MANAGEMENT SCIENCE
DOI 10.1287/mnsc.1070.0691ec
pp. ec1–ec11

e-companion
ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—"Research Note: Strategic Bid-Shading and Sequential Auctioning with Learning from Past Prices" by Robert Zeithammer, *Management Science*, 10.1287/mnsc.1070.0691.

Online Supplement

EC.1. Relationship to eBay and Other Managerial Implications

The model studies the profit-maximization problem of a monopolist, who lives infinitely long in discrete time, can procure one unit of the good in each period at a non-trivial marginal cost, and faces overlapping generations of unit-demand bidders, who live for two periods each. These assumptions are loosely motivated by empirical regularities in the eBay marketplace. In the data about the eBay MP3-player market analyzed by Zeithammer (2006), at least 43 percent of the 22,603 unique bidders participated in two auctions or more (median three) in the data-window of three months, but 93 percent of the eventual winners only won one auction. Therefore, it seems that most eBay bidders on MP3-players have unit demand, and many remain active in future periods whenever they lose. Moreover, the bidders observed participating in multiple auctions tend to be focused on only one model of MP3 player, bid in auctions ending only 3.5 days apart on average (median 1 day), and disappear from the data after 9.2 days on average (median 3.6 days). This short "lifespan" of individual bidders in an eBay consumer-electronics market motivated the model of the demand-side as composed of generations of bidders, with each generation only persisting for a relatively small number of periods. Since new bidders appear constantly on eBay and compete with the remaining "old" bidders, a model with overlapping generations results: in the average MP3 player auction that received at least one bid, 7.5 unique bidders participated, but the long-term average over three months is only 3.2 unique bidders per such auction.[EC1] Several empirical regularities concerning eBay sellers are also captured in the model assumptions: the majority of the category volume is sold by institutional sellers who sell many units. For example, the 340 (8.5 percent) sellers observed selling at least five different units in three months sell an average of 16 units each, and account for 53 percent of the volume. These institutional sellers are much more long-lived than the buyers—they are observed throughout the data, and they space their sales 6 days apart on average. A majority of them specialize on just one model of MP3 player (median 1, average 1.9 out of 30 top-selling models). This motivated the model of a seller as infinitely lived and specialized in one type of good. Zeithammer (2006) discusses more details of this data, and he also provides empirical evidence for forward-looking bid-shading. Bid-shading was also detected in state highway-procurement auctions by Jofre-Bonet and Pesendorfer (2003).

Managerial Implications

This work contributes to the management science literature by exploring general *qualitative* properties of a market institution of large and growing importance—the sequential auction. The model illustrates why and how to take the other side of the market into account when formulating one's own strategy, and what are the properties of the resulting equilibrium. Both managers and management scientists will benefit from the findings. Managers, whether optimizing a selling strategy for their eBay sales channel or formulating a bidding strategy for a procurement-contract auction, will benefit from a deeper understanding of what past outcomes should they take into account, how to carefully interpret the demand information contained in past prices, and why do they need to anticipate the future strategies of the opposite side of the market. Managers who design the market-rules themselves will

[EC1] Participation is partly unobserved on eBay because bidders can only submit their bid if it exceeds the highest bid at the moment. The number of observed participants thus underestimates the number of actual participants.

benefit from the "self-preservation instinct" finding that the possibility of bid-shading does not necessarily reduce the scope of auctions to only very profitable markets. Finally, management scientists will be able to build on the stylized model presented here in developing decision-theoretic statistical models that could be used to analyze the wealth of data generated by auction markets like eBay. The present model emphasizes the need for equilibrium analysis of such markets, and gives guidelines about the phenomena likely to exist as well as the factors that moderate them.

The proposed model abstracts away from several complexities of real-world marketplaces, so it can only provide qualitative properties of optimal buying and selling, not concrete quantitative prescriptions for managerial action. For example, the model predicts that selling strategies will exhibit pulsing near the zero-profit contour of the parameter-space, but remains silent on how to find the optimal length of the pulsing interval on eBay, or on the optimal quantity the pulsing seller should offer for sale. Analogously, the model predicts that bid-shading will be moderated by gains from trade, but remains silent on how to calibrate this effect in procurement settings. Finally, the model predicts that sellers will sell more often after high recent prices, but it does not provide a practical algorithm for calculating optimal closed-loop selling strategies under realistic assumptions about bidders and other sellers.

The strongest assumption employed by the present model is that the seller is a monopolist. While this assumption fits the highway-construction market well (the government is the only "seller" of such contracts), more work is needed to characterize concrete seller strategies in hyper-competitive markets like eBay. This paper shows that the "self-preservation instinct" result persists even with two competing sellers despite the fact that competition limits the equilibrium reduction of bid-shading away from the zero-profit contour. However, the paper does not explicitly investigate selling strategies when there are more than two sellers, and it does not investigate competitive selling strategies for individual sellers with only one item to sell despite the fact that such sellers are quite common on eBay.

It is clear from the paper that even the simplified abstract models are already difficult to solve, so actual implementation of the ideas proposed here remains beyond the scope of this paper. I hope that future research will take up this challenge and finish the work started here by formulating tractable models of equilibrium behavior under more realistic assumptions. While stopping short of this goal, this paper at least considers several alternative specifications of the proposed model, and demonstrates in §3 that the proposed qualitative properties of equilibrium behavior are robust to those changes in the assumptions. Moreover, even the abstract model does capture at least some aspects of practice because its assumptions are designed to capture empirical regularities of markets like eBay, at least regarding buyer and seller lifespans, repeat-bidding behavior, and revealed preferences.

EC.2. Additional Discussion of Properties of Optimal Selling

Figure EC.1 illustrates the qualitative properties of the equilibrium (Proposition 2). The discussion of optimal selling can be extended as follows: To limit the extent of bid-shading, the seller does not have to use a closed-loop strategy that interprets past prices. Suppose that the seller cannot learn from prices, but that she can pre-commit to any pattern of selling over time, for example suppose that she is forced to play an open-loop strategy. Given the overlapping-generations model with lifespan of two periods, only two open-loop strategies can be beneficial: pulsing (selling every other period) and always-selling. Pulsing trivially rules out bid-shading because it sells only once within each generation's lifetime. It also increases the amount of demand-side competition, delivering more profit per auction, but only half as often. Specifically, let R_{2G} be the expected profit with two entire generations

Figure EC.1 Qualitative Equilibrium Characterization

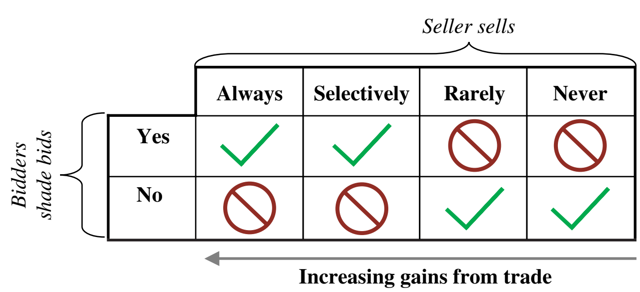

Figure EC.2 Effect of Bid-Shading and Learning from Prices on Optimal Selling Strategy

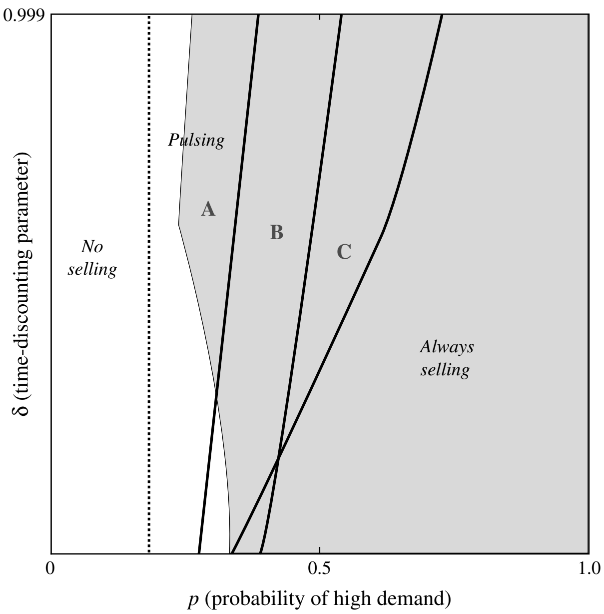

Notes. Remaining parameters are set to $L=0$, $H=3$. The three solid lines are upper bounds in (p,δ)-space of the seller always selling, ordered with the increasing sophistication of the seller: **A**: naïve seller that does not consider bid-shading and does not learn from past prices. **B**: a seller who considers bid-shading, but does not learn from past prices **C**: a seller who both considers bid-shading and learns from past prices (the $\pi=C_0$ constraint in Proposition 2). The shaded area represents the extent of bid-shading in equilibrium of the most sophisticated seller (the second constraint in Proposition 2).

bidding their valuations in a single auction ($R_{2G}=R_Q$ from §2 with $a=0$), and let R_{1G} be the expected profit with only one generation of bidders bidding their valuations in a single auction:

$$R_{2G}=[1-(1-p)^2]H+(1-p)^2L-c>(1-p)H+pL-c=R_{1G}.$$

Then, pulsing delivers the net expected profit $\Pi_{10}=R_{2G}/(1-\delta^2)$, and it is possible that Π_{10} exceeds the profit $\Pi_{11}=\Pi_Q^{all}$ from always-selling.

To assess the profits from always selling, it is essential that the seller understands the phenomenon of bid-shading. Suppose the seller naively does not take bid-shading into account, and assumes instead that all bidders bid their valuations. Then, the open-loop seller ends up selling too often, because she assesses Π_{11} too high. In particular, she thinks $a=0$ which can be shown to imply:

$$\Pi_{11}^{naive}=\frac{(1-\delta(1-p))ER_{2G}+\delta(1-p)ER_{1G}}{1-\delta}>\Pi_Q^{all}=\Pi_{11}.$$

This is illustrated in Figure EC.2 by the shift from curve **A** ($\Pi_{11}^{naive}=\Pi_{10}$) to curve **B** ($\Pi_{11}=\Pi_{10}$). Interestingly, compared to an open-loop seller who understands bid-shading, a closed-loop seller who cannot pre-commit may end up selling too often as well, especially when she is very impatient, i.e. when δ is low. This is illustrated in Figure EC.2 by curve **C** ($\pi=C_0$)—the boundary of always-selling from Proposition 2—crossing curve **B** for low values of δ. The closed-loop seller has a commitment problem because observing $H-a$ reveals a relatively high level of demand, so the bidders do not refrain from bid-shading as much as they do when facing an open-loop seller. Thus, the seller can sometimes prefer to live in an open-loop world. This illustrates why the equilibrium dynamic selling is not just a solution to a dynamic-control problem: the bidder's strategy changes based on the strategy played by the seller. The commitment problem of the closed-loop seller is a general feature of sequential auctions as demonstrated in a separate paper (Zeithammer 2007).

The policy implication for sellers is thus that bid-shading is an important aspect of bidding behavior, and sellers need to understand bid-shading as a strategic response of the bidders to their selling strategies. Adaptive learning from past prices and the ability to commit to a future pattern of selling are both useful strategies for limiting the extent of bid-shading. Taken together, these two selling strategies suggest a general intuition about optimal selling in auction markets, namely the benefit of spacing sales apart from each other in time, especially when the general profitability of the market is low for the seller. A combination of the two selling strategies is beyond the scope of this paper, but it is likely that a seller may want to restrict her decision-frequency to optimally balance the power of the two effects. For example, it is clear that the seller would prefer to be an open-loop seller when p is low and H is high: an open-loop seller would eliminate the $C_Q > \pi > C_{12}$ "wedge" of no selling in Figure 2.

From the point of view of a researcher analyzing a sequential auction market, Figure EC.2 illustrates the importance of understanding both the seller beliefs about the bidder strategy, and the closed-/open-loop nature of the seller's policy: with the same bidders, the different assumptions about the seller obviously lead to very different predictions about optimal selling. For example, researchers trying to identify the seller's costs from observed behavior would arrive to very different conclusions depending on the particular assumptions about learning and strategies chosen.

EC.3. Basic Model with Simpler Bidder Beliefs: Bidders Have Prior Beliefs About Old Competitors

Suppose the model is the same as in §2, but the bidders do not have knowledge of the seller's state. Instead, they assume that $\Pr(\textit{old High}) = p$. Then, Proposition 1 still applies, but clearly with *new High* bidders shading in every period, so $\omega = 1$. Therefore, the shading decrement is greater, let it be denoted by $b = \delta(1-p)(H-L)/(1-\delta p)$, which is a from Proposition 1 with $\omega = 1$. The short- term profits, indicated here by tilde to distinguish them from those in Proposition 2 are modified in that $\widetilde{R}_0 = \widetilde{R}_1$, i.e., the **1** state is not as lucrative as in the basic model, where bidders do not shade in **1**:

$$\widetilde{R}_0 = p(H-b) + (1-p)L - c = \widetilde{R}_1 \qquad \text{and} \qquad \widetilde{R}_Q = pH + p(1-p)(H-b) + (1-p)^2 L - c.$$

Therefore, the seller's learning process from Figure 1 still applies, but with the a replaced with b, and with price $H-b$ instead of H arising in state **1**. The fact that *new High* bidders shade in all states allows one simplification of the assumptions, namely it is no longer necessary to assume that *old* bidders win ties (see Footnote 3, and note that Proposition 1 still holds, but its proof becomes more complicated and is left as a challenge to the reader). Given these preliminaries, an analogue result to the characterization of Proposition 2 results here, without loss of too much generality focusing on the $L = 0$ case:

Proposition A2 (Equilibrium Characterization When Bidders Have Prior Beliefs). *The equilibrium of the auction market depends on the model parameters as follows when* $L = 0$:

- *When the relative gains from trade are so large that* $\widetilde{R}_0 > \delta p(H-c)/(1+\delta p)$, *the seller sells in every period and the bidders shade their bids down. The seller makes* $\Pi_Q^{all} = ([1-\delta(1-p)]\widetilde{R}_Q + \delta(1-p)\widetilde{R}_1)/(1-\delta)$.
- *When the relative gains from trade are medium such that* $-\delta^2 p(1-p)(H-c) < \widetilde{R}_0 \le \delta p(H-c)/(1+\delta p)$, *the seller does not sell in state* **0** *(after price L), sells in* $\{\mathbf{2}, \mathbf{1}, \mathbf{Q}\}$, *and the bidders shade their bids. The seller makes*

$$\widetilde{\Pi}_Q^{not0} = \widetilde{R}_Q + \delta p(1-p)[\widetilde{R}_0 + \delta p(H-c)]/(1-\delta)(1+\delta(1-p)+\delta^2 p(1-p))$$

- *When the relative gains from trade are so small that*

$$-((1-p)[1+\delta(1-p)]/(2-p)(1-\delta p)) < \widetilde{R}_0 \le -\delta^2 p(1-p)(H-c),$$

seller does not sell after prices L or (H-a), and bidders therefore do not shade their bids down. The seller either uses the pulsing strategy of selling every other period, or also sells in other informational states associated with the new learning environment without bid-shading.

- *When the relative gains from trade are even smaller such that* $\widetilde{R}_0 \le -((1-p)[1+\delta(1-p)])/((2-p)\cdot(1-\delta p))$, *the seller never sells.*

Proof. The proof is analogous to the proof of Proposition 2:
Case 1. *Bid-shading, Seller always sells*

$$\widetilde{\Pi}_0 = p(H-b) + (1-p)L - c + \delta[p\widetilde{\Pi}_1 + (1-p)\widetilde{\Pi}_0]$$

$$\widetilde{\Pi}_1 = p(H-b) + (1-p)L - c + \delta[p(H-c+\delta\widetilde{\Pi}_Q) + (1-p)\widetilde{\Pi}_0]$$

$$\widetilde{\Pi}_Q = pH + p(1-p)(H-b) + (1-p)^2 L - c + \delta[p\widetilde{\Pi}_Q + p(1-p)\widetilde{\Pi}_1 + (1-p)^2\widetilde{\Pi}_0]$$

and $\widetilde{\Pi}_0 > \delta\widetilde{\Pi}_Q$, $\widetilde{\Pi}_1 > \delta\widetilde{\Pi}_Q$, and $\widetilde{\Pi}_Q > 0$

Solution:

$$\tilde{\Pi}_Q = \frac{[1-\delta(1-p)]\tilde{R}_Q + \delta(1-p)\tilde{R}_1}{1-\delta}$$

and $\tilde{\Pi}_0 > \delta\tilde{\Pi}_Q$ the binding constraint, satisfied whenever $\tilde{R}_0 > \delta p(H-c)/(1+\delta p)$, implies the other constraints.

Case 2. *Bid-shading, Seller withholds supply in the low state*:

$$\tilde{\Pi}_1 = p(H-b) + (1-p)L - c + \delta[p(H-c+\delta\tilde{\Pi}_Q) + (1-p)\delta\tilde{\Pi}_Q]$$

$$\tilde{\Pi}_Q = pH + p(1-p)(H-b) + (1-p)^2L - c + \delta[p\tilde{\Pi}_Q + p(1-p)\tilde{\Pi}_1 + (1-p)^2\delta\tilde{\Pi}_Q] \quad \text{and}$$

$$\delta\tilde{\Pi}_Q > p(H-b) + (1-p)L - c + \delta[p\tilde{\Pi}_1 + (1-p)\delta\tilde{\Pi}_Q], \quad \tilde{\Pi}_1 > \delta\tilde{\Pi}_Q, \quad \text{and} \quad \tilde{\Pi}_Q > 0$$

The solution is:

$$\tilde{\Pi}_Q = \frac{\tilde{R}_Q + \delta p(1-p)[\tilde{R}_0 + \delta p(H-c)]}{(1-\delta)(1+\delta(1-p)+\delta^2 p(1-p))}$$

and the first two constraints are equivalent to $-\delta^2 p(1-p)(H-c) < \tilde{R}_0 \le \delta p(H-c)/(1+\delta p)$, while the last constraint $\tilde{\Pi}_Q > 0$ does not bind. The implied region of the parameter space is obviously non-empty.

Case 3. *No bid-shading, Seller withholds supply in the* **1** *and* **0** *states*: Now suppose that $\tilde{\Pi}_1 \le \delta\tilde{\Pi}_Q$ in Case 2. Then, the seller does not sell in state **1**. From Proposition 1, the bidders have no incentive to underbid, and the game reverts to the game without bid-shading. The game without bid-shading is hard to analyze because the seller's ability to learn from prices decreases. It is clear that the seller exists completely whenever even pulsing is not profitable, which can be written in terms of $\tilde{R}_0$ as

$$-(1-p)[1+\delta(1-p)]/(2-p)(1-\delta p) < \tilde{R}_0 \le -\delta^2 p(1-p)(H-c)$$

Case 4. *No selling*: When $\tilde{R}_0 \le -((1-p)[1+\delta(1-p)])/((2-p)(1-\delta p))$ even pulsing is unprofitable, the seller does not sell. This constraint is identical to the

$$(1-p)^2L + [1-(1-p)^2]H - c < 0 \quad \Leftrightarrow \quad p(H-L)/(c-L) < 1/(2-p)$$

from Proposition 2. □

Discussion of Proposition A2: Please see Figure EC.3 for an illustration of Proposition A2 with $L = 0$ and $\delta = 0.9$, i.e. the same as in Figure EC.1. The case of simpler beliefs clearly also has a simpler equilibrium than the basic model because $\tilde{\Pi}_1^{not0} > \delta\tilde{\Pi}_Q^{not0}$ implies $\tilde{\Pi}_Q^{not0} > 0$. This holds because state **1** is not as lucrative here, so the seller is never in a position where she would like to sell in **1** but not in **Q**. In other words, the $C_Q > \pi > C_{12}$ "wedge" of no selling that arises in the basic model and is shown in Figure EC.1 does not occur here. Except for this simplification, all the intuition for Proposition A2 is the same as that for Proposition 2. For example, the bidders still shade as long as the seller sells in state **1**.

In comparing the two versions of the basic model, it is clear that the seller is worse off with the "prior-beliefs" whenever the gains of trade are large: when the seller always sells, the bidders with simpler beliefs shade more ($b > a$) and more often (in all states instead of just in **0** and **Q**). On the other hand, when the gains from trade are relatively low but still high-enough to justify selling, the seller is better off with the "prior-beliefs" bidders: the bid-shading ceases at higher levels of H for any given level of p. This is because greater shading requires greater gains to be supported.

EC.4. Basic Model with Independent Bidders Within Each Generation

Suppose each bidder (rather than "each generation" as in the basic model) can be *High* or *Low*, the probability of *High* is p, and two new bidders are drawn independently in each period. For tractability, it is necessary to assume that the bidders always reduce their bid, for example because they have "prior" beliefs about their current *old* competition (see EC.3. in this Supplement). A direct analogue to Proposition 1 holds, with the modification that the *new High* bidders always bid $H - a$ such that

$$a = \delta[(1-p)^2(H-L) + (1-(1-p)^2)a] \quad \Rightarrow \quad a = \delta(1-p)^2(H-L)/1-\delta[1-(1-p)^2]$$

Figure EC.3 Equilibrium Characterization

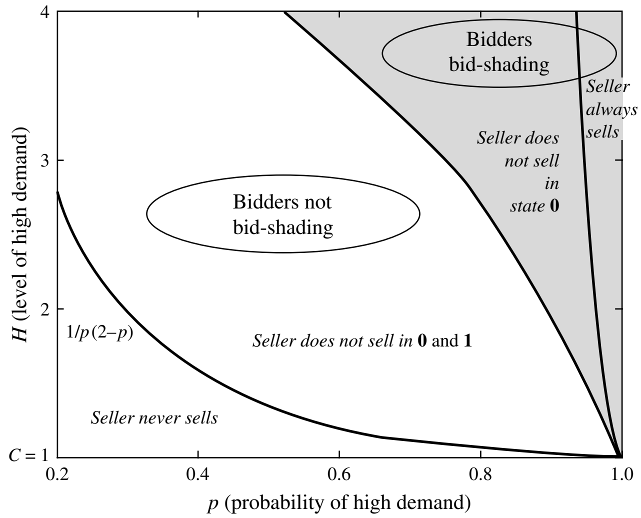

Notes. Illustration of Proposition A2 by fixing reamaining parameters to $L=0$, $\delta=0.9$. To shaded area shows where bidders shade their bids in equilibrium, with the general bidder strategies shown in ovals. The seller strategies are shown in italics.

The price-table is a generalization of Table 1 in that it has one more column H/L to capture the situation when the new bidders are of different types. The seller still cares only about the number of high bidders, because the additional *Low* old bidder in H/L does not affect learning or payoffs. Therefore, the state H/L is the same as **1**, and the price-table over payoff-relevant states is as shown in Table EC.1.

The inference from prices is therefore modified such that:

price $=H-a$: Then the new bidders must have been both *High* or H/L, and so there must be either one remaining *High* bidder (when there were no *old High* bidders and the new bidders were both *High* and when there was one *old High* bidder and the new bidders were actually H/L), or two remaining *High* bidders (there was one *old High* bidder and the new bidders were actually both *High*):

$$(p_0, p_1, p_2) \to \left(0, \frac{p_0p^2+2p_1p(1-p)}{p_0p^2+p_1p(2-p)}, \frac{p_1p^2}{p_0p^2+p_1p(2-p)}\right).$$

So from the $\mathbf{Q}=((1-p)^2, 2p(1-p), p^2)$ state, the price of $H-a$ sends the seller to a "mixed" state $\mathbf{M0}=(0, 5(1-p)/(5-3p), 2p/(5-3p))$. The most important aspect of **M1** is not the exact probability p_2, but the fact that $p_0=0$. This implies that another $H-a$ then tells the seller that there must have been exactly one old bidder, and the subsequent state is just $\mathbf{M1}=(0, 2(1-p)/(2-p), p/(2-p))$. In other words, yesterday's p_1 drops out of the updating equation. Therefore, the seller can only be in five payoff-relevant states in this game: **0**, **Q**, **1**, **M0**, and **M1**, one of which, the **1** state, is only accessible when the seller always sells (please see Figure EC.4 for an illustration of the learning process). This allows a Bellman analysis analogous to that in §3.4.

Table EC.1 Prices in the Auction Market with Bid-Shading

Number of *old High* bidders (belief)	New bidder type; probability		
	$2L$ $(1-p)^2$	$1L$ & $1H$; $2p(1-p)$	$2H$; p^2
0 (p_0)	L	L	$H-a$
1 (p_1)	L	$H-a$	$H-a$
2 (p_2)	H	H	H
	Number of *High* bidders left		
0 (p_0)	0	0	1
1 (p_1)	0	1	2
2 (p_2)	0	1	2

Figure EC.4 Seller's State Transitions in a Model with Independent Bidders

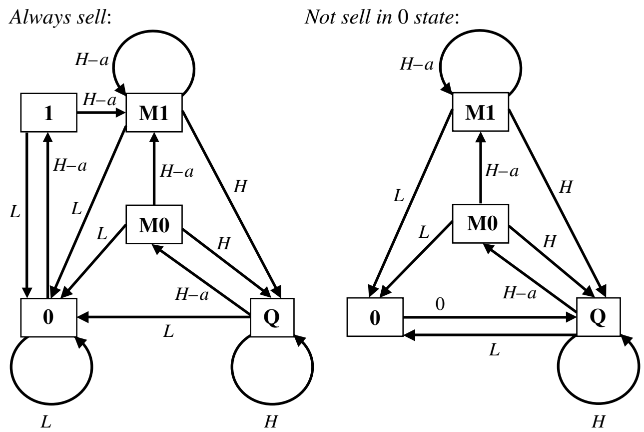

Consider the non-degenerate situation, in which the seller does not sell in state **0**. Since the **M1** state is more lucrative than the **M0** state and **M1** is only accessible through **M0**, the seller either sells in both **M0** and **M1** or uses the pulsing strategy. Whenever withholding supply in **M0** is a credible threat, the bidders will not shade their bids. The solution of $\{\Pi_Q, \Pi_{M0}, \Pi_{M1}\}$ is extremely algebraically involved, but it is possible to argue that withholding supply in **M0** eventually becomes a credible threat as the profitability of the auction-market shrinks, and so the main qualitative conclusion of Propostion 2 will hold.

EC.5. The Basic Model with Supply-Side Competition

Suppose the bidders are the same as in the basic model, and let there be two sellers, each identical to the monopolist considered in the basic model, i.e. each capable of procuring and selling one unit of the good in each period. Let the demand-side of the market remain exactly as in the basic model, i.e. overlapping generations of types L or H, with two buyers per generation with two-period lifespans. When both sellers decide to enter the market, they compete for the same buyers, both obtaining lower sale prices than if selling alone. A simple model of such competition is a third-price Vickrey auction that sells both units of the good to the first and second highest bidders for the third highest bid. Because of dominant strategies of Vickrey auctions, equilibrium bidding in the third-price two-unit auction is almost the same as in the second-price single-unit auction analyzed in Proposition 1. In particular, all *Low* bidders and *old* bidders always bid their valuation while *new High* bidders underbid as long as they have a chance of making a positive surplus should they lose, with the precise magnitude of bid-shading depending on the equilibrium being played. This model of the market thus focuses the impact of seller competition on the seller side of the market, because the buyer-side of the market responds to the aggregate supply of both sellers using essentially the same bidding function that was the best response to the supply of a single seller.

The seller model involves an entry-coordination problem because in some time-periods, the demand-side of the market can only provide positive profits for a sale of a single unit. If the two sellers were asked to choose simultaneously whether or not to sell in the beginning of a given period, mixing strategies would no doubt result, and the market would be inefficient because the coordination problem would not be solved completely. Since the real world does not proceed in discrete time, it is more realistic to assume that the sellers know of each other's most recent selling decisions and avoid overcrowding the market. The eBay sellers probably come to the site frequently (but not continuously), examine the recent demand for their products by looking at the recent prices, and list another unit of the good if the demand seems high and there are not many competing recent listings by other sellers already. In other words, sellers top-up the supply of the good as long as there is sufficient expected demand to make that decision more profitable than the decision to wait. One way to capture such an ongoing dynamic in a discrete-time model is to assume that in the beginning of each period, one of the sellers is chosen by a coin-flip to decide first. The second seller in that period then observes the first seller's entry decision (but not the outcome of that auction) and decides whether or not to enter himself as well. Since the second seller does not observe the outcome of the first seller's auction

Figure EC.5 Entry Game Between the Sellers

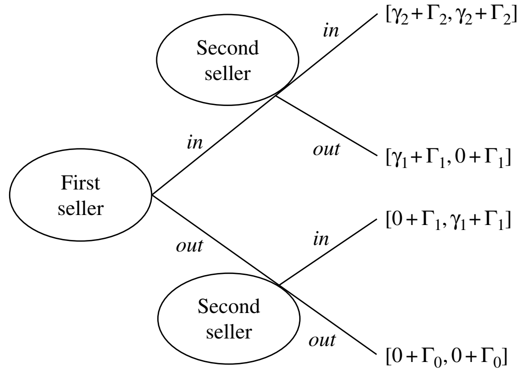

before making her selling decision, both sellers are on equal footing in terms of information about current demand, and the second seller thus does not learn anything new from the first seller's choice; the sequencing merely solves the coordination problem whenever only one seller should be selling.

The entry game in an arbitrary period of the game can be further formalized as follows: the players are symmetric in their knowledge about current demand and in their costs, so the short- and long-term profits depend only the number of sellers who enter. Fix the current belief-state, and let γ_k be the current-period profits when k sellers are selling, and let Γ_k be the corresponding continuation profits to all sellers when k sellers are selling. Since selling today depletes high-valued buyers who may otherwise be present tomorrow, it must be true that $\Gamma_0 \geq \Gamma_1 \geq \Gamma_2$. Since the market price with two sellers selling is the third-highest bid while the price with one seller selling is the second-highest bid, $\gamma_1 \geq \gamma_2$. Finally, not selling means earning zero in the current period. The resulting game is displayed in extensive form in Figure EC.5, with the total profits shown in brackets as [first seller, second seller]. The play (*out*, *in*) never happens, because the first seller can always do better than Γ_1 by playing *in* whenever $\gamma_1 + \Gamma_1 > \Gamma_0$, i.e. when the second player would enter after observing the play *out* by the first player. Therefore, the entry game merely decides which of the two players will be the only player when the market can only bear a single player. Given $\gamma_1 + \Gamma_1 > \Gamma_0$, the first player enters, and the second player joins in if $\gamma_2 + \Gamma_2 > \Gamma_1$. Given the relative magnitudes of Γ_k, it is clear that both sellers enter if γ_2 is positive and big enough whereas there will only be one seller when γ_1 is positive and big enough but γ_2 is negative.

The magnitudes of γ_k and Γ_k are a function of the knowledge state, and the demand-side assumptions provide their structure. As long as the buyers underbid, the same four knowledge-states $\{\mathbf{0}, \mathbf{Q}, \mathbf{1}, \mathbf{2}\}$ as in the basic model lie along the equilibrium path, and the amount of selling in each state depends on the overall profitability of the market. The optimal steady-state selling is captured most parsimoniously by considering the total market profits, i.e. the total market profits of both sellers in a given state Π: $\{\mathbf{0}, \mathbf{Q}, \mathbf{1}, \mathbf{2}\} \to \mathbf{R}$. This parametrization also maintains a close connection to the monopoly model of §2: analogous to the monopolist's state-dependent selling function *sell*: $\{\mathbf{0}, \mathbf{Q}, \mathbf{1}, \mathbf{2}\} \to \{\mathbf{yes}, \mathbf{no}\}$, a state-dependent entry-function *#entrants*: $\{\mathbf{0}, \mathbf{Q}, \mathbf{1}, \mathbf{2}\} \to \{\mathbf{0}, \mathbf{1}, \mathbf{2}\}$ captures the optimal selling behavior in the competitive context. A particular entry-function is a perfect Bayesian Nash equilibrium of the game between the sellers when the channel profits satisfy the appropriate Bellman equations dictated by learning and the number of entrants in each state is an equilibrium of the entry game conditional on the channel profits. Table EC.2 outlines for each knowledge state the current and future profits in the sales channel as a function of the number of sellers selling. The learning from prices can be discerned from the future profits.

To determine the state-dependent γ_k and Γ_k, note that the future channel profits $\Pi_{\mathbf{n}}$ are shared equally between the two sellers because the order-assignment for the entry game is random. In other words, when $\Pi_{\mathbf{n}}$ are the total equilibrium profits of the two sellers together in a future state $\mathbf{n}$, each seller can expect to get $\Pi_{\mathbf{n}}/\mathbf{2}$ for himself, so $\Gamma = \delta\Pi_{\mathbf{n}}/\mathbf{2}$. Note that this makes the model with competition coincide exactly with the monopoly model when there is no state with both sellers selling, i.e. when in all knowledge states, $\gamma_2 + \Gamma_2 < \Gamma_1$. In that case, at most one seller sells in a period, each seller gets half of the monopoly profits in expectation, and all other predictions of the monopoly model hold exactly. This will definitely be the case along the zero-profit contour in the (p, H) parameter space,

Table EC.2 Total Market Profits as a Function of State and Number of Entrants

Current state	Number of sellers	Total current profit	Future profit (discounted by δ)
0	**1**	R_0	$p\Pi_0 + (1-p)\Pi_1$
	2	$2L$	Π_0
1	**1**	R_0	$p\Pi_0 + (1-p)\Pi_2$
	2	$2R_0$	$p\Pi_0 + (1-p)\Pi_1$
2	**1**	H	Π_Q
	2	$2R_0$	$p\Pi_0 + (1-p)\Pi_2$
Q	**1**	R_Q	$p^2\Pi_0 + p(1-p)\Pi_1 + (1-p)\Pi_Q$
	2	$2R_0 - 2p(1-p)(H-a-L) < 2R_0$	$[p + p(1-p)]\Pi_0 + (1-p)^2\Pi_2$
Any	**0**	0	Π_Q

so the basic monopoly model is good in characterizing the behavior of the market near its existence threshold however many sellers there are

The competition among multiple sellers makes it harder but not impossible to withhold supply in order to discourage bid-shading. It is straightforward to demonstrate that *entry*: $[\mathbf{0}, \mathbf{Q}, \mathbf{1}, \mathbf{2}] \to [0, 1, 0, 1]$ is still sometimes the best response to bid-shading and since there is no selling in state **1**, the subsequent best response of the buyers is to stop bid-shading. The entry pattern [0,1,0,1] produces the following Bellman equations of the channel profits:

$$\Pi_0 = \Pi_1 = \delta\Pi_Q, \qquad \Pi_2 = H + \delta\Pi_Q$$
$$\Pi_Q = R_Q + \delta[p^2\Pi_0 + p(1-p)\Pi_1 + (1-p)\Pi_Q]$$

Therefore, $\Pi_Q = R_Q/[(1-\delta)(1+p\delta)]$. The entry pattern is an equilibrium of the game between the sellers if the following incentive constraints hold:

1: no entry $\Leftrightarrow R_0 < \delta/2[\Pi_Q - p\Pi_0 - (1-p)\Pi_1]$
1: no entry $\Leftrightarrow R_0 < \delta/2[\Pi_Q - p\Pi_0 - (1-p)\Pi_2]$
Q: one entrant $\Leftrightarrow R_Q > 0$
Q: not second entrant $\Leftrightarrow R_0 - p(1-p)(H-a-L) < \delta/2[\Pi_Q - [p + p(1-p)]\Pi_0 - (1-p)^2\Pi_2]$
2: one entrant $\Leftrightarrow H > 0$
2: not second entrant $\Leftrightarrow R_0 < \delta/2[\Pi_Q - p\Pi_0 - (1-p)\Pi_2]$

The two critical, potentially binding, constraints that imply all the rest are (**1,** no entry) and (Q, one entrant), equivalent to: $R_Q > 0$ and $2R_0 + \delta(1-p)H < \delta(1-\delta)\Pi_Q \Leftrightarrow 2R_0 + \delta(1-p)H < R_Q/(1+\delta p)$. It can be shown that here exists a region of the parameter-space when both of these constraints hold. Intuitively, there will be no entry in **1** when the short-term profit R_0 is negative, but the profit from waiting positive and sufficiently large.

It is interesting to compare the above incentive constraints with those of a monopolist. A monopolist captures the entire future income stream, so she is more likely to forego selling in **1**, namely whenever $R_0 + \delta(1-p)H < R_Q/(1+\delta p)$. Therefore, bid-shading will happen more often with competition. When the market becomes profitable enough to sometimes accommodate both sellers, the basic dynamics of the monopoly situation remain, but there is more selling and hence more bid-shading.

EC.6. Learning from Prices in the Continuous Model

Suppose $\{b_1(v \mid W_t), b_2(v \mid W_t)\}$ are monotonically increasing in v for every W_t. Everyone starts out (correctly) believing that the old bidders follow H_1, so $W_1 = H_1$. Price p_t is the upper bound on everyone's bids at time t, implying $\bar{v}_{t+1} = b_1^{-1}(p_t \mid W_t)$—the maximum possible valuation of new bidders given p_t. The next-period belief then depends on the age of the winner, who bid p_t. Let $\alpha_t = \alpha(p_t, W_t) = \Pr(old\text{won}_t \mid p_t, W_t)$. Then, the belief evolves according to the following transformation T: $W_{t+1}(x) = T(W_t \mid \bar{v}_{t+1})(x) = \alpha_t H_1(x \mid x < \bar{v}_{t+1}) + (1-\alpha_t)H_2(x \mid \bar{v}_{t+1})$ where H_1 is the truncated distribution of the highest valuation within a generation of bidders, and H_2 is the conditional distribution of the second highest valuation within a generation given the first highest valuation:

$$H_1(x \mid x < \bar{v}) = \left(\frac{F(x)}{F(\bar{v})}\right)^N, \qquad H_2(x \mid \bar{v}) = \left(\frac{F(x)}{F(\bar{v})}\right)^{N-1}.$$

Therefore, the set of all possible beliefs W_t can be parametrized by $\{(\alpha, \bar{v})\} \in [0,1]^2$, and the transition depends on b_1 (through both α and $\bar{v}$) and b_2 (trough α). We can assume WLOG that all agents know not only W, but also its coordinates $(\alpha, \bar{v})$, the second coordinate is just the upper bound of W's support, and the first coordinate is known because $H_1(x \mid x < \bar{v})$ and $H_2(x \mid \bar{v})$ are both known, so the agent can solve for α given W. We will therefore write the bidding strategies as $b_i(v \mid \alpha, \bar{v}) \equiv b_i(v \mid W)$. The transition in terms of the coordinates is:

$$\bar{v}_{t+1} = b_1^{-1}(p_t \mid \alpha_t, \bar{v}_t)$$

$$\alpha_{t+1} = \frac{g_{old}(p_t \mid \alpha_t, \bar{v}_t) G_{new}(p_t \mid \alpha_t, \bar{v}_t)}{g_{old}(p_t \mid \alpha_t, \bar{v}_t) G_{new}(p_t \mid \alpha_t, \bar{v}_t) + g_{new}(p_t \mid \alpha_t, \bar{v}_t) G_{old}(p_t \mid \alpha_t, \bar{v}_t)},$$

where the αtransition follows from Bayes Theorem: g's are the distributions of maximum old-bidder bid $g_{old}(b \mid W_t)$ and maximum new competing-bidder bid $g_{new}(b \mid W_t)$. Clearly, the transition of α is quite involved, with G involving derivatives of both bidding functions:

$$\frac{\alpha_{t+1}}{1-\alpha_{t+1}} = \frac{w(b_2^{-1}(p_t \mid \alpha_t, \bar{v}_t)) H_1(b_1^{-1}(p_t \mid \alpha_t, \bar{v}_t))((db_2^{-1}(p_t \mid \alpha_t, \bar{v}_t))/db)}{h_1(b_1^{-1}(p_t \mid \alpha_t, \bar{v}_t))((db_1^{-1}(p_t \mid \alpha_t, \bar{v}_t))/db) W(b_2^{-1}(p_t \mid \alpha_t, \bar{v}_t))}$$

There are three regions depending on whether or not the *new* or *old* bidders could have won:

$$p_t > b_1(1 \mid \alpha_t, \bar{v}_t) \quad \to \quad \bar{v}_{t+1} = 1, \qquad \alpha_{t+1} = 1$$

$$b_2(\bar{v}_t \mid \alpha_t, \bar{v}_t) < p_t < b_1(1 \mid \alpha_t, \bar{v}_t) \quad \to \quad \bar{v}_{t+1} = b_1^{-1}(p_t \mid \alpha_t, \bar{v}_t), \qquad \alpha_{t+1} = 0$$

$$p_t < \min[b_2(\bar{v}_t \mid \alpha_t, \bar{v}_t), b_1(1 \mid \alpha_t, \bar{v}_t)] \quad \to \quad \bar{v}_{t+1} = b_1^{-1}(p_t \mid \alpha_t, \bar{v}_t), \qquad \alpha_{t+1} = \text{ as above}$$

Therefore, very high prices lead to the starting state (1, 1), and there are prices that indicated that a new bidder won for sure, as stated in the main text of the paper.

Given the above learning, the seller maximizes the net present value of profits Π, starting by assumption in state $\mathbf{Q}$. Therefore, the steady-state profitability of any state W is captured by Bellman equations analogous to those of the basic model:

$$\Pi(W) = \max_{(sell,\, not\ sell)} \{E[p + \delta\Pi(T(W \mid p))] - c, \delta\Pi(\mathbf{Q})\}.$$

To obtain a continuous selling function λ, one can assume that the production cost c is a temporary random shock drawn in the beginning of each period, and a private information of the seller: Given c, every state W is mapped into the selling decision, with selling whenever the expected contemporaneous revenue exceeds the marginal cost of production plus the discounted net effect on future profits, i.e. whenever $E(p) > c + \delta\{\Pi(\mathbf{Q}) - E[\Pi(T(W \mid p))]\}$. The selling probability λ (from the bidders' point of view) then arises as the probability that the future draw of c is low enough for the selling condition to be satisfied. Finally, the demand-supply market equilibrium is characterized jointly as a triple of functions $[\lambda(W), b_1^{\lambda}(v \mid W), b_2^{\lambda}(v \mid W)]$, all satisfying their respective best-response conditions. The exact properties of these functions under specific distributional assumptions can be numerically approximated by simulation because the space of beliefs can be parametrized to the unit square.

EC.7. Proof of Proposition 3

PROPOSITION 3. *For every pair of selling-probability functions λ and μ such that λ involves less selling than μ in every state*: $\lambda(w) < \mu(w)\ \forall\, w \in [0,1]^2$, *the new bidders bid more as their best response to λ than to μ*: $b_1^{\lambda}(v) > b_1^{\mu}(v)\ \forall\, v \in [0,1]$.

PROOF. This proof derives the bidding strategy of *new* bidders, from which the claim follows immediately. For the purposes of the proof, it is convenient to redefine the transformation T in terms of the implied upper bound on surviving bidders' valuations: $T(W_t \mid x) \equiv T(W_t \mid x = b_1^{-1}(p_t \mid W_t))$. Also, let $Y_1 \sim F_1$ be the distribution of the valuation of the highest randomly-selected competing bidder within a generation (highest of $N-1$ iid bidders), and let $(Y_2 \mid Y_1) \sim F_2(\cdot \mid Y_1)$ be the conditional distribution of the valuation of the second highest competing bidder within a generation given the valuation of the first highest competing bidder. Consider a *new* bidder with particular v in state W, and suppose that the other $N-1$ new bidders bid according to $b_1(\cdot \mid W)$, while the remaining old bidders bid according to $b_2(\cdot \mid W)$ (which obviously exists, and is increasing in valuation). Assume that $\lambda(\alpha, \bar{v})$ is continuous and nondecreasing in $\bar{v}$. By playing $b_1(z \mid W)$, the focal bidder (called "I " from here on) with valuation v earns positive expected surplus in three situations that differ in the amount of surplus earned and in the implied state tomorrow:

(1) Surplus $v - b_1(z \mid W)$ if I win now , i.e. when $z > y_1$ & $b_1(z \mid W) > \max(b_2 \mid W)$.

(2) Surplus $\delta\lambda(T(W \mid y_1))[v - b_2(v \mid T(W \mid y_1))]$ if the highest competitor in my generation wins today (and thus exits), and I beat both my second highest competitor tomorrow and tomorrow's *new* bidders, i.e. when: $z < y_1$ & $y_2 < v$ & $b_1(y_1 \mid W) > \max(b_2 \mid W)$ & $b_2(v \mid T(W \mid y_1)) > \max[b_1 \mid T(W \mid y_1)]$.

(3) Surplus $\delta\lambda(T(W \mid b_1^{-1}(b_2(m))))[v - b_2(v \mid T(W \mid b_1^{-1}(b_2(m))))]$ if the highest *old* bidder wins today, and I beat both my highest *new* competitor and tomorrow's *new* competition, i.e. when $b_1(z \mid W) < \max(b_2 \mid W)$ & $v > y_1$ & $b_2(v \mid T(W \mid b_1^{-1}(b_2(m)))) > \max[b_1 \mid T(W \mid b_1^{-1}(b_2(m)))]$.

The objective function in terms of z is:

$$\begin{aligned}\Pi_1(z, v) = W(g(z))F_1(z)[v - b_1(z \mid W)] + \delta \int_z^1 \lambda(T(W \mid x))W(g(x))H_1(g^{-1}(v) \mid T(W \mid x))\\ \cdot F_2(v \mid x)[v - b_2(v \mid T(W \mid x))]\, dF_1(x)\\ + \delta \int_{g(z)}^1 \lambda(T(W \mid g^{-1}(m)))F_1(\min(g^{-1}(m), v))H_1(g^{-1}(v)T(W \mid g^{-1}(m)))\\ \cdot [v - b_2(v \mid T(W \mid g^{-1}(m)))]\, dW(m)\end{aligned}$$

where

$$W(x) = [\alpha H_1(x \mid Y_1 < \bar{v}) + (1-\alpha)H_2(x \mid Y_1 = \bar{v})], \qquad g(z) = b_2^{-1}(b_1(z)) \quad \Rightarrow \quad g^{-1}(z) = b_1^{-1}(b_2(z))$$

The first-order conditions, with $z = v$ applied after maximization of $\Pi_1(z, v)$ to reflect a symmetric equilibrium, lead to a differential equation for $b_1(v \mid W)$ in terms of the probability of winning the first period $\Psi(v \mid W) = W(g(v))F_1(v)$:

$$b_1(v \mid W) = \int_0^v [x - \delta\lambda(T(W \mid x))H_1(g^{-1}(x \mid T(W \mid x)))[x - b_2(x \mid T(W \mid x))]]\, d\Psi(x \mid W,\, x < v)$$

The equation for $b_1(v \mid W)$ describes the bidding strategy only implicitly because g is a function of both b_i. Nevertheless, it is evident that an decrease in λ leads to an increase in $b_1$and vice versa, as claimed in the Proposition. Also notable is the phenomenon of bid-shading: single-shot first-price sealed-bid strategy would just be $\int_0^v x\, d\Psi(x \mid W,\, x < v) > b_1(v \mid W)$. Finally, it should be pointed out that the strategy in Milgrom and Weber (2000) is a special case of $b_1(v \mid W)$ with $\delta = \lambda = H_1 = 1$, and a particularly convenient form of b_2 made possible by the lack of second-period entry. □

References

See references in the main paper.

Zeithammer, R. 2007. Optimal selling in dynamic auctions: Adaptation vs. commitment. *Marketing Sci.* Forthcoming.