

Herding and the Winner's Curse in Markets with Sequential Bids*

Zvika Neeman[†]

*Department of Economics, Boston University, 270 Bay State Road,
Boston, Massachusetts 02215
zvika@BU.edu*

and

Gerhard O. Orosel

*Department of Economics, University of Vienna, Bruenner Strasse 72,
A-1210 Vienna, Austria
orosel@econ.bwl.univie.ac.at*

Received September 24, 1997; revised November 5, 1998

We present a model of social learning in an environment with common values where informational cascades and herding arise in combination with the winner's curse. A seller of an object sequentially obtains bids from potential buyers. We characterize three classes of equilibria that differ widely in their information aggregation properties and in the size of the rent the seller captures from the buyers. We compare the procedure of sequentially soliciting bids from the buyers to conducting an English auction for the object in terms of maximization of seller's revenue and demonstrate the superiority of the former. *Journal of Economic Literature* Classification Numbers: D82, D83. © 1999 Academic Press

1. INTRODUCTION

We present a model of social learning in market-like environments with common values. We show that the phenomena of informational cascades and herding, where the available public information swamps the agents'

* We are grateful to Itzhak Gilboa, Rudolf Kerschbamer, Muriel Niederle, Andreas Ramsauer, and Gerhard Sorger for helpful discussions. A referee and an associate editor contributed valuable comments. The second author thanks the Department of Economics at Boston University for its hospitality.

[†] World Wide Web: <http://econ.bu.edu/neeman/>.

private information and induces them to behave identically, arise in combination with the winner's curse.¹

We consider a model where the owner of an object (henceforth, the "seller") sequentially obtains bids for the object from a finite number of potential buyers. All buyers have the same ex post valuation of the object. They differ, however, in their estimates of this common value. The basic structure of our model is similar to the one considered by the herding literature.² However, our approach differs from the standard herding literature in three respects. First, the order of approaching the buyers is determined by the seller and hence is endogenous to the model. Second, the action spaces of the agents who have private information are continuous sets whereas in the rest of the literature the herding results depend on these sets being discrete (see [5]). Third, and most importantly, the payoffs to the buyers and to the seller depend on the actions of the other agents in the model. That is, besides the informational externality that characterizes the herding literature, there is also a payoff externality, and as a consequence, the winner's curse becomes relevant to the analysis.

The model presented is flexible enough to incorporate a large variety of cases: from an inventor who is searching for a venture capital firm to finance the development of his idea into a marketable product, to a movie director who needs a producer to realize a script, to a university graduate in search of a job. Since the price of the object may be negative, the model also applies to the case of a government agency that is looking for a contractor to perform a public project. Our results help shed light on a number of (casual) empirical observations such as the documented failures of inventors to sell useful inventions. We show that such failures are not necessarily due to the fact that no information about the true quality of the object exists but rather may occur because the market fails to successfully aggregate all the available information. Thus, for example, one of two inventions with similar commercial potential may be sold for an extremely high price, whereas the other may not be sold at all. Similarly, out of a set of comparable university graduates, one may become a "star" in the job market, others may do well, but the rest might fail to get any offers.

¹ An *informational cascade* occurs when the (observable) actions reveal no information about the private signals. Herding occurs when agents behave identically, independently of their private signals.

² [1] and [2] are the classic references. For surveys of different aspects of herding literature see [4], [18], and [3]. Another, especially general, treatment is presented by [17]. [19] considers a herding model that is applied to the sale of IPO shares. It examines a sequential model with one seller (the issuer) and many buyers (the investors). However, while in [19] the seller offers IPO shares to potential buyers at a price that the seller chooses optimally, in our model the buyers bid optimally for the object.

This diversity of outcome in the model arises as a consequence of path dependent informational cascades.

We focus our analysis on three types of equilibria that differ by their information aggregation properties and by the size of the surplus the seller can extract from the buyers. We demonstrate the existence of a “backwards induction equilibrium” which does not exhibit informational cascades or herding and where, except for one particular case, all the available information is used efficiently. In this equilibrium, except for this one particular case, the seller succeeds in capturing all of the buyers’ information rents. Next, we show that there exists an equilibrium where an informational cascade and herding occur with a positive probability that tends to 1 as the number of potential buyers increases. Nevertheless, the seller succeeds in capturing almost the entire surplus in this equilibrium as well. Finally, we demonstrate the existence of a third family of equilibria where no information is revealed and the sale price may be pushed down to the seller’s reservation value or even below, in which case no trade occurs. We discuss the dimensions on which these equilibria can be compared, and the issues of multiplicity and selection of equilibria, at the end of Section 3.

We compare the procedure of sequentially soliciting bids from the buyers to conducting an English auction for the object in terms of the expected payoff to the seller. We show that the sequential procedure presented in this paper is superior. Under three different assumptions on the seller’s behavior, we demonstrate that the seller’s expected payoff under the sequential procedure is either higher than or equal to that generated by an English auction. Under the English auction, buyers who have observed high signals can step forward and bid up the price of the object. On the other hand, under the sequential procedure, the buyers can only bid the price up when approached by the seller. One may expect then that the English auction would be better at aggregating the buyers’ information but our results show the opposite. Because buyers can step forward under the English auction, the winner’s curse is more severe under the English auction, and as a consequence, the sequential procedure is better at aggregating the buyers’ information and at generating higher revenues for the seller. In addition, whereas in the English auction, the seller is obliged to sell the object if buyers bid the price above the initial price, in the sequential procedure, the seller has the advantage of the possibility to refuse to sell the object if he does not think the price is high enough.

Our results relate to those obtained in the literature that studies information aggregation in common value auctions.³ In our model, there exist equilibria where the information contained in all private signals is fully aggregated (except, possibly, for the information contained in one single

³ See, e.g., [20], [8], [9] and [12].

private signal). In the limit, as the number of potential buyers increases, complete information aggregation is achieved even when Milgrom's [8] *necessary* (and sufficient) condition for information aggregation in a first price auction under common values is not satisfied. On the other hand, when the number of buyers is large, there exist more compelling herding equilibria where private signals are never fully aggregated. Since our model does not satisfy Milgrom's [8] sufficient condition for full aggregation of private signals (that requires the existence of signals that are strong enough to overturn any previously held belief), this result is not, perhaps, very surprising.⁴ But Milgrom's condition is very strong, and our model does satisfy the weaker sufficient condition of [12] that consider a variation of Milgrom's model where instead of one, k objects are for sale.⁵

Finally, the sequential procedure that is described in this paper is similar to those described in the search literature.⁶ The important difference is that in the search literature the bids are usually assumed to be independent and no informational or payoff externalities are present. In contrast, in this paper the buyers' bids for the object are correlated for (at least) three reasons: the correlation in buyers' signals, the fact that the object has the same common value to all the buyers, and because buyers condition their bids for the object on publicly observed history. As a consequence, the optimal strategy for the seller cannot be described by a reservation value rule as in the search literature. Another consequence is that even without discounting or search costs, in some equilibria the search may stop before all potential buyers have been approached.

The rest of the paper proceeds as follows. In the next section we present the model and some preliminary results. In Section 3, we explore the range of equilibrium behavior. In Section 4, we compare the procedure of sequentially soliciting bids to standard auctions. All proofs are collected in the Appendix.

2. THE MODEL

We consider the problem of a seller of an object of unknown quality who faces $N \geq 2$ potential buyers. The set of buyers is denoted \mathcal{B} . We denote the quality of the object by q , and assume that it is either high or low. That is, $q \in \{q_L, q_H\}$ where q_H denotes high and q_L denotes low quality. We

⁴ See [17] for an analogous condition that assures that permanent herding on a wrong decision cannot occur.

⁵ In [12] Pesendorfer and Swinkels provide a very weak sufficient condition for asymptotic information aggregation but require both the number of objects that are up for sale, k , and the number of bidders, n , to tend to infinity.

⁶ For an extensive survey, see [7] and the references therein.

assume that it is common knowledge among the seller and buyers that the probability that the object is of a high quality is $0 < \Pr(q_H) < 1$.

We denote the seller's valuation (i.e., his reservation value) of the object by a_0 . Hence, his payoff from selling the object for a price p is $p - a_0$. The seller's payoff if he keeps the object is 0. The payoff to a buyer from purchasing the object at a price p is given by $v(q) - p$ where $v(q_H) > v(q_L)$. The buyers' payoffs when they do not obtain the object (and do not pay) are normalized to zero. We assume that $v(q_H) > a_0$; otherwise, the seller cannot hope to succeed in selling the object for a price that is mutually acceptable to himself and the buyers. Note that we allow for the case where $v(q_L) < a_0$, so that trade need not always be welfare enhancing. The buyers know the seller's reservation value a_0 .

At each point of time $t \in \{1, 2, \dots\}$ the seller may

- (a) search for a "new" buyer $b \in \mathcal{B}$ from whom to solicit a bid for the object (provided, of course, that the seller has not yet exhausted the set \mathcal{B});
 - (b) re-approach a buyer that has already been approached before and ask her to make another bid;
 - (c) accept one of the bids made prior to time t and sell the object;
- or
- (d) end the search without selling the object.

When approached by the seller, each new buyer inspects the object and obtains a private signal about its quality. The buyer's signal expresses her subjective impression of the value of the object. We denote buyer b 's signal by $S^b \in \mathcal{S}$, where $\mathcal{S} = \{s_L, s_H\}$ denotes the signal space. We let S^b denote buyer b 's random signal and s^b its realization. We assume that conditional on the true quality of the object, q , buyers' signals are independent and identically distributed and that for all buyers $b \in \mathcal{B}$,

$$0 < \Pr(S^b = s_H | q_L) < \Pr(S^b = s_H | q_H) < 1.$$

After being approached by the seller and observing their signals and the "history" of the buyers' and the seller's actions (formally defined below), buyers choose an action from the set $\mathbb{R} \cup \{-\infty\}$. We use superscripts to denote buyers' identities and subscripts to denote time. Thus, choosing action $a_t^b \in \mathbb{R}$ implies that, at time t , buyer b is offering to pay a_t^b for the object; we refer to such actions as *bids* or *offers*. Choosing action " $-\infty$ " implies that the buyer declines to make an offer for the object. We refer to a buyer that is re-approached by the seller as an "old" buyer. Old buyers do not receive an additional signal when re-approached. However, as noted above, they observe the actions of all preceding buyers.

The *history* of buyers' actions at time t , $t \in \{1, 2, \dots\}$, is denoted by $h_t = (a_1^{b(1)}, \dots, a_t^{b(t)})$ where $b(\tau)$ denotes the identity of the buyer that is approached by the seller at time τ . At time 0, the history h_0 is given by the empty set. The set of all possible histories is denoted by \mathcal{H} . We denote the maximal bid the seller has received by time t by $p(h_t) = \max\{a_\tau^{b(\tau)} \mid 1 \leq \tau \leq t\}$, and let $p(h_0) = -\infty$. We assume that at any point in time, new buyers must either make a bid that is larger than or equal to $p(h_t)$, or refuse to make a bid. This assumption is based on the idea that a buyer that realizes that her bid will be lower than a previous bid, understands that the seller will never sell her the object and therefore declines to make a bid altogether. When re-approached, old buyers may either repeat their previous offer or make a bid larger than or equal to $p(h_t)$.

A pure strategy for the seller is a function $A^p: \mathcal{H} \rightarrow (\mathcal{B} \times \mathbb{R}) \cup \mathcal{B} \cup \{\text{"stop"}\}$ that maps every history h_t into the actions available to the seller. A choice of $(b, a_\tau^b) \in \mathcal{B} \times \mathbb{R}$ implies that the seller sells the object to buyer b at the price specified in her last bid a_τ^b ; a choice of $b \in \mathcal{B}$ implies that the seller solicits a new bid from buyer b that may be a new or old buyer; and a choice of "stop" implies that the seller gives up the attempt to sell the object and stops soliciting offers from buyers without selling the object. A pure strategy for buyer b is described by a function $A^b: \mathcal{H} \times \mathcal{S} \rightarrow \mathbb{R} \cup \{-\infty\}$ that maps histories and signals into the actions available to the buyer. By assumption, $A^b(h_t, s^b) \in [p(h_t), \infty) \cup \{-\infty\}$ unless b is an old buyer in which case $A^b(h_t, s^b) \in [p(h_t), \infty) \cup \{a_\tau^b\}$ where a_τ^b is the action chosen by the buyer when she was previously approached by the seller at time $\tau < t$. Mixed strategies for the seller and buyers are defined in the usual way.

It is useful to explicitly introduce the players' beliefs about the relationship between the buyers' actions and their observed signals (types). A belief for player i is given by a function $\Gamma^i: \mathcal{H} \rightarrow \bigcup_{t \geq 1} [0, 1]^t$ that maps every history h_t into a vector of probabilities $(\gamma_1^i, \dots, \gamma_t^i)$ where γ_τ^i is the probability that player i assigns to the event that buyer $b(\tau)$ has observed the signal s_H .

We rely on the notion of perfect Bayesian equilibrium (see, e.g., [11, p. 231–233]) as our equilibrium notion.

DEFINITION 1. A profile of strategies and beliefs is a perfect Bayesian equilibrium (PBE) of the sequential bidding game above if (1) for every possible history h_t the seller's strategy maximizes the seller's expected payoff given his beliefs and the buyers' strategies; (2) for every possible history h_t , buyer b 's strategy maximizes her expected payoff conditional on her observed signal given her beliefs and the seller's and other buyers' strategies; and (3) whenever possible, beliefs are updated according to Bayes' rule.

Given the seller's and buyers' strategies, a player may be able to infer some (possibly all) of the previous buyers' signals from the history of actions. Note that if buyers adopt pure strategies, then their actions cannot partially reveal their signals. Either an action completely reveals the respective buyer's signal, or the action is completely uninformative about this signal. Suppose that buyers adopt pure strategies. Let $n_L(h_t)$ denote the number of low signals and let $n_H(h_t)$ denote the number of high signals that can be inferred from the history h_t . Let $P_H(n_L, n_H)$ denote the probability that the object is of high quality conditional on the fact that, out of $n_L + n_H$ signals, n_L are low and n_H are high.

LEMMA 1. $P_H(n_L, n_H)$ is strictly decreasing in n_L and strictly increasing in n_H . In addition, there exists a unique integer $k \geq 1$ such that

$$P_H(n_L + k, n_H + 1) \leq P_H(n_L, n_H) < P_H(n_L + k - 1, n_H + 1)$$

for every n_L and n_H . This integer is given by

$$k = \min \left\{ \ell \in \{1, 2, \dots\} \mid \frac{\Pr(s_H | q_H)}{\Pr(s_H | q_L)} \leq \left(\frac{\Pr(s_L | q_L)}{\Pr(s_L | q_H)} \right)^\ell \right\}.$$

For many specifications of the signal structure, the integer k that is defined in the previous lemma is a small number. We demonstrate this in the following two examples.

EXAMPLE 1. When the signal structure is symmetric, that is, when $\Pr(s_H | q_H) = \Pr(s_L | q_L) = \alpha$ and $\Pr(s_L | q_H) = \Pr(s_H | q_L) = 1 - \alpha$, where $\frac{1}{2} < \alpha < 1$, $k = 1$.

EXAMPLE 2. Assume $\Pr(s_H | q_H) = \Pr(s_L | q_H) = \frac{1}{2}$, $\Pr(s_L | q_L) = (n-1)/n$ and $\Pr(s_H | q_L) = \frac{1}{n}$, where $n > 2$. Then, $k = 2$ for $n = 4$; $k = 3$ for $n = 10$; and $k = 6$ for $n = 100$.

Let

$$V(h_t) = E[v(q) \mid (A^p, A^b)_{b \in \mathcal{B}}, (\Gamma^p, \Gamma^b)_{b \in \mathcal{B}}, h_t]$$

denote the expected valuation of the object given the (public) information that is contained in the history h_t .

LEMMA 2. $V(h_t)$ is a martingale. That is,

$$E[V(h_{t+1}) \mid V(h_t)] = V(h_t)$$

for every perfect Bayesian equilibrium profile of strategies $(A^p, A^b)_{b \in \mathcal{B}}$, beliefs $(\Gamma^p, \Gamma^b)_{b \in \mathcal{B}}$ and history $h_t \in \mathcal{H}$.

We introduce the notions of an informational cascade and of herding.

DEFINITION 2. An informational cascade occurs at time T if no information about the quality of the object is revealed at time T or after.

DEFINITION 3. Herding occurs at time T if all buyers, if and when approached, refuse to bid up the price at time T or after.

Although there is no necessary relationship between the notions of an informational cascade and herding, these two phenomena may reinforce each other.⁷ Buyers may refuse to bid up the price because they do not receive any information that justifies a higher price and refusing to bid up the price prevents information about the quality of the object from being transmitted to other buyers. In efficiency terms, the significance of informational cascades and herding is that together they may induce the seller to inefficiently give up on selling the object despite the fact that it has a high quality.

It is important to note that in contrast to other herding models, in this model the informational cascades and herding may remain invisible to an outside observer because as soon as the seller realizes that future buyers will refuse to reveal additional information about the object or bid up the price he may refuse to approach any new buyers.

3. EQUILIBRIUM BEHAVIOR

The main problem that the seller faces is the difficulty to get buyers to bid the price up to the expected value $V(h_t)$ of the object conditional on the public information and their private signal. The problem is that buyers know that if they bid $V(h_t)$ for the object, their bid gives the seller a valuable option. Specifically, a buyer b that considers bidding the price up to $V(h_t)$ realizes that the seller may subsequently approach a new buyer. If this new buyer observes a high signal, she may outbid buyer b , but if she observes a low signal, she will decline to make a bid for the object and the seller may sell the object to buyer b at the price $V(h_t)$ which, because an additional low signal has been observed, is higher than the object's (updated) value $V(h_{t+1})$. Thus, the fact that buyers' bids have an option

⁷ Note that buyers may indicate their signals without bidding up the price (by, say, matching the price if they observe a high signal and refusing to bid if they observe a low signal) and buyers may bid the price up in a way that is independent of their signal.

value for the seller gives rise to a *winner's curse*: being selected to buy the object may imply that the seller has approached additional buyers who observed low signals and therefore the object is worth less in expectation than the price that is paid for it.

In the next three subsections, we present three types of pure strategy Bayesian perfect equilibria that illustrate the range of equilibrium behavior. In the fourth subsection, we comment on the plausibility of these equilibria and the way it depends on the number of potential buyers.

3.1. *The Backwards Induction Equilibrium*

The first PBE we consider is one where all the information about the quality of the object is revealed (except, possibly, for one single private signal). Furthermore, the price at which the object is sold converges to the true value of the object conditional on buyers' information as the number of buyers increases.

Note that although the number of buyers is finite, the fact that the seller can re-approach old buyers implies that the game may still be infinite. Nevertheless, the game has a perfect Bayesian equilibrium which closely resembles a backwards induction equilibrium. Along the equilibrium path of this "backwards induction" perfect Bayesian equilibrium, the seller approaches the buyers sequentially and every buyer that is approached (except for the last one) takes account of the winner's curse by making a bid equal to the expected valuation of the object conditional on public information, her own signal, and the assumption that all the buyers that have not yet been approached by the seller observe low signals. After obtaining bids from all the buyers, the seller sells the object to the buyer who has made the highest bid, provided this bid is higher than his reservation value.

We introduce the following notation. Let

$$p^N = E[v(q) \mid \text{all } N \text{ buyers observe low signals}]$$

denote the lowest possible valuation of the object conditional on buyers' information, and let

$$\begin{aligned} &W_N(h_{t-1}, s^{b(t)}) \\ &= E \left[v(q) \mid (A^b)_{b \in \mathcal{B}}, (\Gamma^b)_{b \in \mathcal{B}}, h_{t-1}, s^{b(t)}, \text{ and all the buyers who} \right. \\ &\quad \left. \text{have not yet revealed their signals observe low signals} \right] \end{aligned}$$

denote the expected value of the object conditional on public information at time t , buyer $b(t)$'s own signal and all the buyers whose signals are still unknown observing low signals.

PROPOSITION 1. *There exists a perfect Bayesian equilibrium in pure strategies where along the equilibrium path: (i) all buyers except the last one bid up the price to $W_N(h_{t-1}, s^b)$ when they observe a high signal, and decline to make a bid when they observe a low signal; (ii) the last buyer behaves like previous buyers unless $a_0 < p^N$ and all previous buyers have observed low signals, in which case she bids p^N independently of her signal; (iii) after soliciting bids from all the buyers, the seller sells the object to the buyer who has made the highest bid, provided it is larger or equal to a_0 ; he sells the object to the last buyer if the highest bid is p^N and $a_0 < p^N$; and (iv) unless $a_0 < p^N$, all buyers except the last one observe low signals and the last buyer observes a high signal, the seller sells the object at a price equal to the value of the object conditional on all the buyers' information.*

The proposition shows that if $a_0 \geq p^N$ or at least one buyer except for the last one observes a high signal, the sale price perfectly aggregates all the buyers' information. When it does not, the sale price still incorporates the information of the first $N - 1$ buyers. In this case, the last buyer earns a positive rent when she observes a high signal. She pays p^N although the object is worth $E[v(q) | N - 1 \text{ signals are low, 1 signal is high}] > p^N$ to her. This is the only case where a buyer succeeds in capturing a positive rent from the seller.

The equilibrium outcome is always efficient. The object is sold if and only if its value to the buyers conditional on all the buyers' information is larger than the seller's reservation value a_0 .

To understand the equilibrium and the special role played by the last buyer, recall that as discussed above, the seller has to overcome the difficulty of inducing buyers to reveal the fact they have observed high signals. By insisting on selling only to buyers who have revealed themselves to observe high signals, the seller can solve this problem for all the buyers except for the last one. But, if the seller's reservation value is sufficiently low (i.e., if $a_0 < p^N$) and if all the buyers before the last buyer have revealed that they have observed low signals, the last buyer has an incentive to pretend she has observed a low signal even when her signal is high. In this case, the seller cannot extract the informational rent from the last buyer and has no choice but to sell her the object at a price equal to her bid, p^N . This observation relies on two subtle points. First, for the seller, the alternative strategy of insisting that the last buyer bids the price up to her valuation when she observes a high signal is not subgame perfect because independently of her signal it is always optimal for the buyer to refuse to bid up the price (if she bids up the price she reveals that her signal is high and consequently her rent will be zero). Second, since the last buyer does not reveal her signal, the expected value of the object conditional on public information is higher than the expected value conditional on all N buyers

observing low signals, which is the price at which the object is sold. Nevertheless, no buyer can benefit by deviating and bidding up the price because if she does, the last buyer will bid the price even higher if she observes a high signal but decline to bid if she observes a low signal and thus expose the deviating buyer to the winner's curse.

We obtain the following result on the relationship between the sale price and the true value of the object.

PROPOSITION 2. As the number of buyers N increases, the transaction price under the perfect Bayesian equilibrium described in Proposition 1 converges to the true value of the object with probability 1.

3.2. The Herding Equilibrium

In this subsection we present perhaps the most interesting of the three equilibria we consider. Recall that an informational cascade means that from some point in time onwards no information about the buyers' signals is revealed and that herding is defined as a situation where all buyers refuse to bid up the price. Accordingly we say that a buyer herds if she declines to bid up the price independently of her signal.

We present a PBE where, initially, buyers who observe high signals bid up the price, buyers who observe low signals decline to bid, and herding and an informational cascade occur with a positive probability that approaches one as the number of buyers, N , tends to infinity. When herding occurs, the seller sells the object at a price equal to its expected value conditional on public information provided it is not below his reservation value; otherwise, if it is below his reservation value, the seller ends the search and retains the object.

The intuition underlying this equilibrium is easier to explain under the provisional assumption that there are infinitely many buyers. This implies that the seller can always approach a new buyer. Since there are many buyers and only one seller, the price competition among the buyers induces them to bid up the price to the value conditional on history, their private signal (which is revealed by their bid) and on the information implied by winning the object (i.e., the winner's curse). When the number of buyers is infinite, there exists a PBE where buyers who observe a low signal decline to make a bid for the object; unless herding has already occurred, buyers who observe a high signal bid up the price while avoiding the winner's curse by bidding the expected value of the object conditional on past history, their signal, and k additional low signals where k is the integer identified in Lemma 1; and a high signal followed by a sequence of k low signals generates herding and an informational cascade.⁸ Suppose that

⁸ See Proposition 9 in the appendix, in the proof of Proposition 3.

buyer b is approached at time t and has observed a high signal. Given the strategies of the other buyers, she will win the object only if her high signal is followed by an uninterrupted sequence of k low signals. Because of Lemma 1, a buyer that is approached after buyer b and observes a high signal before k new low signals have been inferred, will bid up the price; but after k low signals have been observed, every buyer will herd. This follows since after the k additional low signals the bid of buyer b exactly equals $V(h_{t+k})$. The next buyer (i.e., the buyer that is approached after k low signals have been observed), even if she has observed a high signal, will refuse to bid up the price because bidding the value of the object conditional on the history h_{t+k} , her private signal $s^{b(t+k+1)}$ and the winner's curse (i.e., k additional low signals) results in a lower bid than $V(h_{t+k})$. Consequently, she will herd. The same argument implies that every subsequent new buyer that is approached by the seller will herd as well. Taking this into account, the seller will not bother to solicit offers from additional buyers after a sequence of k low signals is observed.

The PBE presented in the next proposition relies on the intuition developed for the case of an infinite number of buyers to motivate and rationalize the strategies of the first $N - 2k - 1$ buyers approached by the seller in a model with a finite number of buyers. The strategies of the last $2k + 1$ buyers are modified to take into account the fact that the number of buyers is finite.

PROPOSITION 3. *If $N \geq 2k + 3$, there exists a perfect Bayesian equilibrium in pure strategies where along the equilibrium path: (i) An informational cascade and herding are generated if and only if at least one of the first $N - 2k - 1$ buyers observes a high signal. Thus, the probability of an informational cascade and herding is positive and approaches 1 as $N \rightarrow \infty$. (ii) The seller continues to approach new buyers until an informational cascade and herding occur. If herding does not occur, the seller approaches all the buyers. (iii) After herding has occurred or the seller has approached all the buyers, the seller sells the object to the buyer who made the highest bid provided her bid is larger or equal to a_0 . The seller retains the object otherwise. Finally, (iv) whenever $a_0 \geq p^N$ or at least one buyer except for the last one observes a high signal, the sale price is equal to the value of the object conditional on public information.*

As in the PBE described in Proposition 1, in the PBE described above in Proposition 3 the last buyer may capture some rent from the seller, but only if all buyers before her have observed low signals, she observed a high signal, and $a_0 < p^N$.

The PBE presented in Proposition 3 combines the intuition about buyers' behavior when there is an infinite number of buyers with the

intuition developed in the backwards induction equilibrium described in Proposition 1. The last $2k + 1$ buyers adopt strategies that are identical to the strategies they adopt in the PBE of Proposition 1. However, they are approached by the seller only in case herding does not occur. The first $N - 2k - 1$ buyers behave as they do in the PBE described in Proposition 9 where the number of buyers is infinite. An informational cascade and herding occur if and only if at least one of the first $N - 2k - 1$ buyers observes a high signal. Once this happens, an informational cascade and herding occur as soon as either the next k new buyers approached by the seller decline to bid for the object thereby revealing themselves to have observed low signals, or buyer $N - 2k$ is approached by the seller, whichever occurs first. If herding has not occurred by the time buyer $N - 2k$ is approached by the seller, that is, all previous buyers have observed low signals, then play reverts to the backwards induction equilibrium described in Proposition 1.

The important difference between the PBE described in Proposition 3 and the one described in Proposition 1 is that the former may be inefficient. Herding and an informational cascade may occur after a bid that is below the seller's reservation price a_0 . The seller's equilibrium strategy calls for not selling the object in spite of the fact that its real value for the buyers (conditional on the information of all the buyers) may be higher than a_0 . It may also be the case that herding and an informational cascade occur after a bid that is higher than a_0 and the seller sells the object, but the object's valuation conditional on the signals of all the buyers is lower than a_0 .

3.3. Uninformative Equilibria

In this subsection we describe a third type of equilibrium where no information is revealed and the highest bid can be as low as the seller's reserve price or even lower. In the latter case, no trade occurs and the seller retains the object. If, in addition, $a_0 < p^N$, this is inefficient.

The intuition behind this equilibrium is as follows. Suppose the current bid is below p^N . Buyers refuse to bid up the price because equilibrium strategies are such that if they bid the price up but still below p^N , other buyers "retaliate" by bidding the price even further, so they have no chance of buying the object at a price equal to their bid. And if they bid the price up above p^N they expose themselves to the winner's curse, so that, at best, they can only hope to buy the object at a price equal to its expected value. In this equilibrium, no information is ever revealed and an informational cascade and herding occur immediately.

PROPOSITION 4. *Fix a number $p \leq p^N$. There exists a perfect Bayesian equilibrium in pure strategies where along the equilibrium path: (i) the first*

buyer bids p regardless of her signal; (ii) if $p \geq a_0$ the seller sells the object to the buyer and if $p < a_0$ the seller stops the search without selling the object.

The reason that the seller does not even attempt to solicit offers from other new buyers is that he knows that they will decline to increase the existing bid. If they do, they will be believed to have observed high signals and play will revert to the backwards induction PBE described in Proposition 1, where no buyer except for the last one can ever capture a positive rent. The deviating buyer is surely not the last buyer and the equilibrium is sustained. Note that along the equilibrium path, if trade occurs, the first buyer captures a positive rent.⁹

3.4. Discussion and Evaluation

The equilibria described in Sections 3.1–3.3 span a wide range of equilibrium outcomes. In the first PBE described in Proposition 1, herding and informational cascades never occur and the seller extracts almost the entire buyers' surplus. This is the only PBE among the three considered where the outcome is always efficient. In the second PBE that is described in Proposition 3, observing a high signal followed by k successive low signals (sufficiently early) triggers herding and generates an informational cascade. The seller succeeds in capturing some surplus from the buyers, but as the next proposition shows, not as much as in the first PBE. This is not because the buyers succeed in securing some rent for themselves, but rather because the latter PBE is less efficient. Finally, in the PBE described in Proposition 4, herding and an informational cascade occur immediately and the buyers capture all the rent. The seller's expected payoff is the lowest in this case.

PROPOSITION 5. *The expected payoff to the seller in the PBE of Proposition 1 is larger than or equal to his expected payoff in the PBE of Proposition 3.*

We have described only one equilibrium in Sections 3.1 and 3.2, respectively, but by changing the equilibrium strategies outside the equilibrium path, many similar equilibria can be constructed. Many equilibria similar to those described in Section 3.3 exist. In addition, other equilibria can be created by “combining” the equilibria described in Sections 3.1–3.3. This multiplicity of equilibria has two reasons: in general, price competition among buyers implies that the buyer who wins the object pays exactly the expected value of the object. Therefore, each buyer is indifferent between winning the object and paying the associated price, and not winning the

⁹ The situation described here is similar to the one described by [14] where in a complete information model, a non competitive equilibrium arises in a frictionless environment.

object but not paying. The other reason is that in order to avoid the winner's curse each buyer conditions her bid on the information implied by her winning the object. However, the event of winning the object depends on other buyers' strategies, and the constraints on the strategy profiles implied by the winner's curse do not lead to a unique outcome. There are many mutually consistent ways in which buyers can incorporate their concerns about the winner's curse into their bids.

For reasons similar to those that make the backwards induction outcome less plausible in the chain-store paradox [15] and in the Centipede game [13], we believe that the backwards induction equilibrium described in Proposition 1 is more plausible when the number of potential buyers is small, but the herding equilibrium described in Proposition 3 is more plausible when this number is large. In addition, when the number of potential buyers is large, the PBE of Proposition 3 has the advantage that the first $N - 2k - 2$ buyers need not know the total number of potential buyers N . For these "early" buyers, it is sufficient to know that the seller can approach at least $2k + 1$ additional new buyers.

The uninformative equilibria described in Proposition 4 are unappealing because buyers rely on weakly dominated strategies. Specifically, if the seller continues to solicit offers from new buyers, a buyer who deviates and bids p^N may end buying the object for this price, capturing a rent even if her signal is low.

Finally, note that of the three equilibria described, only the backwards induction PBE described in Proposition 1 is robust to a modification of the game into one where the seller is allowed to make a take-it-or-leave-it offer to the buyers after herding occurs. Thus, which equilibrium emerges as the likely outcome in this more general game depends on the ability of the seller to commit to such take-it-or-leave-it offers. The problem is that after such an offer is accepted by one of the buyers, herding and the informational cascade break down and the seller could only benefit from approaching new buyers who may bid the price even higher if they observe high signals, thus exposing the buyer who accepted the offer to the winner's curse.

4. COMPARISON TO AUCTIONS

In this section we address the following question. The seller can either solicit bids from buyers sequentially, as described here, or conduct an auction. Which procedure is better for the seller?

The environment we consider is a common values environment where buyers' signals are affiliated. Milgrom and Weber [10] have shown that in such environments the English auction generates higher expected revenues

for the seller than the other commonly used auctions such as the Dutch and the sealed bid first and second price auctions. We therefore compare the procedure of sequentially soliciting offers from buyers to an English auction. We obtain sharp results. We show that the sequential procedure, even under the herding PBE of Proposition 3, is superior to an English auction; depending on the seller's reservation value, the seller's expected payoff under the sequential procedure is either larger or equal to that generated by the English auction.

The English auction has many variants. For our purposes it is most convenient to focus on the variant that is referred to by Milgrom and Weber [10] as the Japanese version of the English auction. It is idealized as follows. Before the auction begins, the bidders are given the opportunity to inspect the object and to observe their private signals. Upon observing their signals, the bidders choose whether to be active at the start price p_0 . As the auctioneer raises the price, bidders drop out one by one. No bidder who has dropped out can become active again. After any bidder quits, all remaining active bidders know the price at which she quit. The auction ends as soon as no more than one bidder remains active. The remaining bidder gets the object for the prevailing price. If several bidders dropped out simultaneously, ending the auction, one of these bidders is chosen randomly and gets the object at the price at which she quit.

Define

$$p^{1, N-1} = E[v(q) \mid \text{out of } N \text{ signals, } 1 \text{ is high and } N-1 \text{ are low}].$$

The following proposition describes the outcome under the symmetric equilibrium that is characterized by Milgrom and Weber [10].

PROPOSITION 6. *The equilibrium outcome under the English auction is as follows:*

(i) *If $p_0 \leq p^N$ the sale price is equal to the expected valuation of the object conditional on all the buyers' information except when there is only one high signal, in which case it is equal to p^N .*

(ii) *If $p^N < p_0 \leq p^{1, N-1}$ the sale price is equal to the expected valuation of the object conditional on all the buyers' information except when there is only one high signal, in which case it is equal to p_0 . If all signals are low, the seller retains the object.*

(iii) *If $p^{1, N-1} < p_0$ the seller retains the object regardless of the realization of buyers' signals. No bid is made.*

In case (iii) above, the winner's curse prevents the buyers from bidding for the object and the object is not sold. This is because a buyer who observes a high signal knows that if she bids the price up, she may either suffer

the winner's curse if she is the only buyer to have observed a high signal, or, if some other buyers have observed high signals, competition will drive her rent to zero. As a consequence, she refuses to bid up the price, no information is revealed, and the seller retains the object.

We distinguish among three cases that differ by the seller's degree of sophistication and his ability to commit not to sell the object. First, we consider the case of a naive seller who sets an initial price $p_0 = a_0$. Next, we consider the case of a sophisticated seller that lacks commitment power. Finally, we consider the case of a sophisticated seller that has commitment ability.

Proposition 6 illustrates why it is naive for the seller to set the initial price p_0 equal to the reservation price a_0 . If $a_0 > p^{1, N-1}$, the seller retains the object in spite of the fact that had he started with a lower initial price, he could have sold the object and obtained a positive payoff. Proposition 6 also implies that a sophisticated seller may always set a start price that is either $p_0 = p^N$ or $p_0 = p^{1, N-1}$. It can be shown that when $a_0 < p^N$ but sufficiently close to p^N , $p_0 = p^{1, N-1}$ is the initial price that maximizes the seller's expected payoff. This is because the seller is better off trading the higher revenue associated with $p^{1, N-1}$ (vs p^N) against the smaller probability of selling (the seller retains the good if all buyers observe low signals). Note, however, that to do this successfully, the seller must be able to commit not to sell the object to some buyer for the price p^N after he failed to sell it when he set the initial price $p^{1, N-1}$ and the buyers have all observed low signals. If the seller cannot commit, that is, if the buyers believe that if all buyers observe low signals, the seller will sell the object to one of them for the price p^N , then the buyers have an incentive to pretend to have observed low signals. A sophisticated seller who lacks commitment power and understands the buyers' incentives has no choice but to set an initial price $p_0 = p^N$ when $a_0 < p^N$.

We compare the expected revenue to the seller under the sequential procedure, specifically under the PBE described in Proposition 3, to the seller's expected revenue under the English auction for each seller type and show that the former generates a higher expected payoff for the seller. Proposition 5 implies that the conclusion holds a fortiori for the PBE described in Proposition 1.

PROPOSITION 7. *The expected payoff to a naive seller, and to a sophisticated seller who lacks commitment power, under the sequential procedure with the PBE of Proposition 3 is higher than or equal to his expected payoff under an English auction.*

To prove Proposition 7, we distinguish among three cases. (1) If $a_0 \in [p^N, p^{1, N-1}]$, a sophisticated seller sets $p_0 = p^{1, N-1}$, and the English

auction, the PBE described in Proposition 1 and the PBE described in Proposition 3 are all equivalent in terms of expected revenues they generate for the seller. By Proposition 6, a naive seller who sets $p_0 = a_0$ obtains a lower or equal expected revenue. (2) If $a_0 < p^N$, in the English auction the seller loses rent to the winning buyer whenever there is exactly one buyer who has observed a high signal, whereas in the PBE under the sequential procedure described in Propositions 1 and 3, the seller loses rent only if the *last* buyer is the only one to have observed a high signal. Finally, (3) if $a_0 > p^{1, N-1}$, the seller either sets an initial price $p_0 > p^{1, N-1}$, in which case, by Proposition 6, buyers refuse to bid and the seller retains the object, or the seller sets an initial price $p_0 = p^{1, N-1}$ and with a positive probability will end up selling the object at a price that is below his reservation value a_0 .¹⁰

To fairly compare the performance of the sequential procedure to that of the English auction, we must assume that the seller has the power to commit under the sequential procedure if we assume that he has this ability under the English auction. We therefore assume that the seller has the power to insist on the price $p^{1, N-1}$ even when $a_0 < p^N$ and that he can do so under both procedures (sequential or auction). In the sequential procedure, this would give rise to PBEs that are similar to those described in Propositions 1 and 3. We name them as the *modified* PBE of Proposition 1 or 3, respectively.

PROPOSITION 8. *The expected payoff to a sophisticated seller who has commitment power under the sequential procedure with the modified PBE of Proposition 3 is higher or equal to his expected payoff under an English auction.*

The intuition for Proposition 8 is similar to that of Proposition 7 except that the interval of reservation values for which the English auction generates the same expected revenue to the seller as the PBE described in Proposition 3 is now $[\alpha, p^{1, N-1}]$ instead of $[p^N, p^{1, N-1}]$ for some $\alpha < p^N$. The reason for this is that, as explained above, for a sophisticated seller with commitment power the optimal initial price is $p^{1, N-1}$ whenever a_0 is smaller than but close to p^N , whereas for a seller without commitment

¹⁰ Lopomo [6] considers a variant of the English auction where the seller sets a low initial price, but may set a reserve price after all but the last bidder have dropped out from the auction. That is, the seller may make the last remaining bidder a take-it-or-leave-it offer for the object. This variant of the English auction performs equally well as the sequential procedure under the PBEs described in Propositions 1 and 3 in the case where $p^N \leq a_0 \leq p^{1, N-1}$. It performs equally well as the PBE described in Proposition 1 in the case where $a_0 > p^{1, N-1}$ because it eliminates the buyers' concerns about the winner's curse. But in the case where $a_0 < p^N$, the seller still loses rent to the buyers under this variant of the English auction more frequently than under the sequential procedure under both the PBEs described in Propositions 1 and 3.

power and a reservation value below p^N , any initial price that is smaller or equal to p^N is an optimal initial price. When either $a_0 < \alpha$ or $a_0 > p^{1, N-1}$, because of the reasons given above for a seller that lacks commitment power, the English auction generates a lower expected revenue to the seller than the PBE described in Proposition 3.

Finally, welfare comparisons between the sequential procedure and the English auction are complicated by the fact that the PBE of Proposition 3 may give rise to inefficiency. However, we can make the following straightforward observation. The PBE of Proposition 1 is always efficient whereas the English auction is efficient only if $a_0 \leq p^{1, N-1}$ and the seller has no commitment power.

APPENDIX: PROOFS

Proof of Lemma 1. By Bayes' law,

$$P_H(n_L, n_H) = \frac{\Pr(s_L | q_H)^{n_L} \Pr(s_H | q_H)^{n_H} \Pr(q_H)}{\left[\Pr(s_L | q_H)^{n_L} \Pr(s_H | q_H)^{n_H} \Pr(q_H) + \Pr(s_L | q_L)^{n_L} \Pr(s_H | q_L)^{n_H} \Pr(q_L) \right]}.$$

Because $0 < \Pr(s_L | q_H) < \Pr(s_L | q_L) < 1$, $P_H(n_L, n_H)$ is strictly decreasing in n_L and because $1 > \Pr(s_H | q_H) > \Pr(s_H | q_L) > 0$, $P_H(n_L, n_H)$ is strictly increasing in n_H .

It is straightforward to verify that $P_H(n_L + i, n_H + j)$ equals

$$\frac{\Pr(s_H | q_H)^j \Pr(s_L | q_H)^i P_H(n_L, n_H)}{\left[\Pr(s_H | q_H)^j \Pr(s_L | q_H)^i P_H(n_L, n_H) + \Pr(s_H | q_L)^j \Pr(s_L | q_L)^i (1 - P_H(n_L, n_H)) \right]}$$

for every i, j, n_L and n_H . Thus, $P_H(n_L + k, n_H + 1) < P_H(n_L, n_H)$ if and only if

$$\frac{\Pr(s_H | q_H)}{\Pr(s_H | q_L)} \leq \frac{\Pr(s_L | q_L)^k}{\Pr(s_L | q_H)^k}$$

and $P_H(n_L, n_H) < P_H(n_L + k - 1, n_H + 1)$ if and only if

$$\frac{\Pr(s_L | q_L)^{k-1}}{\Pr(s_L | q_H)^{k-1}} < \frac{\Pr(s_H | q_H)^k}{\Pr(s_H | q_L)}.$$

By our assumptions $\Pr(s_H | q_H)/\Pr(s_H | q_L)$ and $\Pr(s_L | q_L)/\Pr(s_L | q_H)$ are strictly larger than 1. As a consequence, there exists a unique k such that

$$\frac{\Pr(s_L | q_L)^{k-1}}{\Pr(s_L | q_H)^{k-1}} < \frac{\Pr(s_H | q_H)}{\Pr(s_H | q_L)} \leq \frac{\Pr(s_L | q_L)^k}{\Pr(s_L | q_H)^k}. \quad \blacksquare$$

Proof of Lemma 2. Given any profile of pure strategies $(A^p, A^b)_{b \in \mathcal{B}}$ and $(\Gamma^p, \Gamma^b)_{b \in \mathcal{B}}$ let $P_H(h_t)$ denote the probability that the object is of high quality given the signals that can be inferred from the history h_t . $V(h_t)$ can be written as

$$\begin{aligned} V(h_t) &= P_H(h_t) v(q_H) + (1 - P_H(h_t)) v(q_L) \\ &= P_H(h_t) [v(q_H) - v(q_L)] + v(q_L). \end{aligned}$$

Therefore, it is sufficient to show that $P_H(h_t)$ is a martingale, i.e., that $E[P_H(h_{t+1}) | P_H(h_t)] = P_H(h_t)$. Given the PBE $(A^p, A^b)_{b \in \mathcal{B}}$ and $(\Gamma^p, \Gamma^b)_{b \in \mathcal{B}}$ and a history h_t , the action of buyer $b(t+1)$ at $t+1$ depends on her random signal and thus is a random variable with realizations $a_{t+1}^{b(t+1)} \in \mathcal{A}_{t+1}^{b(t+1)}$ where the set $\mathcal{A}_{t+1}^{b(t+1)}$ of admissible actions depends on $p(h_t)$ and on whether the buyer is old or new. Given h_t , let $dF(a_{t+1}^{b(t+1)} | h_t)$ denote the density or probability, respectively, of $a_{t+1}^{b(t+1)}$; and let $dF(a_{t+1}^{b(t+1)} | h_t, q_H)$ denote the density of probability, of $a_{t+1}^{b(t+1)}$ conditional on quality q_H . Applying the definition $h_{t+1} = (h_t, a_{t+1}^{b(t+1)})$ and Bayesian updating we get

$$\begin{aligned} E[P_H(h_{t+1}) | h_t] &= E[P_H(h_t, a_{t+1}^{b(t+1)}) | h_t] \\ &= \int_{\mathcal{A}_{t+1}^{b(t+1)}} \Pr(q = q_H | h_t, a_{t+1}^{b(t+1)}) dF(a_{t+1}^{b(t+1)} | h_t) \\ &= \int_{\mathcal{A}_{t+1}^{b(t+1)}} \frac{dF(q = q_H, A_{t+1}^{b(t+1)} = a_{t+1}^{b(t+1)} | h_t)}{dF(a_{t+1}^{b(t+1)} | h_t)} \\ &\quad \times dF(a_{t+1}^{b(t+1)} | h_t) \\ &= \int_{\mathcal{A}_{t+1}^{b(t+1)}} \frac{dF(a_{t+1}^{b(t+1)} | h_t, q_H) P_H(h_t)}{dF(a_{t+1}^{b(t+1)} | h_t)} dF(a_{t+1}^{b(t+1)} | h_t) \\ &= P_H(h_t) \int_{\mathcal{A}_{t+1}^{b(t+1)}} dF(a_{t+1}^{b(t+1)} | h_t, q_H) \\ &= P_H(h_t). \quad \blacksquare \end{aligned}$$

Proof of Proposition 1. Let b_N denote the last new buyer that is approached by the seller. For any $t \geq 2$, let

$$\begin{aligned} \hat{W}_N(h_{t-1}) &= E \left[v(q) \left| \begin{array}{l} (A^b)_{b \in \mathcal{B}}, (\Gamma^b)_{b \in \mathcal{B}}, h_{t-1}, \text{ and all the buyers who have} \\ \text{not yet revealed their signals observe low signals} \end{array} \right. \right], \end{aligned}$$

which given the history is the minimal expected value of the object.

The following profile of strategies constitutes a symmetric pure strategy PBE.

If buyer $b \in \mathcal{B}$ is a new buyer different from b_N , then for every history $h_{t-1} \in \mathcal{H}$,

$$A^b(h_{t-1}, s_L) = -\infty$$

and

$$A^b(h_{t-1}, s_H) = \begin{cases} W_N(h_{t-1}, s_H) & \text{if } W_N(h_{t-1}, s_H) \geq p(h_{t-1}) \\ -\infty & \text{otherwise} \end{cases}.$$

If $a_0 \geq p^N$, then the last buyer, b_N , bids like other buyers; if $a_0 < p^N$, the last buyer employs the same bidding strategy as other buyers except in the case where all preceding buyers have revealed themselves to have observed low signals in which case the last buyer bids p^N regardless of her signal.

If buyer $b' \in \mathcal{B}$ is an old buyer, then for every history $h_{t-1} \in \mathcal{H}$ and private signal $s^{b'} \in \mathcal{S}$,

$$A^{b'}(h_{t-1}, s^{b'}) = \hat{W}_N(h_{t-1}).$$

The seller's strategy is

$$A^s(h_t) = \begin{cases} b & \text{if the number of new buyers approached is smaller than } N \\ b' & \text{if the number of new buyers approached is equal to } N, p^N < p(h_t) < V(h_t), \text{ and } b' \text{ has revealed that she had observed a high signal} \\ b_N & \text{if the number of new buyers approached is equal to } N \text{ and either } p(h_t) < p^N \text{ or } p(h_t) = p^N \text{ but the offer of buyer } b_N \text{ is less than } p^N \\ \text{accept } p(h_t) & \text{if the number of new buyers approached is equal to } N \text{ and either } p(h_t) \geq V(h_t) \geq a_0 \text{ and the buyer who has offered } p(h_t) \text{ has revealed that she had observed a high signal or } p(h_t) = p^N > a_0 \text{ and } p^N \text{ has been offered by } b_N \\ \text{"stop"} & \text{if the number of new buyers approached is equal to } N \text{ and either } p^N < V(h_t) \leq p(h_t) < a_0 \text{ or } p(h_t) = p^N \leq a_0 \end{cases}$$

for every $h_t \in \mathcal{H}$ where b is a new buyer and b' is an old buyer. If more than one buyer has offered to pay $p(h_t) > p^N$ for the object and $p(h_t) \geq V(h_t) \geq a_0$, the seller selects the buyer who first made this bid provided she has revealed that she has observed a high signal.

If $p(h_t) = p^N > a_0$, the seller selects buyer b_N .

Finally, each player's beliefs are consistent with other players' strategies. If a new buyer b chooses an action a_t^b that cannot be generated by the buyers' equilibrium strategies, the other players infer that she has observed a low signal if she declines to bid, and that she has observed a high signal otherwise. The players' beliefs about old buyers signals are independent of their actions.

To see that this profile of strategies and beliefs forms a PBE, note that new buyers cannot benefit from pretending to observe low signals when they observe high signals because they will not have an opportunity to buy the object in this case. When they reveal themselves to observe high signals, they cannot benefit from not bidding the price up to its valuation conditional on public information and the winner's curse because if they do, the seller will approach an old buyer who will.

The last buyer b_N captures a rent when $a_0 < p_N$ and all buyers before her have revealed low signals. As explained in the body of the paper, the seller cannot extract this rent from the last buyer. It is straightforward to observe that the old buyers' (who are not approached by the seller along the equilibrium path) and the seller's strategies are best responses as well. ■

Proof of Proposition 2. For every the number of buyers N , the transaction price under the PBE described in Proposition 1 is equal to $V(h_N)$ with probability $1 - \Pr(S^b = s_L | q)^{N-1}$ and to p^N with probability $\Pr(S^b = s_L | q)^{N-1}$. Hence the transaction price converges to $V(h_N)$ as $N \rightarrow \infty$. By Lemma 2, $V(h_N)$ is a martingale, and therefore, by the martingale convergence theorem (see, e.g., [16, p. 508]) converges to a random variable V with probability 1. We show that $V = v(q)$ where q is the true quality of the object with probability 1.

Conditional on the true quality of the object, the signals $S^{b(1)}, S^{b(2)}, \dots$ are i.i.d. Let λ_N denote the proportion of high signals from among the signals $\{S^{b(1)}, \dots, S^{b(N)}\}$. By the strong law of large numbers, conditional on the true quality q , $\lambda_N \rightarrow \Pr(s_H | q)$ with probability 1. Because $\Pr(s_H | q_H) > \Pr(s_H | q_L)$, observing the signals identifies the true quality of the object with probability 1 at the limit. The convergence of λ_N therefore implies that $V(h_N) \xrightarrow{N \rightarrow \infty} v(q)$ where q is the true quality of the object with probability 1. ■

Proof of Proposition 3. It is helpful to first consider the case where the number of buyers is countably infinite. For every $t \in \{1, 2, \dots\}$, let

$$W_k(h_{t-1}, s_H) = E \left[v(q) \left| \begin{array}{l} (A^b)_{b \in \mathcal{B}}, (\Gamma^b)_{b \in \mathcal{B}}, h_{t-1}, s^{b(t)} = s_H, \\ \text{and } k \text{ additional low signals} \end{array} \right. \right]$$

denote a new buyer's valuation of the object at time t given what she can learn from past history, the fact that she observed a high signal, and the assumption that k additional new buyers will observe low signals. For every possible history h_t , let $m(h_t)$ denote the number of inferred low signals that were observed by new buyers after the latest inferred high signal. For every $t \in \{1, 2, \dots\}$, let

$$W_k(h_{t-1}, s_L) = E \left[v(q) \left| \begin{array}{l} (A^b)_{b \in \mathcal{B}}, (\Gamma^b)_{b \in \mathcal{B}}, h_{t-1}, s^{b(t)} = s_L, \text{ and} \\ \max\{k - m(h_{t-1}) - 1, 0\} \text{ additional low signals} \end{array} \right. \right]$$

denote a new buyer's valuation of the object at time t given what she can learn from past history, the fact that she observed a low signal, and the assumption that an additional number of $\max\{k - m(h_{t-1}) - 1, 0\}$ low signals will be observed. The following proposition describes a particularly compelling equilibrium for the case where the number of buyers is infinite.

PROPOSITION 9. *When the number of buyers is countably infinite, the following strategies constitute a symmetric perfect Bayesian equilibrium in pure strategies.*

If buyer $b \in \mathcal{B}$ is a new buyer, then for every history $h_{t-1} \in \mathcal{H}$, and private signal $s^b \in \mathcal{S}$,

$$A^b(h_{t-1}, s^b) = \begin{cases} W_k(h_{t-1}, s^b) & \text{if } W_k(h_{t-1}, s^b) > p(h_{t-1}) \\ -\infty & \text{otherwise.} \end{cases}$$

If buyer $b' \in \mathcal{B}$ is an old buyer, then for every history $h_{t-1} \in \mathcal{H}$ and private signal $s^{b'} \in \mathcal{S}$,

$$A^{b'}(h_{t-1}, s^{b'}) = \begin{cases} W_k(h_{t-1}, s_L) & \text{if } W_k(h_{t-1}, s_L) > p(h_{t-1}) \\ (h_{t-1}) & \text{if } W_k(h_{t-1}, s_H) \leq p(h_{t-1}) < V(h_{t-1}) \\ a_\tau^{b'} & \text{otherwise,} \end{cases}$$

where $a_\tau^{b'}$ is the previous action of buyer b' .

The seller's strategy is

$$A^p(h_t) = \begin{cases} b & \text{if } p(h_t) < V(h_t) \text{ and } W(h_t, s_H) > p(h_t) \\ b' \neq b(t) & \text{if } W(h_t, s_H) \leq p(h_t) < V(h_t) \text{ and } b' \text{ has} \\ & \text{revealed that she had observed a high signal} \\ \text{accept } p(h_t) & \text{if } p(h_t) \geq V(h_t) \geq a_0 \text{ and the buyer who has} \\ & \text{offered } p(h_t) \text{ has revealed that she had} \\ & \text{observed a high signal} \\ \text{stop} & \text{if } V(h_t) \leq p(h_t) < a_0 \end{cases}$$

for every $h_t \in \mathcal{H}$ where (i) b is a new buyer, (ii) b' is an old buyer, (iii) “accept $p(h_t)$ ” implies that the seller sells the object for the price $p(h_t)$ to the buyer who first submitted this bid provided she has revealed that she has observed a high signal, and (iv) “stop” implies that the seller ends the search for a buyer without selling the object.

Finally, player's beliefs are consistent with other players' strategies. If a new buyer b chooses an action a_t^b that cannot be generated by the buyers' equilibrium strategies, other players infer that she has observed a low signal if she declines to bid, and that she has observed a high signal otherwise. The players' beliefs about old buyers signals are independent of their actions.

Proof. Along the equilibrium path, only buyers who have observed a high signal bid for the object. The seller only sells to buyers who have revealed high signals. Therefore, buyers cannot benefit from pretending to observe a low signal when their signal is high. A buyer who has revealed herself to observe a high signal cannot benefit from bidding below the object's expected valuation conditional on public information and the winner's curse because the seller will approach an old buyer who will bid the price up. The old buyers' and the seller's strategies are best responses as well. ■

We now consider the case where the number of buyers N is finite and larger or equal to $2k + 3$. There exists a PBE where up to buyer $N - (2k + 1)$ the players employ the strategies described in the PBE of Proposition 9. If herding has not occurred by then, the players revert to the backwards induction equilibrium strategies and the seller approaches all N buyers. Specifically, the following strategies and beliefs constitute a perfect Bayesian equilibrium in pure strategies.

- As long as $N - (2k + 1)$ or less buyers have revealed their signals, the strategies of the buyers and the seller prescribe the same actions as the strategies specified in Proposition 9.

- The strategies for the last $2k + 1$ new buyers are as follows. A new buyer $b(t)$ that is approached when all but $n \in \{2, \dots, 2k + 1\}$ buyers have revealed their signals bids $W_N(h_{t-1}, s_H)$ if she has observed a high signal and $W_N(h_{t-1}, s_H) > p(h_{t-1})$, and declines to bid otherwise. The strategy of the last new buyer is identical to the strategy of the last buyer in the PBE of Proposition 1.

- If an old buyer is re-approached by the seller at some t after exactly $N - 2k - 1$ buyers have revealed their signals, she bids as old buyer in

Proposition 9. If an old buyer is re-approached by the seller at some t after $N - 2k$ or more buyers have revealed their signals, she bids $\hat{W}_N(h_{t-1})$ (defined in the proof of Proposition 1).

- After $N - 2k$ or more buyers have revealed their signals the seller's strategy prescribes the same actions as the seller's strategy in the PBE of Proposition 1.

- If at t exactly $N - (2k + 1)$ buyers have revealed their signals (and the seller has not yet sold the object) the seller's strategy prescribes the following actions. If $p(h_{t-1}) < W_N(h_{t-1}, s_H)$, he solicits an offer from a new buyer (and continues to do so until he has approached all buyers). If $p(h_{t-1}) \geq \max[V(h_{t-1}), a_0]$, he sells the object for the price $p(h_{t-1})$. If $a_0 > p(h_{t-1}) \geq V(h_{t-1})$ or if $a_0 > V(h_{t-1}) > p(h_{t-1}) \geq W_N(h_{t-1}, s_H)$, he ends the search without selling the object. If $a_0 \leq V(h_{t-1})$ and $V(h_{t-1}) > p(h_{t-1}) \geq W_N(h_{t-1}, s_H)$, he re-approaches the buyer who, among those who have revealed a high signal (because of $p(h_{t-1}) \geq W_N(h_{t-1}, s_H)$ there is at least one), has made the highest offer and demands a bid $V(h_{t-1})$. If this buyer declines to bid $V(h_{t-1})$, the seller approaches any other buyer who has revealed her signal and sells her the object if she bids $V(h_{t-1})$. If this other buyer bids less than $V(h_{t-1})$, the seller continues to approach buyers who have revealed their signal until one offers $V(h_{t-1})$, in which case she sells the object.¹¹

- Finally, each player's beliefs are consistent with other players' strategies. If a new buyer b chooses an action a_t^b that cannot be generated by the buyers' equilibrium strategies, the other players infer that she has observed a low signal if she declines to bid and that she has observed a high signal otherwise. The players' beliefs about old buyers signals are independent of their actions.

The game ends when either herding has occurred or the seller has approached all N buyers. Herding will occur if and only if at least one among the first $N - 2k - 1$ buyers observes a high signal. This is because herding occurs if either buyer $N - 2k - 1$ observes a high signal or buyer $N - 2k - 1$ observes a low signal but some buyer before her has observed a high signal.

We show that the specified strategies are best replies. For the first $N - 3k - 1$ new buyers the arguments of the proof of Proposition 9 carry over. Since buyer $N - 2k$ conditions her bid on the history, on her own signal and on $2k$ additional low signals, these arguments apply analogously

¹¹ Note that the buyers' strategies are such that in the case where $V(h_{t-1}) > p(h_{t-1}) \geq W_N(h_{t-1}, s_H)$, any buyer who has revealed her signal is willing to bid $V(h_{t-1})$. Thus, along the equilibrium path the seller will never re-approach a buyer that has not revealed a high signal.

for the new buyers $N - 3k$ to $N - 2k - 1$.¹² For the last $2k + 1$ new buyers the arguments of the proof of Proposition 1 hold analogously. Similarly, for old buyers the arguments of the proof of Proposition 9 (if less than $N - 2k$ buyers have revealed their signals) or Proposition 1 (if $N - 2k$ or more buyers have revealed their signals), respectively, apply analogously. Finally, the seller's strategy is also optimal. Whenever there is herding at some t an offer $V(h_{t-1})$ is the best he can get, and by re-approaching old buyers (if necessary) he obtains a bid equal to $V(h_{t-1})$. The optimality of the seller's strategy in the case where herding does not occur follows from the same arguments as in the proof of Proposition 1. ■

Proof of Proposition 4. The strategies and beliefs that support this PBE are as follows. The first buyer bids p . If the first buyer bids p , all other buyers decline to make a bid. If the first buyer declines to make a bid or bids differently from p , (i) the seller and all the buyers (including the first buyer when re-approached) revert to playing the PBE of Proposition 1; (ii) the seller and all the other buyers believe that the first buyer has observed a high signal. If after the first buyer has made the bid p any other buyer, when approached by the seller, makes a bid, (i) the seller and all the buyers revert to playing the PBE of Proposition 1; (ii) the seller and all the buyers except the deviating one believe that the deviating buyer has observed a high signal. Since conditional on one (or more) high signals no buyer gets a positive rent in the PBE of Proposition 1, no buyer has an incentive to make a higher bid than p and the first buyer's strategy is optimal. Given the buyers' strategies, it is optimal for the seller either to accept p when $p \geq a_0$ or to stop the search without selling the object when $p < a_0$. ■

Proof of Proposition 5. The proof relies on the fact that the seller's payoff, $\max\{V(h_t) - a_0, 0\}$ is a submartingale. We prove a stronger result. If $a_0 \leq p^{1, N-1}$ or $a_0 \geq E[v(q) | N \text{ high signals}]$ then the seller's expected payoff under the PBEs of Proposition 1 and 3 is identical; otherwise, it is strictly higher under the PBE of Proposition 1. If $a_0 \geq E[v(q) | N \text{ high signals}]$, the seller obtains a payoff of zero under both equilibria. Suppose that $a_0 < E[v(q) | N \text{ high signals}]$. Denote the maximum bid the seller obtains by p_{\max} . The seller's payoff is $\max\{p_{\max} - a_0, 0\}$. The expected maximum bid, p_{\max} , is identical under the PBEs of Propositions 1 and 3. Because $V(h_t)$ is a martingale (Lemma 2), $E[V(h_t)]$ is the same in both equilibria. The maximum bid in both equilibria is equal to $V(h_t)$ except in

¹² Note that conditioning on (only) k additional low signals is a best reply only if the new buyer who is approached after k additional low signals have been realized conditions again on at least k low signals, in addition to history and her own signal. This implies that for a new buyer conditioning on (only) k additional low signals can be a best reply only if there are at least $2k + 1$ additional new buyers left.

the case where the first $N - 1$ buyers have observed a low signal. But this case, $p_{\max} = p^N$ in both equilibria. Therefore, if $a_0 \leq p_{\max}$, and thus whenever $a_0 \leq p^{1, N-1}$ the expected payoff to the seller is the same in both equilibria. Consider now the case where $a_0 > p^{1, N-1}$. Assume herding has occurred at some price $p = E[v(q) | h_t]$ at $t + 1$ after a history h_t that reveals exactly one high signal. This implies $p > p^{1, N-1}$ because less than N buyers have revealed their signals at $t + 1$. The seller's payoff under the PBE of Proposition 3 is $\max\{p - a_0, 0\}$. In the PBE of Proposition 1 the expected payoff conditional on h_t is $E[\max\{E[v(q) | S^1, \dots, S^N] - a_0, 0\} | h_t]$. Note that the event $E[(v(q) | S^1, \dots, S^N) | h_t] = p^{1, N-1}$ has positive probability. Because of this, the fact that $a_0 > p^{1, N-1}$ and the martingale property (Lemma 2) it follows that

$$\begin{aligned} E[\max\{E[v(q) | S^1, \dots, S^N] - a_0, 0\} | h_t] \\ &> \max\{E[E[v(q) | S^1, \dots, S^N] | h_t] - a_0, 0\} \\ &= \max\{p - a_0, 0\}. \end{aligned}$$

If more than one high signal can be inferred from the history h_t , the strict inequality changes to a weak inequality. The proposition follows. ■

Proof of Proposition 6. A pure strategy for buyer b specifies, as a function of the private signal and the observable activity of the other buyers, (i) whether or not to be active at the initial price p_0 , and (ii) for any price $p \geq p_0$, whether or not to remain active at p , provided buyer b has been active up to p . The argument given in [10] can be adapted to show that the following profile of strategies constitutes a symmetric equilibrium. If $p_0 \leq p^N$ every bidder is active at p_0 . All the bidders who have observed low signals drop out at p_0 , and all the bidders who have observed high signals remain active. Thus, "immediately after" the price has been raised above p_0 all signals are revealed. If no buyer has observed a high signal, the auction ends at p_0 and the object is sold to a randomly selected buyer for the price p_0 . If exactly one buyer has observed a high signal, the auction ends at p_0 and the object is sold to this buyer for the price p_0 . If two or more buyers have observed a high signal, they remain active until the price reaches the value of the object conditional on all the bidders' signals and the object is sold to one of them (randomly selected) at this price. If $p_0 \in (p^N, p^{1, N-1}]$, a bidder chooses to be active at p_0 if and only if she has observed a high signal. If there is only one such bidder, she gets the object for the price p_0 ; if there are two or more, the auction proceeds like in the case where $p_0 \leq p^N$. Finally, if $p_0 > p^{1, N-1}$, because of the winner's curse no bidder is willing to become active at the initial price $p_0 > p^{1, N-1}$. This follows from the fact that whenever the expected value conditional on the observed signals is at least p_0 , there must be at least two bidders who have observed

a high signal and thus they will bid up the price to its expected value conditional on all the signals. Since there is a positive probability that a bidder who chooses to be active at p_0 is the only one who has observed a high signal, a bidder that chooses to be active at p_0 exposes herself to the winner's curse. ■

Proof of Proposition 7. We first prove the proposition for the case of the unsophisticated seller who sets an initial price $p_0 = a_0$ in the English auction. If $a_0 > p^{1, N-1}$, Proposition 6 implies that the seller retains the object under the English auction. Since the seller may sell the object for a higher price than a_0 under the PBE of Proposition 3, his expected payoff is higher or equal under the sequential procedure. Suppose that $a_0 \leq p^{1, N-1}$. It follows from the proof of Proposition 5 that in this case the expected payoff to the seller is identical under the PBEs of Propositions 1 and 3. We can therefore compare the expected payoff under the English auction with the one in the PBE of Proposition 1. If $a_0 = p^{1, N-1}$ the English auction and the PBE of Proposition 1 lead to an identical outcome. If $a_0 < p^{1, N-1}$ the outcome of the auction and of the equilibrium in the PBE of Proposition 1 are identical whenever at least two buyers have observed a high signal. But whenever any buyer who is not the last buyer b_N in the PBE of Proposition 1 is the only one to have observed a high signal, the seller's revenue in the auction is $\max\{p^N, a_0\}$ whereas it is $p^{1, N-1} > \max\{p^N, a_0\}$ in the PBE of Proposition 1. If buyer b_N is the only one to have observed a high signal, the seller's revenue in the PBE of Proposition 1 is greater than in the auction if $a_0 \in [p^N, p^{1, N-1})$ and identical if $a_0 < p^N$.

We now prove the proposition for the case of a sophisticated seller who lacks commitment power. We first derive the optimal initial price p_0 for the English auction. If

$$a_0 > E[v(q) \mid \text{out of } N \text{ signals at least one is high}]$$

the seller is better off not auctioning the object at all (or, equivalently, choosing $p_0 > p^{1, N-1}$); if $a_0 \leq E[v(q) \mid \text{out of } N \text{ signals at least one is high}]$, it is optimal for the seller to choose

$$p_0 = \begin{cases} p^N & \text{if } a_0 < p^N \\ p^{1, N-1} & \text{if } a_0 \geq p^N \end{cases}$$

The former statement follows from the fact that with $p_0 = p^{1, N-1}$ the expected maximum bid in the English auction conditional on selling is $E[v(q) \mid \text{out of } N \text{ signals at least one is high}]$. The latter statement follows immediately from Proposition 6 and the inability to commit, which implies

$p_0 \leq p^N$ whenever $a_0 < p^N$. If $a_0 \in [p^N, p^{1, N-1}]$, the English auction and the PBE of Proposition 1 generate identical payoffs for the seller for all signal realizations and thus (see proof of Proposition 5) the expected payoff is the same as in the PBE of Proposition 3. Next, consider the case where $a_0 < p^N$ and thus $p_0 = p^N$. Arguments analogous to those employed in the first half of the proof imply that the PBE of Proposition 3 generates a higher expected payoff for the seller than the English auction. Consider now the case where $p^{1, N-1} < a_0 \leq E[v(q) \mid \text{out of } N \text{ signals at least one is high}]$ and thus $p_0 = p^{1, N-1}$. Clearly, if in the PBE of Proposition 3 herding does not occur, the seller is never worse off than in the English auction, and he is strictly better off whenever the object's value conditional on all the buyers' signals is below a_0 , which happens with positive probability because $a_0 > p^{1, N-1}$. If herding has occurred at some price $p = E(v \mid h_t)$ at $t + 1$ after a history h_t (which necessarily reveals at least one high signal), the seller's payoff is $\max\{p - a_0, 0\}$. Because of the martingale property (Lemma 2), the expected payoff in the English auction conditional on the signals that are revealed by h_t is $p - a_0 \leq \max\{p - a_0, 0\}$ with a strict inequality whenever $p < a_0$. Taking the expectation over all possible outcomes of the PBE of Proposition 3 gives the proposition. ■

Proof of Proposition 8. If $a_0 \geq E[v(q) \mid N \text{ high signals}]$ the seller retains the object under both procedures and therefore obtains the same expected payoff. We therefore assume that $a_0 < E[v(q) \mid N \text{ high signals}]$. Note that if $a_0 \geq p^N$, the ability to commit is not an issue and the result therefore follows from the previous proposition. Suppose then that $a_0 < p^N$. In this case, the optimal initial price under the English auction is as follows. There exists a unique $\beta < p^N$ such that it is optimal for the seller to choose

$$p_0 = \begin{cases} p^N & \text{if } a_0 \leq \beta; \\ p^{1, N-1} & \text{if } a_0 > \beta, \end{cases}$$

β is determined by the condition that when $a_0 = \beta$ the seller's expected payoff is identical under the two initial prices, $p_0 = p^N$ and $p_0 = p^{1, N-1}$. This follows from the fact that for $p_0 = p^N$ the seller's expected payoff decreases with a_0 whereas for $p_0 = p^{1, N-1}$ it increases with a_0 .

If the seller has the power to commit, the PBEs described in Propositions 1 and 3 change in a way similar to the change in the auction's initial price p_0 . There exists a unique $\alpha \in (\beta, p^N)$ such that a seller with a reservation value $a_0 > \alpha$ does not accept a bid p^N . The number α is determined by the condition that when $a_0 = \alpha$ the seller's expected payoff is identical when he sells when the maximal bid is p^N and when he refuses to sell. The inequality $\beta < \alpha$ is due to the fact that in the English auction the seller loses informational rent to the winning buyer whenever *any of the N buyers* is the only buyer who has observed a high signal, whereas in the modified PBEs of

Proposition 1 and 3, the seller loses rent only if *one particular buyer* (the “last buyer” b_N) is the one who has observed a high signal. If $a_0 \notin (\alpha, p^N)$ the PBEs of Propositions 1 and 3 are not affected by the seller’s ability to commit.

Clearly, if $a_0 \notin (\beta, p^N)$ Proposition 7 applies, and thus the sequential procedure is equivalent for $a_0 \in [p^N, p^{1, N-1}]$ and better otherwise. Moreover, for $a_0 \in (\alpha, p^N)$ the English auction and the modified PBE of Proposition 1 have identical outcomes for all signal realizations and thus (using an argument similar to the one in the proof of Proposition 5) the seller’s expected payoff is identical under the English auction and the modified PBE of Proposition 3. If a_0 is smaller than α , the seller’s expected payoff rises in the modified PBE of Proposition 3 whereas in the English auction it decreases for every $a_0 \geq \beta$ and rises for $a_0 < \beta$. However, the argument that implied $\beta < \alpha$ implies that for $a_0 < \beta$ the seller’s expected payoff is strictly higher in the modified PBE of Proposition 3 than in the English auction. ■

REFERENCES

1. A. Banerjee, A simple model of herd behavior, *Quart. J. Econ.* **107** (1992), 797–817.
2. S. Bikhchandani, D. Hirshleifer, and I. Welch, A theory of fads, fashion, custom, and cultural change as in formational cascades, *J. Polit. Econ.* **100** (1992), 992–1026.
3. A. Devenow and I. Welch, Rational herding in financial economics, *Europ. Econ. Rev.* **40** (1996), 603–615.
4. D. Gale, What have we learned from social learning? *Europ. Econ. Rev.* **40** (1996), 617–628.
5. I. H. Lee, On the convergence of informational cascades, *J. Econ. Theory* **61** (1993), 395–411.
6. G. Lopomo, Optimality and robustness of the English auction, *J. Econ. Theory*, to appear.
7. J. McMillan and M. Rothschild, Search, in “Handbook of Game Theory with Economic Applications, II” (R. J. Aumann and S. Hart, Eds.), Elsevier, Amsterdam, 1994.
8. P. R. Milgrom, A convergence theorem for competitive bidding with differential information, *Econometrica* **47** (1979), 679–688.
9. P. R. Milgrom, Rational expectations, information acquisition, and competitive bidding, *Econometrica* **49** (1981), 921–943.
10. P. R. Milgrom and R. J. Weber, A theory of auctions and competitive bidding, *Econometrica* **50** (1982), 1089–1122.
11. M. J. Osborne and A. Rubinstein, “A Course in Game Theory,” MIT Press, Cambridge, MA, 1994.
12. W. Pesendorfer and J. M. Swinkels, The loser’s curse and information aggregation in common value auctions, *Econometrica* **65** (1997), 1247–1282.
13. R. Rosenthal, Games of perfect information, predatory pricing, and the chain-store paradox, *J. Econ. Theory* **25** (1981), 92–100.
14. A. Rubinstein and A. Wolinsky, Decentralized trading, strategic behavior and the Walrasian outcome, *Rev. Econ. Stud.* **57** (1990), 67–78.
15. R. Selten, The chain-store paradox, *Theory and Decision* **9** (1978), 127–159.

16. A. N. Shiryaev, "Probability," 2nd ed., Springer-Verlag, New York, 1996.
17. L. Smith and P. Sørensen, Pathological outcomes of observational learning, mimeo, MIT and Nuffield College, Oxford, (1997).
18. X. Vives, Social learning and rational expectations, *Europ. Econ. Rev.* **40** (1996), 589–601.
19. I. Welch, Sequential sales, learning, and Cascades, *J. Finance* **47** (1992), 695–732.
20. R. Wilson, A bidding model of perfect competition, *Rev. Econ. Stud.* **44** (1977), 511–517.