bids that also win with probability zero without affecting neither the utility of the agents nor the revenue of the auctioneer.

[20] The idea in the second argument appears in Example 1 of Dekel, Fudenberg and Levine (2004). The parallel result in the identification literature is given by Athey and Haile (2002, Theorem 6(i)). The difference is that in Proposition 2, bidders' identities need not be observed and bids need not have a common support.

[21] Relatedly, in the pure common value model, Hendricks, Pinkse, and Porter (2003) obtain identification when the econometrician observes ex-post the common value of the object. The setting in Proposition 3 is more general but also requires each bidder to observe ex-post the value of her own valuation.

4 Characterization of non-Nash SCE behavior

In this section I characterize SCE in those cases where players need not have correct beliefs in equilibrium. From Section 3, these are the cases where only the winning bid is observed (in the private values model), or where the two highest bids are observed but bidders only observe their own valuation in the auctions that they win (in the interdependent values model). While a useful characterization of SCE can be obtained without making further assumptions, an unambiguous comparison of equilibria (and therefore a comparison of different feedback policies) requires that we specialize the setting to the symmetric model of Milgrom and Weber (1982).

□ **Characterization of SCE**. Define

$$\rho_i^R(\beta_{-i}) \equiv \left\{ \widehat{\beta}_i : \begin{array}{c} \text{for a.e. } s_i,\ \widehat{\beta}_i(s_i) \in \arg\max_{b_i \geq \widehat{\beta}_i(s_i)} \pi_i(b_i \mid G^0_{s_i,\beta_{-i}}) \\ \text{and } \pi_i(\widehat{\beta}_i(s_i) \mid G^0_{s_i,\beta_{-i}}) \geq 0 \end{array} \right\}$$

and define $\rho_i^L(\beta_{-i})$ in exactly the same way except that the maximization problem takes places with respect to $b_i \leq \widehat{\beta}_i(s_i)$ rather than $b_i \geq \widehat{\beta}_i(s_i)$.

Lemma 4 (characterization of SCE) *(i) β is a SCE of a private values first price auction where only the winning bid is observed if and only if $\beta_i \in \rho_i^R(\beta_{-i})$ for all i; (ii) β is a SCE of an interdependent values first price auction where only the two highest bids and winners' own valuations are observed if and only if $\beta_i \in \rho_i^L(\beta_{-i})$.*

Consider Lemma 4 for the private values model. The only if part confirms the intuition that there is in some sense overbidding when only the winning bid is revealed. It says that if β is a SCE then, for every bidder i and for every type s_i, it would actually be optimal to bid $\beta_i(s_i)$ if restricted to choose some bid that is higher than $\beta_i(s_i)$. In addition, bidding $\beta_i(s_i)$ must provide at least expected utility of zero, since that is the utility that bidders know that they can at least obtain by bidding zero. Hence, either $\beta_i(s_i)$ is a best response to other bidders' strategies or it involves overbidding relative to (any) best response. Intuitively, overbidding is possible since a winning bidder does not get feedback about how much money was left on the table (i.e. the difference between her bid and the second highest bid).

On the other hand, part (ii) suggests that there is in some sense underbidding in the interdependent values model: either $\beta_i(s_i)$ is a best response to other bidders' strategies or it involves underbidding relative to (any) best response. To see this, note that, as previously determined, observing the two highest bids implies that in equilibrium bidders have correct beliefs about their probability of winning. Furthermore, since a bidder of a particular type observes both her own valuation and the maximum of the opponents' bids (which is below her equilibrium bid) when she wins, by consistency she must have correct beliefs about her expected valuation had she chosen any bid below her equilibrium bid. Hence, in equilibrium bidders have correct beliefs about their expected profits from choosing bids below their equilibrium bids.

However, Lemma 4 only states that over/underbidding may occur with respect to the best response to others' strategies, but this result does not imply that in (a self-confirming) equilibrium bidders will overbid relative to Nash equilibrium behavior. An important lesson from the modern comparative statics literature (e.g. Milgrom and Roberts, 1994) is that obtaining a monotone ordering of best responses is in general not sufficient for equilibria to be compared. When only the winning bid is revealed, it is indeed possible to construct examples with both (self-confirming) equilibrium underbidding and overbidding relative to Nash equilibrium.[22]

An interesting question is whether it is possible to obtain sharp predictions about the effect of different information policies in those particular auction settings where auction theory also obtains unambiguous comparative statics results.[23] I use Lemma 4 together with a property of symmetric, affiliated first price auctions (see below) to show that the previous statements regarding over/underbidding also hold in equilibrium.

In addition, Lemma 4 implies that in a setting where only the winning bid is revealed a SCE can still be characterized as a fixed point of a (generalized) "best response" correspondence. This characterization can be useful for finding a SCE (e.g. by numerical computation in a discrete bid setting), for testing whether bidders play a SCE in the data and for estimating the primitives of a structural model under the assumption that bidders play a SCE (see Section 5), and for showing that a candidate profile constitutes a SCE, as illustrated below.

☐ **Example.** Consider a first price private values auction with $N = 2$ bidders and i.i.d. (for simplicity) valuations $S_i \sim U[0,1]$, but where bidders do not know that valuations are independent. The unique Nash equilibrium is symmetric and is given by $\beta^{NE}(s) = .5s$ for $s \in [0,1]$. For any $k \in [.5, 1)$, $\beta^{SCE}(s) = ks$ is a SCE. To see this, note that the best response to β^{SCE} is $.5s$ and that this best response is lower than $\beta^{SCE}(s)$. Together with the concavity of the payoff function, it follows that β^{SCE} is the best response when restricted to $b \geq \beta^{SCE}$. Hence $\beta^{SCE} \in \rho^R(\beta^{SCE})$, and Lemma 4 implies that β^{SCE} is a SCE.

A multiplicity of beliefs can sustain the previous SCE. For example, G_s can be such that type s believes her probability of winning with a bid b is given by

$$P(b \mid s) = \begin{cases} b/k & \text{if } b \in (ks, k) \\ \frac{(1-k)sb}{k(s-b)} & \text{if } b \in [0, ks] \end{cases}$$

☐ **The monotone-symmetric model.** Now I restrict attention to the symmetric setup of Milgrom and Weber (1982), where the following assumptions are added to the setup in Section 2: (i) $V_i = u_i(S_1, ..., S_N)$, where $u_i(s) = u(s_i, \{s_{-i}\})$ and u is nonnegative, twice continuously differen-

---

[22] Equilibria could be unambiguously compared if the game had strategic complementarities, but it is well known (see Athey 2001) that this is not the case for first price auctions.

[23] The standard proof of many comparative statics results in symmetric auctions characterizes Nash equilibrium as the solution to a differential equation and studies how this solution changes as a parameter of interest is varied. This approach is in principle not feasible here because of the need to rank two different equilibrium concepts rather than the equilibrium of a parameterized model.

tiable, nondecreasing in s_{-i} and increasing in s_i; (ii) $S_1, ..., S_N$ are affiliated random variables with joint density f that is symmetric, continuously differentiable on its support $[\underline{s}, \overline{s}]^N$, and bounded away from zero and infinity; and (iii) bidders are restricted to play nondecreasing strategies. I will refer to this setting as the monotone-symmetric model. Assumptions (i) and (ii) imply that bidders are symmetric in terms of their signals and valuations, so that I drop the i subscript for this symmetric model, and now β represents a strategy rather than a strategy profile. Assumptions (ii) and (iii) imply a monotone ordering of expected valuations in terms of the offered bid.

In the monotone-symmetric model there exists a unique Nash equilibrium and it is in symmetric, continuous, increasing, and differentiable strategies (Milgrom and Weber (1982), McAdams (2007)). Similar results hold for the set of SCE under a natural restriction on beliefs.

In addition, I make the following restriction on beliefs, which is immediately satisfied in the private values model.[24]

Definition 5 (monotone beliefs) *Beliefs are monotone if for every player* i *and type* s_i*, the belief* $G_{s_i} \in \mathbb{G}_i$ *is such that* $E(V_i \mid b_i \geq \max_{j \neq i} B_j)$ *is nondecreasing in* b_i*, where the expectation is taken according to the distribution* G_{s_i}*.*

Lemma 5 *Consider the monotone-symmetric first price auction where beliefs are monotone and the winning bid is revealed. Let* β^{SCE} *be a symmetric SCE. Then* β^{SCE} *is increasing and satisfies* $\lim_{s \to \underline{s}} \beta^{SCE}(s) = \beta^{NE}(\underline{s})$*.*

The following lemma partly characterizes best responses in the monotone-symmetric model and is the crucial result explaining why unambiguous comparative statics results are obtained in such a setting. Strategies are compared according to the partial order $\geq$, where $\beta_A \geq \beta_B$ if and only if $\beta_A(s) \geq \beta_B(s)$ for a.e. $s \in [\underline{s}, \overline{s}]$. Also, define $\beta_A \not\geq \beta_B$ as $\beta^A \notin \{\beta : \beta \geq \beta_B\}$ and let μ denote the standard Lebesgue measure.

Lemma 6 *Let* β^{NE} *be the (symmetric) Nash equilibrium of the monotone-symmetric first price auction. Suppose opponents play a symmetric profile* β *such that* β *is increasing and* $\lim_{s \to \underline{s}} \beta(s) = \beta^{NE}(\underline{s})$*.*

- *(a) If* $\beta \not\geq \beta^{NE}$ *then* $\mu\left(s : \beta(s) \notin \arg\max_{b \geq \beta(s)} \pi(b \mid G^0_{s,\beta})\right) > 0$*.*
- *(b) If* $\beta \not\leq \beta^{NE}$ *then* $\mu\left(s : \beta(s) \notin \arg\max_{b \leq \beta(s)} \pi(b \mid G^0_{s,\beta})\right) > 0$*.*

---

[24] Given the restriction that bidders play nondecreasing strategies and given the primitives of the economy, *monotone beliefs* requires players to believe something that is actually true. In particular, one implication is that a Nash equilibrium is still a SCE.

In words, part (a) roughly says that for a strategy profile that is (i) a potential SCE, and (ii) not higher than the Nash equilibrium strategy, there is a positive measure of types s that, facing β, would obtain higher utility by choosing some bid that is higher than $\beta(s)$ rather than by choosing $\beta(s)$.

To gain some intuition for this result, let β^{NE} be the NE and consider a strategy β that is lower than β^{NE} and a particular type s at which the slope of β is flatter than the slope of β^{NE}. By definition of Nash equilibrium, $\beta^{NE}(s)$ is optimal for type s given that others play β^{NE}, so that at $\beta^{NE}(s)$ the marginal payoff change from increasing the bid is zero. Consider the same type s but now suppose others play β. Suppose type s were to bid $\beta(s)$. Then both the probability of winning and the expected valuation of the object would be the same as when opponents were playing β^{NE} and she was best responding with $\beta^{NE}(s)$. The difference is that now the marginal payoff change from increasing her bid is positive since: (i) $\beta(s) < \beta^{NE}(s)$ implies surplus is higher, (ii) the slope of β is flatter than the slope of β^{NE} at s, and therefore she would now outbid more types by slightly increasing her bid, and (iii) the expected value of the objects she wins does not decrease when she increases her bid. Hence, type s would obtain higher utility by choosing some bid larger than $\beta(s)$ than by choosing $\beta(s)$.

□ **Comparing feedback policies in the monotone-symmetric model.** In the monotone-symmetric model with private values, the policy of revealing only the winning bid can be unambiguously compared to the policy of revealing at least the two highest bids.

Proposition 7 *Consider the monotone-symmetric first price auction with private values where the winning bid is observed. Let* β^{NE} *be the unique (symmetric) Nash equilibrium and suppose that* β^{SCE} *is a symmetric SCE. Then* $\beta^{SCE} \geq \beta^{NE}$.

Essentially, starting from a symmetric strategy profile that is below the symmetric Nash equilibrium strategy, Lemma 6(i) implies that a positive measure of types would gain by increasing their bid. However, as Lemma 4(i) indicates, in a SCE bidders cannot gain by increasing their bid. On the other hand, a strategy below the NE is not necessarily ruled out as a SCE since, while Lemma 6(ii) implies that a positive measure of types would gain by decreasing their bid, in a SCE bidders need not have correct beliefs about what would happen in equilibrium if they lower their bid.

Proposition 7 immediately implies the following result.[25]

Corollary 8 *In the monotone-symmetric first price auction with private values, the policy of revealing only the winning bid results in weakly higher revenues than the policy of revealing at least the two highest bids when bidders play a symmetric (self-confirming) equilibrium.*

---

[25] The qualification regarding symmetric equilibria is important since there do exist asymmetric self-confirming equilibria in the monotone-symmetric model, and it appears to be hard to obtain unambiguous results for asymmetric equilibria as well. However, players' beliefs are not necessarily restricted to be symmetric.

The next result establishes the existence of non-Nash SCE, so that from the previous result bidding is strictly more aggressive for such SCE compared to the NE.[26]

Proposition 9 *Consider the monotone-symmetric first price auction with private values where the winning bid is observed. Let* β^{NE} *be the unique (symmetric) Nash equilibrium and let the function* $\alpha : [\underline{s}, \overline{s}] \rightarrow \mathbb{R}_+$ *be differentiable, nondecreasing, and satisfy* $\alpha(\underline{s}) = 0$ *and* $\alpha(s) < s - \beta^{NE}(s)$ *for all* s*. Then* $\beta^{SCE} = \beta^{NE} + \alpha$ *is a symmetric SCE.*

The idea of the proof proceeds in two steps: first, look for a strategy β such that its best response is lower than itself and such that expected utility is nonnegative for every type; and second, check that for every s, $\pi(b \mid G^0_{s,\beta})$ is quasiconcave in b, so that for every s, $\beta(s)$ is a better response to β than any bid higher than $\beta(s)$.

Analogous arguments can be made for the interdependent values model. The next result is a direct consequence of Lemmas 4(ii) and 6(ii).[27]

Proposition 10 *Consider the monotone-symmetric first price auction where beliefs are monotone and where only the two highest bids and winners' own valuations are observed. Let* β^{NE} *be the unique (symmetric) Nash equilibrium, and suppose that* β^{SCE} *is a symmetric SCE. Then* $\beta^{SCE} \leq \beta^{NE}$*.*

Proposition 10 implies that, when bidders observe the two highest bids and observe their own valuation only when they win, symmetric equilibrium bidding is less aggressive relative to NE (or to the case where bidders also observe their realized valuations in auctions where they lose). An implication is that an auctioneer might increase revenues by providing information about the ex post value of the object. For example, if the object is exploited to obtain future revenues, the auctioneer might want to require the bidder to report her revenues. This is actually the case in US offshore oil and gas lease sales, where the federal government publishes monthly production of oil and gas.

Finally, consider the case where only the winning bid is observed, rather than the two highest bids. Regarding beliefs about valuations, since now only the winning bid is revealed, a bidder may be uncertain about the expected value of the object were she to win with either a higher or a lower bid. Regarding beliefs about strategies, a bidder may be uncertain about her probability of winning with a lower bid. Although the second effect takes a specific direction, the first effect can take any direction, leading to no unambiguous ranking of best responses and, consequently, of equilibria. Therefore, it is not necessarily true that revealing only the winning bid yields more revenue when valuations are interdependent.[28]

---

[26] There are other non-Nash SCE which are not characterized in Proposition 9.

[27] Existence of non-Nash SCE can also be established following the logic in Proposition 9.

[28] An interesting question in the interdependent values setting is what would happen in the presence of bidders who use information to learn about the environment, as in this paper, but who fail to realize that winning provides

□ **Equilibrium refinements.** Since the concept of SCE often leads to a multiplicity of equilibrium outcomes, in the future it may be worth exploring several refinements. One possibility is that players decide to experiment in such a way that they end up with correct beliefs, as in a Nash equilibrium. However, bidders might not always have incentives to fully experiment, since by experimenting they sacrifice current payoffs in the hope of obtaining higher future payoffs (see Fudenberg and Levine (1993b) for a formal model). Also, even full experimentation by one bidder may not lead to correct beliefs if the others are not already using NE strategies, and one may even interpret the set of SCE as the possible steady states once experimentation has taken place.

An alternative refinement consists of placing further a priori restrictions on the beliefs that players are allowed to have in equilibrium. For example, in certain settings requiring bidders to believe that the support of the maximum of opponents' bids does not depend on their type would eliminate SCE in strategies that are not continuous.[29] Moreover, requiring bidders to believe that, for each of their types s_i, G_{s_i} is such that the perceived profit function $\pi_i(\cdot \mid G_{s_i})$ is differentiable would eliminate the SCE provided in the example in this section as well as those characterized in Proposition 9.[30] However, even if one were able to justify that players have differentiable beliefs, this latter refinement may be too strong. The spirit of SCE is that players form their beliefs using past data that is presumably finite. To make the analysis tractable, consistency implicitly assumes that an unlimited amount of data is available. However, with finite data it may not be possible to reject the hypothesis that the (right hand side) slope of the profit function at the equilibrium bid is different than zero, even if this were actually the case.[31]

More generally, further work is needed to understand both the dynamics of experimentation and the actual estimation strategies of real-life bidders. Advances in this respect would provide better guidance about the refinements that would be appropriate in different circumstances. Nevertheless, as long as there is room left for beliefs to be incorrect in the presence of partial data, the findings in this section remain relevant.

5 Implications for the empirical auction literature

Once the issue of bidder identification has been studied, a natural question is what are the implications for the testing and estimation of a structural auction model where the econometrician takes into account that bidding behavior depends on bidders' observation of past outcomes. The

---

information about the value of the object. Esponda (forthcoming) defines the notion of a behavioral equilibrium and applies it to general contexts where players ignore such selection problems when learning. In symmetric first price auctions, the presence of such naive bidders generates underbidding in equilibrium.

[29] In the context of the example in this section, $\beta(s) = \frac{1}{2}s$ for $s \in [0, .5)$ and $\beta(s) = \frac{1}{2}s + \frac{1}{8}$ for $s \in [.5, 1]$ is a symmetric SCE strategy that is not continuous at $s = 0.5$, but where bidders must believe that the support of the opponent's bids changes with their own type.

[30] It is an open question whether a non-Nash SCE exists in the monotone-symmetric model under the differentiability refinement, since a bidder may still choose to bid at a local, rather than a global, maximum.

[31] If bidders use, for example, Kernel estimation, then even unlimited data may result in an imprecise estimate of the right hand side derivative of the profit function at the equilibrium bid. The reason is that the Kernel estimate of this function is not necessarily consistent near this boundary point (e.g. Zhang, Karunamuni, and Jones, 1999).

distinction between bidder and econometrician identification highlighted in this paper is important because in circumstances where the econometrician has managed to collect enough data to Nash-identify the model, it is still possible that the data available to bidders at the time of bidding were not sufficient to justify that *players* play a Nash equilibrium. In a private values setting, for example, the econometrician may have observed all bids (therefore obtaining the probability of winning faced by each player), but inferring players' valuations (and therefore the primitives of the model) assuming that players themselves have correct beliefs about their probability of winning is likely to be inadequate if the auctioneer only reveals the winning bid.[32] Similarly, in a pure common value setting, it is important for identification that the econometrician obtains ex-post estimates of the common value of the object (Hendricks, Pinkse, and Porter (2003)), but in addition he must make sure that bidders themselves learn by accessing this information. The latter may not always be true if, for example, bidders only obtain this information after several years and in a context where different market conditions render extrapolation from the past inadequate. On a more positive side, in cases where the econometrician has enough data and players have correct beliefs in a SCE, the standard identification and estimation approach is robust to the use of the less restrictive and more realistic concept of self-confirming equilibrium.

The characterizations of SCE provided in the paper can also be applied to test whether bidders do play a SCE, and to estimate the parameters of a structural model under the assumption that they do play a SCE. To illustrate how to establish testable restrictions, let $\mathcal{P}$ be the set of probability distributions with bounded support on $\mathbb{R}_+^N$. A distribution $P \in \mathcal{P}$ of observed bids $B_1, ..., B_N$, with support supp P, is *Nash-rationalized* by a private values model if P is the NE distribution of bids when valuations are distributed according to some primitive F_S. Similarly, P is *SCE-rationalized* by a private values model where the winning bid is observed if P is the SCE distribution of bids when valuations are distributed according to some primitive F_S.

For a fixed P, let $P_i(b_i \mid b_i^t) \equiv \Pr(b_i \geq \max_{j \neq i} B_j \mid B_i = b_i^t)$ denote the probability that bidder i would win with bid b_i in those cases where she chooses to bid b_i^t. If the data come from play of a NE, optimality requires that for every observed bid b_i^t with associated valuation s_i^t,

$$(s_i^t - b_i^t) P_i(b_i^t \mid b_i^t) \geq (s_i^t - b_i) P_i(b_i \mid b_i^t) \tag{1}$$

for all b_i. To compare this expression with the one resulting under the assumption that bidders play a SCE and observe only the winning bid, define[33]

$$H_i(b_i^t, b_i) \equiv b_i^t + \frac{(b_i^t - b_i) P_i(b_i \mid b_i^t)}{P_i(b_i^t \mid b_i^t) - P_i(b_i \mid b_i^t)}.$$

---

[32] For another example, Laffont, Ossard, and Vuong (1995) parametrically estimate a Dutch auction by assuming that bidders play a NE. From Proposition 2, this assumption is justified when bidders know that valuations are independent and benefit from learning across types.

[33] When (b_i^t, b_i) is such that $P_i(b_i^t \mid b_i^t) = P_i(b_i \mid b_i^t) > 0$, define $H_i(b_i^t, b_i) = +\infty$. When $P_i(b_i^t \mid b_i^t) = P_i(b_i \mid b_i^t) = 0$, define $H_i(b_i^t, b_i) = b_i$.

The optimality condition in (1) can then be written in two steps:

$$s_i^t \geq \sup_{b_i < b_i^t} H_i(b_i^t, b_i) \tag{2}$$

$$s_i^t \leq \inf_{b_i > b_i^t} H_i(b_i^t, b_i) \tag{3}$$

The expression $H_i(b_i^t, b_i)$ can be obtained from the data, and in an experimental setting where s_i^t is known it is then possible to test whether bidders do play a NE by checking that (2) and (3) are satisfied. In a non-experimental setting valuations are usually not observed, but it is still possible to test for NE behavior by asking whether for each observed b_i^t there exists a corresponding s_i^t that satisfies (2) and (3).

Now consider the same private values setting where the auctioneer only reveals the winning bid. By Lemma 4, equation (1) is now required to hold only for all $b_i \geq b_i^t$. Hence, in terms of data, equation (3) is required to hold, but not (2). In addition, since expected payoffs must be positive, it follows that

$$s_i^t \geq b_i^t \tag{4}$$

must also hold in a SCE. Hence, (3) and (4) can now be used to test for SCE behavior.[34]

Proposition 11 *(i) A distribution of bids $P \in \mathcal{P}$ is Nash-rationalized by a private values model if and only if for almost every $b^t = (b_1^t, ..., b_N^t) \in \text{supp}\, P$, for all $i \in N$,*

$$\sup_{b_i < b_i^t} H_i(b_i^t, b_i) \leq \inf_{b_i > b_i^t} H_i(b_i^t, b_i).$$

(ii) A distribution of bids $P \in \mathcal{P}$ is SCE-rationalized by a private values model where only the winning bid is observed if and only if for almost every $b^t = (b_1^t, ..., b_N^t) \in \text{supp}\, P$, for all $i \in N$,

$$b_i^t \leq \inf_{b_i > b_i^t} H_i(b_i^t, b_i).$$

By noting that $b_i^t \leq \sup_{b_i < b_i^t} H_i(b_i^t, b_i)$, it follows that, as expected, SCE behavior imposes weaker restrictions on observable bids.

If an observable distribution can be either Nash or SCE rationalized, it is then possible to use the set of s_i^t that satisfy the corresponding inequalities (2), (3), and (4) to estimate the primitive F_S. This estimation method has been pioneered by Guerre, Perrigne, and Vuong (2000) under the assumption of Nash equilibrium. In general, for each b_i^t there will be multiple s_i^t's consistent with SCE behavior, so that methods for bounding the primitive will need to be applied (see Haile and Tamer (2003) for an application of such methods in the context of English auctions).

---

[34] The proof for the Nash equilibrium case is closely related to Paarsch and Robert's (2003) proof in a discrete bid environment. It differs slightly from the original proof in Guerre, Perrigne, and Vuong (2000) since it makes neither regularity assumptions on the distribution of types nor restrictions on the strategy space.

6 Conclusion

Information feedback may affect competitive outcomes through its influence on players' perceptions about their strategic environment. This role of feedback determined both the setting and the choice of self-confirming equilibrium as a solution concept. The findings indicate that relaxing the assumption of Nash equilibrium in a realistic way can provide important insights on previously unexplored issues while still restricting behavior in equilibrium. In fact, unambiguous results were obtained precisely in those settings where auction theory obtains unambiguous comparative statics results. Finally, the approach followed in this paper may be fruitfully applied to obtain novel results in other settings about the effect of information feedback on players' beliefs about their strategic environment.

Appendix

Let B^1, B^2, and $B_i^{\max}$ denote the highest bid, the second highest bid, and the highest opponent bid faced by bidder i given the random variables $B_1, ..., B_N$. Denote their realizations by b^1, b^2, and $b_i^{\max}$.

Proof of Proposition 1. Let β be a SCE. Since $I_i = \{B_i, B^1, B^2\}$, consistency implies that type s_i who bids $b_i = \beta_i(s_i)$ has correct beliefs about the joint distribution of (b_i, B^1, B^2). Since $b_i = b^1$ implies $b_i^{\max} = b^2$ and $b_i \neq b^1$ implies $b_i^{\max} = b^1$, type s_i must also have correct beliefs about $B_i^{\max}$, and consequently about $\Pr(\widehat{b}_i \geq \max_{j \neq i} B_j)$ for all $\widehat{b}_i$. Hence, β is also a NE.

Proof of Proposition 2. Let β be a *SCE*. Let $\underline{b}_i$ denote the infimum of the support of bidder i's equilibrium bids. Suppose, without loss of generality, that $\underline{b}_N \leq ... \leq \underline{b}_2 \leq \underline{b}_1$.

Let $P_i(b_i \mid G_{s_i}) \equiv \Pr(b_i \geq \max_{j \neq i} B_j)$ where B_{-i} is distributed according to G_{s_i}, which is type s_i's belief in the SCE. Since $I_i = \{B_i, B^1\}$, consistency implies that type s_i who bids $\beta_i(s_i)$ has correct beliefs about the joint distribution of $(\beta_i(s_i), B^1)$. Whenever $b_i^{\max} \geq \beta_i(s_i)$, $b_i^{\max} = b^1$. Hence, type s_i must have correct beliefs about $B_i^{\max}$ in the range $b_i^{\max} \geq \beta_i(s_i)$, i.e. $P_i(b_i \mid G_{s_i}) = P_i(b_i \mid G^0_{s_i,\beta_{-i}})$ for $b_i \geq \beta_i(s_i)$. In addition, by the independence assumption it follows that $G^0_{s_i,\beta_{-i}}$ and $P_i(b_i \mid G_{s_i})$ do not depend on s_i – then denote the correct distribution and the belief by $G^0_{\beta_{-i}}$ and G respectively. It follows from consistency and independence that $P_i(b_i \mid G) = P_i(b_i \mid G^0_{\beta_{-i}})$ for $b_i \geq \underline{b}_i$. The proof is completed by considering two cases.

i) Suppose $\underline{b}_2 = \underline{b}_1$. Then for all i, $P_i(b_i \mid G) = P_i(b_i \mid G^0_{\beta_{-i}}) = 0$ for $b_i < \underline{b}_1$, so that all bidders have correct beliefs about their probability of winning with any bid, and therefore β is a NE.

ii) Suppose $\underline{b}_2 < \underline{b}_1$. Then for all $i \neq 1$, $P_i(b_i \mid G) = P_i(b_i \mid G^0_{\beta_{-i}}) = 0$ for $b_i < \underline{b}_1$, so all bidders except 1 have correct beliefs. If bidder 1 also had correct beliefs, then then proof would be completed. So suppose bidder 1 does not have correct beliefs about her probability of winning in the range $b_1 < \underline{b}_1$. Then construct a profile $\widehat{\beta}_{-1}$ that differs from β_{-1} only for those bids that win with probability zero in equilibrium, i.e. $b < \underline{b}_1$. For this case, modify the strategies for bidders $i \neq 1$ so that bidder 1's beliefs are actually correct. Note that these new bids still win with probability zero and so are optimal. Then $\widehat{\beta}$ constitutes a Nash equilibrium profile, since all players are playing best responses, and their beliefs are correct. This implies that $SCE \subset NE$ for profiles that are restricted to those bids that win with strictly positive probability.

Proof of Proposition 3. Let β be a SCE. Since $I_i = (B_i, B^1, B^2, V_i)$, consistency implies that type s_i who bids $b_i = \beta_i(s_i)$ has correct beliefs about the joint distribution of (b_i, B^1, B^2, V_i). Following the argument in the proof of Proposition 1, type s_i must have correct beliefs about the joint distribution of $(B_i^{\max}, V_i)$. From this joint distribution, type s_i can then correctly compute both its probability of winning (as argued in Proposition 1), and its expected value conditional on winning with any bid $\widehat{b}_i$, $E_{B_i^{\max},V_i}\left(V_i \mid \widehat{b}_i \geq B_i^{\max}\right)$. Hence, β is also a NE.

Proof of Lemma 4.

Part (i). *Only if.* Let β be a SCE. Let G_{s_i} be type s_i's belief in equilibrium. As argued in the proof of Proposition 2, since the winning bid is observed it follows that for every s_i, $P_i(b_i \mid G_{s_i}) = P_i(b_i \mid G^0_{s_i,\beta_{-i}})$ for $b_i \geq \beta_i(s_i)$. Hence, $\pi_i(b_i \mid G_{s_i}) = \pi_i(b_i \mid G^0_{s_i,\beta_{-i}})$ for $b_i \geq \beta_i(s_i)$. Therefore, $\beta_i(s_i)$ gives higher expected payoff than choosing any other bid $b_i \geq \beta_i(s_i)$. In addition, players know that they can get a payoff of at least zero by bidding zero, so that $\pi_i(\beta_i(s_i) \mid G^0_{s_i,\beta_{-i}}) \geq 0$. Together, these conditions imply that $\beta_i \in \rho^R_i(\beta)$.

If. Suppose that $\beta_i \in \rho^R(\beta_{-i})$ for all i. For each s_i, consider beliefs G_{s_i} such that $P_i(b_i \mid G_{s_i}) = P_i(b_i \mid G^0_{s_i,\beta_{-i}})$ for $b_i \geq \beta_i(s_i)$, and

$$(s_i - b_i)P_i(b_i \mid G_{s_i}) \leq (s_i - \beta_i(s_i))P_i(\beta_i(s_i) \mid G^0_{s_i,\beta_{-i}}) \tag{A1}$$

for $b_i < \beta_i(s)$, where the LHS of (A1) is nonnegative (which is possible since the RHS is nonnegative by assumption). Since $\beta_i \in \rho^R(\beta_{-i})$, it then follows that $\beta_i(s_i)$ is optimal given beliefs G_{s_i}. It remains to check that such beliefs are consistent when only the winning bid is observed. As argued in the proof of Proposition 2, consistency requires beliefs about expected profit to be correct for bids above the equilibrium bid, and this requirement is satisfied by G_{s_i} constructed above. Finally, since when $b_i^{\max} < b_i$ nothing is observed about other players' bids, consistency places no restriction on beliefs about the probability of winning for bids below the equilibrium bid (except that P_i must be a distribution and in particular cannot be negative, a restriction which can be shown to be satisfied together with (A1) and the fact that that the LHS of (A1) is nonnegative). By optimality and consistency, β is then a SCE.

Part (ii). The proof is analogous to part (i), except that it remains to show that, in the interdependent values model, when only the two highest bids and winners' own valuations are observed, $\pi_i(b_i \mid G_{s_i}) = \pi_i(b_i \mid G^0_{s_i,\beta_{-i}})$ for $b_i \leq \beta_i(s_i)$ in a SCE. To see this, note that as argued in the proof of Proposition 1, type s_i must have correct beliefs about $\Pr(\widehat{b}_i \geq \max_{j \neq i} B_j)$ for all $\widehat{b}_i$ when the two highest bids are observed. In addition, since $(b_i^{\max}, v_i)$ is observed if $b_i^{\max} \leq \beta_i(s_i)$, type s_i has correct beliefs about the joint distribution of $(B_i^{\max}, V_i)$ in the range $b_i^{\max} \leq \beta_i(s_i)$. Hence, type s_i has correct beliefs about $E_{B_i^{\max},V_i}(V_i \mid b_i \geq B_i^{\max})$ for $b_i \leq \beta_i(s_i)$, which establishes the desired result.

Proof of Lemma 5. Consider first the statement that β^{SCE} is increasing. To obtain a contradiction, suppose that β^{SCE} has a flat segment, i.e. $\beta^{SCE}(s) = b \geq 0$ for $s \in [s_1, s_2]$, where $s_1 < s_2$. Consider the behavior of types s_1 and s_2 who respectively make profits $\pi_1 \geq 0$ and $\pi_2 \geq 0$ when everyone plays β^{SCE}, and note that it must be true that $0 \leq \pi_1 < \pi_2$ (since types s_1 and s_2 win under exactly the same conditions on opponents' types, but $u(s_i, \{s_{-i}\})$ is increasing in s_i). Since $\pi_2 > 0$, there exists a small enough $\varepsilon > 0$ such that by deviating to $b + \varepsilon$ type s_2 obtains a discrete jump in her probability of winning by giving up at most an ε cost in terms of lost surplus (at most ε since the expected value of the object will not decrease by increasing her bid). Since the winning bid is revealed and since beliefs are monotone, type s_2 would (correctly) believe that she could obtain a greater expected payoff by increasing her bid to $b + \varepsilon$. Hence, β^{SCE} has no flat segments and by assumption that β^{SCE} is nondecreasing it follows that β^{SCE} must be increasing.

Regarding the initial condition, Milgrom and Weber (1982) show that the valuation of the lowest type $\underline{s}$ when every other opponent is also of type $\underline{s}$ is $\beta^{NE}(\underline{s})$. Suppose that $\lim_{s \to \underline{s}} \beta^{SCE}(s) > \beta^{NE}(\underline{s})$. Then there is a small enough ε' such that type $\underline{s}+\varepsilon'$ wins with strictly positive probability, pays more than her valuation, and therefore makes negative expected utility, which is not possible in a SCE. Suppose instead that $\lim_{s \to \underline{s}} \beta^{SCE}(s) < \beta^{NE}(\underline{s})$. Then a small ε increase in her bid would raise type $\underline{s}$'s expected payoff from zero to some small positive number. Since the winning bid is revealed and since beliefs are monotone, type $\underline{s}$ would then (correctly) believe that she could do better by deviating to such a higher bid.

Proof of Lemma 6. Preliminaries: Denote by

$$v(s,t) = E[u \mid S_i = s, \max_{j \neq i} S_j = t]$$

the expected utility of winning the object for a bidder of type s whose opponents' maximum type is t. Expected utility of type s from bidding b when all other bidders play a strategy β that is continuous and increasing is given by

$$\pi(b \mid G^0_{s,\beta}) = \int_{\underline{s}}^{\phi(b)} (v(s,\alpha) - b) dF_m(\alpha \mid s)$$

where ϕ is the inverse of β and F_m is the conditional distribution of the maximum of opponents' types given own type s, with density f_m. When choosing a best response to β, it is enough to consider bids in the range of β, so that $\beta(x)$ is a best response to β for type s if and only if $x \in \arg\max_{x \in [\underline{s}, \bar{s}]} \pi(\beta(x) \mid G^0_{s,\beta})$. Assuming differentiability, let

$$\pi_b(\beta(x) \mid G^0_{s,\beta}) = \frac{f_m(x \mid s)}{\beta'(x)} \left[v(s,x) - \beta(x) - \frac{F_m(x \mid s)}{f_m(x \mid s)} \beta'(x) \right] \tag{A2}$$

denote the derivative of expected utility with respect to b, evaluated at $\beta(x)$. Then if β^{NE} is a Nash equilibrium, it follows that (A2) is zero when $x = s$ and therefore

$$v(s,s) - \beta^{NE}(s) - \frac{F_m(s \mid s)}{f_m(s \mid s)} \beta^{NE'}(s) = 0. \tag{A3}$$

There are three cases to consider for the increasing function β that satisfies $\lim_{s \to s} \beta(s) = \beta^{NE}(s)$.

Case 1. β is continuous on $[s_1, s_2]$, where $s_1 < s_2$, $\beta(s_1) = \beta^{NE}(s_1)$, and $\beta(s_2) < \beta^{NE}(s_2)$.

Case 2. β is continuous on $[s_1, s_2]$, where $s_1 < s_2$, $\beta(s_1) = \beta^{NE}(s_1)$, and $\beta(s_2) > \beta^{NE}(s_2)$.

Case 3. β is discontinuous at $s_0 > \underline{s}$, $\lim_{s \uparrow s_0} \beta(s) < \beta^{NE}(s_0)$, and $\lim_{s \downarrow s_0} \beta(s) > \beta^{NE}(s_0)$.

Claim 6.1: If $\beta \ngeq \beta^{NE}$, then case 1 holds.

Proof: Suppose that $\beta \ngeq \beta^{NE}$. If there exists $s_2 > \underline{s}$ such that $\beta(s) < \beta^{NE}(s)$ for all $s \in (\underline{s}, s_2)$, then case 1 holds. Else, β must eventually cross β^{NE} and stay below β^{NE} for an interval of types, so that case 1 also holds.

Claim 6.2: If $\beta \nleq \beta^{NE}$, then either case 2 or case 3 hold.

Proof: Suppose that $\beta \nleq \beta^{NE}$. If there exists $s_2 > \underline{s}$ such that $\beta(s) > \beta^{NE}(s)$ for all $s \in (\underline{s}, s_2)$, then case 2 holds. Else, two situations are possible. In one, β eventually crosses β^{NE} and stays above β^{NE} for an interval of types, so that case 2 holds. In the other, β must jump at some point s_0 from a bid below $\beta^{NE}(s_0)$ to a bid above $\beta^{NE}(s_0)$, so that case 3 holds.

Claim 6.3: If case 1 holds, then $\mu\left(s : \beta(s) \notin \arg\max_{b \geq \beta(s)} \pi(b \mid G^0_{s,\beta})\right) > 0$.

Proof: Since β is increasing and by continuity, it follows that β and $\pi(b \mid G^0_{s,\beta^{NE}})$ are differentiable a.e. in $[s_1, s_2]$ and that there exists a set $S^* \subset (s_1, s_2)$ with strictly positive measure such that for all $s^* \in S^*$, (i) $\beta(s^*) < \beta^{NE}(s^*)$ and (ii) $\beta'(s^*) < \beta^{NE\prime}(s^*)$. It follows from (A2) and (A3) that, because of (i) and (ii), the sign of $\pi_b(b \mid G^0_{s^*,\beta})$ is

$$v(s^*, s^*) - \beta(s^*) - \frac{F_m(s^* \mid s^*)}{f_m(s^* \mid s^*)}\beta'(s^*) > 0.$$

Hence, $\beta(s^*)$ cannot be the optimal bid of type s^* given β when restricted to bids higher or equal than $\beta(s^*)$. The claim follows since S^* is a set with strictly positive measure.

Claim 6.4: If either case 2 or 3 holds, then $\mu\left(s : \beta(s) \notin \arg\max_{b \leq \beta(s)} \pi(b \mid G^0_{s,\beta})\right) > 0$.

Proof: If case 2 holds, then the proof is analogous to Claim 6.3 and is therefore omitted. Suppose that case 3 holds. Then there exists $s' > s_0$ such that, for every $s \in (s_0, s')$, reducing the bid from $\beta(s)$ to $\lim_{s \uparrow s_0} \beta(s)$ produces (i) a negligible loss in the probability of winning, (ii) a negligible loss in the expected value of the object conditional on winning (since u is continuous in s_{-i}), and (iii) a non-negligible decrease in the amount being paid for the object. Hence, there is a bid lower than $\beta(s)$ that provides higher utility and the claim follows.

Proof of Lemma 6: Part (a) follows from Claims 6.1 and 6.3, and part (b) follows from Claims 6.2 and 6.4.

Proof of Proposition 7. From Lemma 5, β^{SCE} is increasing and $\lim_{s \to \underline{s}} \beta^{SCE}(s) = \beta^{NE}(\underline{s})$. To obtain a contradiction, suppose that $\beta^{SCE} \ngeq \beta^{NE}$. By Lemma 6, $\beta^{SCE}(s) \notin \arg\max_{b \geq \beta^{SCE}(s)} \pi(b \mid G^0_{s,\beta^{SCE}})$ for a positive measure of types. But then β^{SCE} cannot be a SCE, as characterized by the "only if" part of Lemma 4(i).

Proof of Proposition 9. Suppose everyone plays β^{SCE}. Then, it follows from equation (A2) that the sign of $\pi_b(\beta^{SCE}(x) \mid G^0_{s,\beta})$ is given by the sign of

$$v(s, x) - \beta^{NE}(x) - \alpha(x) - \frac{F_m(x \mid s)}{f_m(x \mid s)}\left[\beta^{NE\prime}(x) + \alpha'(x)\right]. \tag{A4}$$

By (A3), at $x = s$ the expression in (A4) reduces to

$$-\alpha(s) - \frac{F_m(s \mid s)}{f_m(s \mid s)}\alpha'(s) \leq 0.$$

In addition, by the affiliation assumption, $F_m(x \mid s)/f_m(x \mid s)$ is nonincreasing in s (Milgrom and Weber (1982), Theorem 2), so that the expression in (A4) is increasing in s. Hence, for all $x > s$ the expression is negative, implying that when everyone chooses β^{SCE} a choice of $x = s$ (or, $\beta^{SCE}(s)$) yields a higher payoff than a choice of $x > s$ (or $b > \beta^{SCE}(s)$). Finally, since $\alpha(s) < s - \beta^{NE}(s)$, expected payoffs are strictly positive, and therefore β^{SCE} is a SCE by Lemma 4.

Proof of Proposition 10. From Lemma 5, β^{SCE} is increasing and $\lim_{s \to \underline{s}} \beta^{SCE}(s) = \beta^{NE}(\underline{s})$. To obtain a contradiction, suppose that $\beta^{SCE} \nleq \beta^{NE}$. By Lemma 6, $\beta^{SCE}(s) \notin \arg\max_{b \leq \beta^{SCE}(s)} \pi(b \mid G^0_{s,\beta^{SCE}})$ for a positive measure of types. But then β^{SCE} cannot be a SCE, as characterized by the "only if" part of Lemma 4(ii).

Proof of Proposition 11. Part (i). By straighforward algebra from equation (1), b_i^t is optimal if and only if

$$s_i^t \in \left[\sup_{b_i < b_i^t} H_i(b_i^t, b_i), \inf_{b_i > b_i^t} H_i(b_i^t, b_i) \right].$$

If the inequality condition holds, then the above interval is nonempty and we can define a function $f : \mathbb{R}_+^N \longrightarrow \mathbb{R}_+^N$ such that $f(b_1^t, ..., b_N^t) = (s_1^t, ..., s_N^t)$, where for all i, s_i^t belongs to the interval above. Let F_S denote the joint distribution of $f(B_1, ..., B_N)$, where $B_1, ..., B_N$ are distributed according to P. Then by construction P is the NE distribution of bids when the primitive is F_S. Consider now what happens when the condition stated in (i) fails, so that the above interval is empty for a set of bids of strictly positive measure. Then no s_i^t exists for which it is optimal to bid b_i^t, and therefore P is not a Nash equilibrium distribution of bids. The proof of part (ii) is established along similar lines.

References

[1] Athey, Susan. “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information.” *Econometrica*, Vol. 69 (2001), pp. 861-889.

[2] Athey, Susan and Bagwell, Kyle. “Optimal Collusion with Private Information.” *RAND Journal of Economics*, Vol. 32 (2001), pp. 428-465.

[3] Athey, Susan, Bagwell, Kyle, and Sanchirico, Chris. “Collusion and Price Rigidity.” *Review of Economic Studies*, Vol. 71 (2004), pp. 317-349.

[4] Athey, Susan and Haile, Philip A. “Identification in Standard Auction Models.” *Econometrica*, Vol. 70 (2002), pp. 2107-2140.

[5] Athey, Susan and Haile, Philip A. “Nonparametric Approaches to Auctions.” In J. Heckman and E. Leamer, eds., *Handbook of Econometrics*, Volume 6, forthcoming.

[6] Athey, Susan and Haile, Philip A. “Empirical Models of Auctions.” Ninth World Congress of the Econometric Society, London, 2006.

[7] Battigalli, Pierpaolo. “Comportamento Razionale ed Equilibrio nei Giochi e Nelle Situazioni Sociali.” Unpublished dissertation, Universita Bocconi, Milano, 1987.

[8] Cox, James C., Roberson, Bruce, and Smith, Vernon L. “Theory and Behavior of Single-Object Auctions.” In Vernon L. Smith, ed., *Research in Experimental Economics* Greenwhich, CT: JAI Press, 1982.

[9] Crawford, Vincent. P. and Iriberri, Nagore. “Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions?” *Econometrica*, Vol. 75 (2007), pp. 1721-1770.

[10] Dekel, Eddie, Fudenberg, Drew, and Levine, David K. “Learning to Play Bayesian Games.” *Games and Economic Behavior*, Vol. 46 (2004), pp. 282-303.

[11] Dufwenberg, Martin and Gneezy, Uri. “Information Disclosure in Auctions: an Experiment.” *Journal of Economic Behavior and Organization*, Vol. 48 (2002), pp. 431-444.

[12] Engelbrecht-Wiggans, Richard and Katok, Elena. “Regret in Auctions: Theory and Evidence.” *Economic Theory*, Vol. 33 (2007), pp. 81-101.

[13] Esponda, Ignacio. “Behavioral Equilibrium in Economies with Adverse Selection.” *American Economic Review*, forthcoming.

[14] Esponda, Ignacio. “Rationalizable Conjectural Equilibrium in Games with Uncertainty.” Working paper, New York University, 2007.

[15] Filiz, Emel and Ozbay, Erkut Y. “Auctions with Anticipated Regret: Theory and Experiment.” American Economic Review, Vol. 97 (2007), pp. 1407-1418.

[16] Friedman, Lawrence. “A Competitive Bidding Strategy.” *Operations Research*, Vol. 4 (1956), pp. 104-112.

[17] Fudenberg, Drew and Kreps, David M. “Learning in Extensive Form Games. I. Self Confirming Equilibrium.” *Games and Economic Behavior*, Vol. 8 (1995), pp. 20-55.

[18] Fudenberg, Drew and Levine, David K. “Self-Confirming Equilibria.” *Econometrica*, Vol. 61 (1993a), pp. 523-45.

[19] Fudenberg, Drew and Levine, David K. “Steady State Learning and Nash Equilibrium.” *Econometrica*, Vol. 61 (1993b), pp. 547-573.

[20] Guerre, Emmanuel, Perrigne, Isabelle, and Vuong, Quang. “Optimal Nonparametric Estimation of First-Price Auctions.” *Econometrica*, Vol. 68 (2000), pp. 525-574.

[21] Haile, Philip and Tamer, Elie. “Inference with an Incomplete Model of English Auctions.” *Journal of Political Economy*, Vol. 111 (2003), pp. 1-51.

[22] Hendricks, Kenneth, Pinkse, Joris, and Porter, Robert. “Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions.” *Review of Economic Studies*, Vol. 70 (2003), pp. 115-145.

[23] Jehiel, Philippe. “Manipulative Auction Design.” Working paper, 2007.

[24] Kagel, John H. “Auctions: A survey of Experimental Research.” In John H. Kagel and Alvin E. Roth, eds., *Handbook of Experimental Economics* Princeton: Princeton University Press, 1995.

[25] Isaac, R. Mark and Walker, J. M. “Information and Conspiracy in Sealed Bid Auctions.” *Journal of Economic Behavior and Organization*, Vol. 6 (1985), pp. 139-159.

[26] Laffont, Jean-Jacques. “Game Theory and Empirical Economics: the Case of Auction Data.” *European Economic Review*, Vol. 41 (1997), pp. 1-35.

[27] Laffont, Jean-Jacques, Ossard, Herve, and Vuong, Quang. “Econometrics of First Price Auctions.” *Econometrica* Vol. 63 (1995), pp. 953-980.

[28] Maskin, Eric and Riley, John. “Equilibrium in Sealed High Bid Auctions.” *Review of Economic Studies,* Vol. 67 (2000), pp. 439-454.

[29] McAdams, David. “Uniqueness in Symmetric First Price Auctions with Affiliation.” *Journal of Economic Theory*, Vol. 136 (2007), pp. 144-166.

[30] Milgrom, Paul and Roberts, John. "Comparing Equilibria." *American Economic Review*, Vol. 84 (1994), pp. 441-459.

[31] Milgrom, Paul R. and Weber, Robert J. "A Theory of Auctions and Competitive Bidding." *Econometrica*, Vol. 50 (1982), pp.1089-1122.

[32] Neugebauer, Tibor and Perote, Javier. "Theory and Misbehavior of First-Price Auctions: the Importance of Information Feedback in Experimental Markets." *Experimental Economics*, forthcoming.

[33] Neugebauer, Tibor and Selten, Reinhard. "Individual Behavior of First-Price Auctions: the Importance of Information Feedback in Computerized Experimental Markets." *Games and Economic Behavior*, Vol. 54 (2006), pp. 183-204.

[34] Ockenfels, Axel and Selten, Reinhard. "Impulse Balance Equilibrium and Feedback in First Price Auctions." *Games and Economic Behavior*, Vol. 51 (2006), pp. 155-170.

[35] Paarsch, Harry J. and Robert, Jacques. "Testing Equilibrium Behaviour at Frist-price, Sealed-bid auctions with Discrete Bid Increments." Working paper, University of Iowa, 2003.

[36] Park, William R. *Construction Bidding for Profit*. John Wiley & Sons, 1979.

[37] Porter, Robert. "The Role of Information in US Offshore Oil and Gas Lease Auctions." *Econometrica*, Vol. 63 (1995), pp. 1-27.

[38] Schaik, F.D.J. van and Kleijnen, J.P.C. "Sealed Bid Auctions: Case Study." Tilburg University Discussion Paper nr. 73, 2001.

[39] Skrzypacz, Andrzej and Hopenhayn, Hugo. "Tacit Collusion in Repeated Auctions." *Journal of Economic Theory*, Vol. 114 (2004), pp. 153-169.

[40] Stigler, George J. "A Theory of Oligopoly." *Journal of Political Economy*, Vol. 72 (1964), pp. 44-61.

[41] Thomas, Charles J. "Information Revelation and Buyer Profits in Repeated Procurement." Working paper, Clemson University, 2006.

[42] Zhang, S., Karunamuni, R. J., and Jones, M. C. "An Improved Estimator of the Density Function at the Boundary." Journal of the American Statistical Association, Vol. 94 (1999), pp. 1231-1241.