MARKETING SCIENCE
Vol. 26, No. 6, November–December 2007, pp. 859–867
ISSN 0732-2399 | EISSN 1526-548X | 07 | 2606 | 0859

DOI 10.1287/mksc.1070.0269

Research Note

Optimal Selling in Dynamic Auctions: Adaptation Versus Commitment

Robert Zeithammer
Anderson School of Management, University of California–Los Angeles, 110 Westwood Plaza, Los Angeles, California 90095, rzeitham@ucla.edu

This paper analyzes optimal selling strategies of a monopolist facing forward-looking patient unit-demand bidders in a sequential auction market. Such a seller faces a fundamental choice between two selling regimes: *adaptive* selling that involves learning about remaining demand from early prices, and *commitment* selling that forgoes such learning and makes all selling decisions in the beginning of the game. A model of the game between the seller and the bidders is proposed to characterize the optimal regime choice. The model implies that the relative profitability of the two regimes depends on the expected gains from trade: when the expected gains from trade are low, commitment dominates adaptation, and vice versa.

Key words: auctions; game theory; durable-goods monopoly; optimal selling; commitment; adaptive dynamic pricing
History: This paper was received July 11, 2005, and was with the author 3 months for 4 revisions; processed by Ram C. Rao.

1. Introduction

Unit-demand goods are often auctioned in a sequence, one unit at a time: In the B2C context, unit-demand durables such as books and electronics are sold on Internet sites such as eBay. Nondurable unit-demand goods are also auctioned to consumers: For example, tickets to sports events on eBay, and travel tickets on Priceline (Fay 2004). In either case, the consumers have unit demand over time because today's buyers remain satiated in the future. In the B2B context, governments and large companies use sequences of auctions to procure public infrastructure and production inputs, respectively. Each procurement firm has limited demand for contracts over time because of its capacity constraints: Fulfilling one contract ties up the firm's resources for awhile (Jofre-Bonet and Pesendorfer 2003). All of the above sequential auction markets are growing in economic importance,[1] but little is known about optimal selling strategies in these markets when multiple units of the good can be produced and auctioned sequentially. This paper characterizes when a seller prefers to pre-commit to all her selling decisions, and when she prefers to sell adaptively, i.e., learn about demand from today's auction before making tomorrow's selling decision.

To analyze the profitability of these two selling regimes, the following model is developed: a monopolist seller lives for two periods, today and tomorrow, and can produce and auction off one unit of the good per period. The seller faces patient bidders who will bid again tomorrow if they lose today, so their bids today likely contain useful information about demand tomorrow. Therefore, it might seem that the seller should always sell adaptively to capture the rent associated with more informed selling in the second period. However, this single-agent dynamic-programming intuition is incomplete because the patient bidders are forward looking and *shade* their bids (e.g., bid less aggressively) today when they expect another auction tomorrow.[2] Because of bid shading, adaptation's information rent obtained tomorrow thus comes with a cost in the form of reduced revenue today. The precommitted seller cannot benefit from learning, but she can control bid shading better: When she rations the good by only selling a unit today and not tomorrow, the bidders

[1] In 2004: the relevant part of eBay transactions totaled about $17 billion (company report), Priceline transacted about $12 billion (company report), the federal government spent about $37 billion on bridge and highway construction (U.S. Department of Transportation report), numbers that are likely dwarfed by the procurement auctions of (more secretive) firms. eBay transaction volume has been growing about 20% per year.

[2] Such bid shading is not merely a theoretical construct: it has been detected in several real-world markets, for example in the eBay market for electronics (Zeithammer 2006) and in the California highway procurement market (Jofre-Bonet and Pesendorfer 2003).

respond by not shading their bids today. The adaptive seller cannot credibly ration the good—she is known to be tempted to sell again tomorrow by high realizations of demand, so profitable high-value bidders will shade today's bids more often when facing her. This paper analyzes the resulting trade-off between adaptation and commitment, and concludes that the overall potential gains from trade (expected bidder valuation minus seller production cost) determine which selling regime is optimal for the seller: commitment dominates adaptation when gains from trade are so low that the market can bear only one unit, and adaptation thus cannot increase profits enough to compensate for bid shading. When the gains from trade are high and the market can bear two units, adaptation in turn dominates commitment.

It is well known that in an exogenously given sequence of auctions for unit-demand goods, strategic forward-looking bidders shade their early bids to account for the option value of later auctions (Milgrom and Weber 1999, Engelbrecht-Wiggans 1994, Jeitschko 1999, and many others). Because commitment selling gives rise to a sequence of auctions that can be considered as given by the bidders, the commitment profits are straightforward to analyze using existing models. The main technical contribution of this paper is the analysis of adaptive selling, where future auctions depend endogenously on the early bids. When the seller sells adaptively, the bidders shade their bids today only when it is the best response to the seller, who will learn about demand from those bids and decide whether or not to sell again. This is different from Jeitschko (1999), who studies an uncertain but still exogenous future sale.

The adaptive seller has a *commitment problem*: Her selling decision tomorrow is susceptible to manipulation by the bidders today. An analogous commitment problem has been analyzed in the context of a monopolist, who has only one unit to sell, can set positive reserve prices, but cannot commit not to reauction the good in the future when no current bids exceed the reserve price and the good thus remains unsold. McAfee and Vincent (1997) study such a seller, and they demonstrate that as the speed of potential reauctioning becomes infinite, the auction with a sequentially optimal (optimal given reauctioning) reserve price produces the same expected revenue as an auction with a reserve price equal to the seller's valuation of the good. Even when resale is not instantaneous, the reserve's ability to increase revenue is greatly reduced compared to the classic no-reauctioning case in Myerson (1981).[3] Skreta (2006) strengthens the conclusions of McAfee and Vincent by showing that their auction is the optimal two-period mechanism for selling one unit of a good, so McAfee and Vincent actually compare simple auctions without strategic reserves to optimal auctions, and conclude that there is not much difference in revenue. For simplicity, the model analyzed here does not model the possibility of resale of unsold units explicitly, but approximates the profit outcomes by not allowing the seller to set reserve prices above her valuation. To keep the playing field between regimes level and to focus on the trade-off between bid shading and learning, even the commitment seller is also restricted to reserves at (or below) her valuation. In summary, both types of seller are assumed to be able to instantaneously reauction unsold units.

The restriction of the reserve price to the seller's valuation of the good is internally consistent with the assumed seller's production technology: The seller values the good at zero (a normalization without loss of generality [WLOG]), and she has a positive marginal production cost, which has to be sunk before the auction. Furthermore, there is no other market for the good, so the seller can recover the cost only by selling to the bidders. Because the production cost is sunk before the auction, the seller's opportunity cost at auction time is zero (her valuation of the good), and because the possibility of instantaneous resale makes positive reserves not credible, the seller uses an auction with zero reserve. Unlike in the models of McAfee and Vincent (1997) and Skreta (2006), which treat the sale of a single unit, the seller in this paper can produce and sell multiple units over time. The market power of the adaptive monopolist is thus weakened not only by the diminished effectiveness of a reserve price (which is assumed to weaken the commitment seller equally, as explained in the previous paragraph) but also by the additional bid shading. The adaptive seller thus has a commitment problem even without positive reserve prices because she is known to be tempted by high realizations of demand to sell again tomorrow.

Not allowing the adaptive seller to use arbitrary reserve prices is critical for her ability to learn about demand, because the bidders are in turn willing to reveal their valuations in the first period of the game. If the adaptive seller could set an arbitrary reserve price in the second period, she would set it to the highest-remaining bidder-valuation (revealed perfectly by first-period price), and extract all remaining consumer surplus.[4] In response to such ratcheting

[3] For example, a seller facing two bidders and a common per period interest rate of 5% can only gain 10% of nonreauctioning-seller's incremental revenue due to the reserve price (4% with five bidders).

[4] Note that this strategy would work even with the possibility of immediate reauctioning because the seller would know the highest remaining bidder valuation exactly. Knowing the highest remaining value exactly is fundamentally different from having only probabilistic beliefs over a continuum of possible highest remaining

(Freixas et al. 1985, Caillaud and Mezzetti 2004), pooling or complicated mixed bidding strategies would emerge in the first period to hide the demand information from the seller. If the early bids no longer contained information, commitment would prevail over adaptation because it could still reduce the extent of bid shading, while adaptation would no longer have an information rent.

The proposed model can be interpreted as a model of perfectly durable goods without secondhand markets. Because perfectly durable goods provide the same utility in each period after purchase, today's buyers continue to enjoy their durable goods forever. Because they are forever satiated by just one purchase, buyers of perfectly durable goods have unit demand over time identical to the buyers in this paper. Several existing papers also use the assumption of unit demand over time without secondhand markets to model durable goods (for example Stokey 1979, Besanko and Winston 1990, Dhebar 1994, and McAfee and Vincent 1997). The lack of secondhand markets restricts the durable-good interpretation to those durable goods that can be enjoyed by the same individual over time. In particular, the model does not capture products such as baby carriages, which are economically durable on a timescale of years only when there is a secondhand market (Bulow 1982, p. 318). The model proposed here also does not apply well to goods for which secondhand markets seem to play an important role—for example, cars (analyzed by Desai and Purohit 1998, Bruce et al. 2006, and others). The model predictions might change substantially with such secondhand resale markets because incentives for speculative buying by low-valuation bidders might arise (Garratt and Tröger 2006). However, note that in the real-world auction markets discussed in the first paragraph, near-future resale by winning bidders is limited by several factors: Obsolescence and shipping-related transaction costs limit resale of consumer electronics on eBay, nontransferability limits resale of sports event and travel experiences, and laws prevent resale of travel tickets and government procurement contracts.

Given the durable-goods interpretation of the model, it is important to compare the present results to the results of durable-good analyses in posted-price markets. The reduction of the adaptive seller's market power due to additional bid shading is a commitment problem analogous to the Coase conjecture in posted-price markets for durables. Coase (1972) argues and Stokey (1979) formally demonstrates that a long-lived price-setting monopolist selling unit-demand durables would always benefit from ex ante credible commitment not to produce more units in the future. This unqualified preference for commitment stems from the fact that the canonical models of posted-price durable-good markets do not involve sellers' demand uncertainty: Stokey (1979) and others after her assume there is a continuum of bidders that gives rise to a downward-sloping demand curve known to the seller.[5] In contrast, the auction-market model analyzed here assumes a finite number of bidders with valuations uncertain to the seller, so adaptive selling has an idiosyncratic learning advantage. With each selling regime having an idiosyncratic advantage, their relative profits are unclear, and this paper characterizes when one regime prevails over the other.

2. Model

There are two periods, 1 and 2. Everyone discounts second-period outcomes by factor $\delta < 1$.

Seller. A monopolist risk-neutral seller lives for both periods. She can produce one unit of an indivisible good per period, at a marginal cost of production $c > 0$, where 0 is WLOG her valuation of the unit.[6] When the seller decides to sell a unit, she first has to sink the cost c, and then sells the unit by a second-price sealed-bid auction with a reserve of zero. In the auction, ties are resolved by randomization. When the seller decides not to sell, she earns zero.

Bidders. All bidders have unit demand and independent private valuations for one unit of the good drawn from ($L = Low < c < H = High$) Bernoulli distribution with $\Pr(H) = p$.[7] For simplicity, let $L = 0$, and WLOG let $H = 1$ to set the scale of the utility function (labels L and H are retained wherever they make exposition clearer). There are two patient bidders who live for both periods. Because of the unit-demand assumption, a patient winner of the first auction does not bid in the second auction. In each period, there is also one impatient bidder (different individual in each period), who can only derive utility from the good in his period. The impatient bidders play no strategic

values, because the take-it-or-leave-it offer at the exactly known highest value is always credible (bidders cannot make the seller lower the price by waiting). See Levine and Pesendorfer (1995) for an in-depth exposition of the difference.

[5] Under that assumption, the seller anticipates the quantity demanded in the first period exactly, so there is nothing to learn about remaining demand from the first-period outcome, and adaptation has no advantage.

[6] The capacity constraint merely simplifies the analysis; it is not crucial to the result, as explained in the discussion.

[7] In the case of durables that provide a flow of benefits in each period, *valuation* is the net present value of the total utility derived from the good as in Stokey (1979), i.e., the net present value of the good if obtained in Period 1.

role; they merely provide sufficient competition to the patient bidders.[8]

Selling regimes. Two selling regimes are considered: adaptation and commitment. The adaptive seller decides whether or not to sell in the beginning of each period, and hence plays a Perfect Bayesian Nash equilibrium strategy: Before deciding about the second-period sale, she updates her beliefs about remaining demand using the first-period transaction price r_1. The commitment seller, on the other hand, decides her entire selling strategy in the beginning of the game and plays a Nash equilibrium strategy.

Information. The seller knows p, but not the individual valuations of the bidders. The bidders know the selling regime the seller is using and the seller cost c. There is a third-party auctioneer (such as eBay or Agentrics). The auctioneer collects all the sealed bids in each period and announces only the winner and the transaction price r_1.

2.1. Bidding Strategies

The impatient bidders have a dominant strategy to bid their valuation because they are effectively participating in a single-shot, second-price, sealed-bid auction. The surviving patient bidders bid their valuations in the second period for the same reason. In the first period, patient bidders sometimes shade their bids below their valuations because winning the first unit involves an opportunity cost of missing the second auction. (The winner of the first auction does not bid in the second auction because of the unit-demand assumption.) In particular, the patient bidders bid less than valuation when there is a chance of winning the second unit for a lower price than the first unit. Only *High* bidders have such a chance in this game because the *Low* bidders never make a positive surplus. The resulting first-period bidding strategy is the following lemma.

Lemma 1. *Let $a = \delta(1-p)(H-L)$, and let $\lambda(r_1)$ be the probability of a second sale as perceived by bidders in the first period. Suppose λ: $r_1 \to \{0, 1\}$ is a nondecreasing step function. Then there is an unique symmetric pure-strategy bidding equilibrium in the first period, and the patient High bidders bid $H - a\lambda(H-a)$ in the equilibrium. For all λ, patient Low bidders bid L in the first period*. (See the appendix for proof.)

It follows that $L < H - a < H$. Intuitively, each patient *High* bidder assumes that, should he lose the first auction, he would have to be the highest remaining bidder in order to get any expected surplus from a potential second auction. Therefore, he would have to be the price-determining (second-highest) bidder in the first auction, and so, the seller would observe price r_1 equal to his bid. The lemma is not novel in finding shading; it merely highlights the role of the seller in the game (through λ), and shows how the phenomenon of bid shading manifests itself under the present distributional assumptions. In analogous models with continuous distributions of valuations, all but the lowest bidders shade their bids down from valuation (Milgrom and Weber 1999, Jeitschko 1999, and many others). In that sense, the *High* bidders are the generic bidders here.

2.2. Selling Strategies: Commitment Seller

There are four possible strategies $(sell_1, sell_2)$, where 1 denotes a sale and 0 denotes no sale: $\{(0,0), (0,1), (1,0), (1,1)\}$. The waiting strategy $(0,1)$ is dominated by the rationing strategy $(1,0)$, because both involve three bidders bidding their valuations, and $\delta < 1$. Two nontrivial strategies remain: $(1,0)$ and $(1,1)$. In the case of $(1,0)$, the probability of second sale is $\lambda \equiv 0$, so there is no bid-shading (Lemma 1), and the profits thus amount to the profits from three iid bidders bidding their valuation: $\Pi_{10} = p^2(3-2p) - c$. When $(1,1)$, the probability of second sale is $\lambda \equiv 1$, so patient *High* bidders shade their bids down by a (Lemma 1), resulting in the expected revenue loss $a \Pr(r_1 = H - a) = [\delta(1-p)][p^2(3-2p)]$ in the first period (see Figure 1). However, the $(1,1)$ strategy also makes an additional profit $p^3(4-3p) - c$ on the second unit. Combining these two differences between Π_{11} and Π_{10} yields

$$\Pi_{11} = \Pi_{10} + \delta[p^2(-5p^2 + 9p - 3) - c]. \qquad (1)$$

From Figure 1, it is clear that the net revenue effect of always selling is always positive when $p > 1/2$. When $p < 1/2$, the net effect on revenues is ambiguous because three *High* first-period bidders make it positive, while only two *High* first-period bidders make it negative. Equation (1) expresses the net effect precisely as $\delta p^2(-5p^2 + 9p - 3)$, which is positive and increasing for $p > (9 - \sqrt{21})/10 \approx 0.44$. Equation (1) also characterizes optimal selling as

$$\text{always sell unless} \begin{cases} \text{ration when} \\ \quad \Pi_{10} > \Pi_{11} \Leftrightarrow c > p^2(-5p^2 + 9p - 3) \\ \qquad\qquad \text{(line C2 in Figure 2);} \\ \text{exit when} \\ \quad \Pi_{10} < 0 = \Pi_{00} \Leftrightarrow c > p^2(3-2p) \\ \qquad\qquad \text{(line IR in Figure 2).} \end{cases}$$

There is an interesting analogy with a static monopolist: Both the dynamic commitment seller and a static monopolist sell more units for a lower per unit price

[8] Assuming more than two patient bidders would be cumbersome under the Bernoulli distribution because there would be too many ties, and symmetric bidding equilibria in pure strategies would not exist.

Figure 1 **Extensive Form of the Commitment-Selling Game**

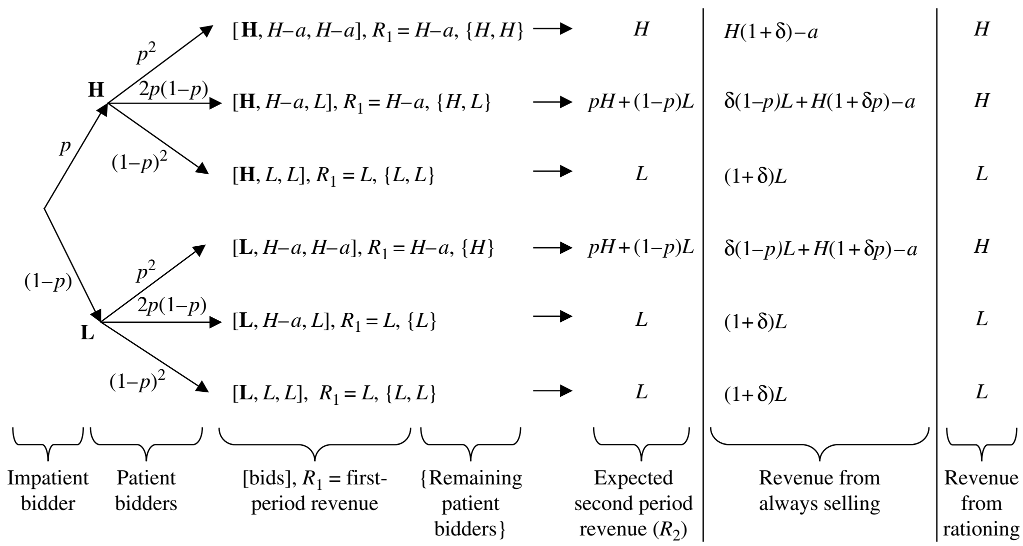

when demand is higher relative to cost, i.e., when the gains from trade are greater. Figure 1 implies immediately that the expected price from selling just one unit exceeds the expected unit price when two units are sold: The expected revenue from rationing exceeds both the expected price $E(R_2)$ of the second unit and the expected price $E(R_1)$ of the first unit with always selling. Moreover, Equation (1) and Figure 2 make it clear that restricting the quantity sold to one unit via rationing is beneficial for higher p or lower c, or both, i.e., for lower expected gains from trade. However, the isomorphism with the static monopoly is not trivial because when the commitment seller sells two units, $E(R_1)$ is not the same as $E(R_2)$: while there are more patient bidders in the first period, they shade their bids down, and the net effect is not a wash.[9]

Figure 2 **Optimal Selling Strategies for Both Regimes in** (p, c) **Space**

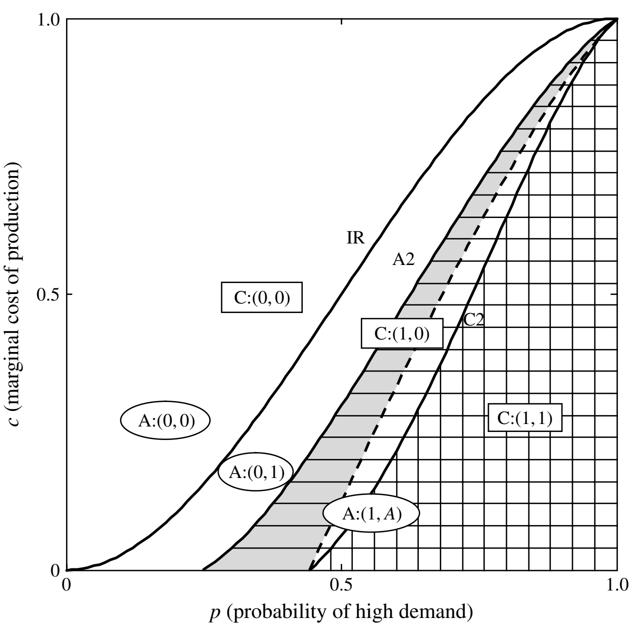

Notes. The optimal adaptive strategy is shown in ovals and preceded by "A:," optimal commitment strategy is shown in rectangles and preceded by "C:". Line IR is the boundary of selling for both selling regimes. Lines A2 and C2 are the boundaries of multiunit selling of the adaptive and commitment sellers, respectively. The dashed line is the line above which the seller would prefer commitment to adaptation. The shaded area represents the region in which the adaptive seller is worse off despite actively collecting the information rent. The vertical lines show the region of the parameter space, in which *High* bidders shade their first-period bids when facing a committed seller, the horizontal lines show the corresponding region for the adaptive seller. $\delta = 0.7$, increasing δ shifts line A2 toward the dashed line. The dashed line, as well as lines IR and C2, does not depend on δ.

Figure 1 also allows us to return to the durable-goods interpretation of the model with more concreteness. A key property of durability without secondhand trading is that early buyers drop out of the market, and these early buyers have higher valuations so future demand is reduced by an early sale. This key property holds in the present model because a first-period sale does skim off one high-valuation buyer with a positive probability: The first period has two patient bidders with expected valuation $E(v) = pH + (1-p)L$, while the second period has fewer than two patient bidders with a lower expected valuation of $pE(v) + (1-p)[p^2H + (1-p^2)L] < E(v)$. In this sense, second-period buyers are more price sensitive than first-period buyers whenever there is a sale in the first period, as should be true in a model of durable goods. Of course, bid shading prevents simple intertemporal discrimination; see footnote 9.

[9] Note that $H > H - a > pH + (1-p)L$, so the difference between the two expected revenues is ambiguous. It is $E(R_1) - E(R_2) = p^2(1-p)(H-L)[3(1-\delta) - p(3-2\delta)]$, which (conditional on a particular p) is positive for small δ and negative for large δ. Therefore, intertemporal price discrimination might or might not occur in this model.

2.3. Selling Strategies: Adaptive Seller

The adaptive seller may choose not to sell in the first period, collecting the waiting profit $\Pi_{01} = \delta\Pi_{10}$, but the rationing strategy $(1, 0)$ is unavailable to her. Therefore, to eliminate bid shading, the adaptive seller has to abstain from selling in the first period, and only obtain rationing profits in the second period. Alternatively, she might sell in the first period, and thus observe first-period price r_1 before deciding the second-period sale. Denote this strategy $(1, A)$, and let Π_{1A} be the resulting ex ante expected profit. Suppose the *High* patient bidders shade their first-period bids to $H - a$, so $r_1 \in \{L, H - a\}$. Price $r_1 = L$ implies zero second-period revenue because the remaining patient bidder is *Low*, so $\lambda(L) = 0$. On the other hand, $r_1 = H - a$ leaves one or two patient *High* bidders, implying $\lambda(H - a) = 1$.[10] Off the equilibrium path, any $r_1 > L$ implies one or two patient remaining patient *High* bidders, so $\lambda(r_1) = 1$, because it is always irrational to bid more than one's valuation, so the price-determining bidder bidding r_1 must have been *High*. To complete the seller's strategy, assume $r_1 < L$ implies a belief that the remaining bidder is *Low*. Hence λ is a step function with image $\{0, 1\}$, Lemma 1 applies, and *High* patient bidders always shade their bids as suggested. Therefore, the first-period profit of the adaptive seller is the same as that of an always-selling commitment seller (a coincidence of this model; continuous bidder types would involve slightly higher first-period revenue as discussed in Zeithammer 2007). While she faces the same bid shading, the adaptive seller saves on second-period production costs by producing the second unit only when it is profitable to do so:

$$\begin{aligned}\Pi_{1A} &= \Pi_{11} + \delta c[1 - p^2(3 - 2p)] \\ &= \Pi_{10} + \delta p^2[(-5p^2 + 9p - 3) - (3 - 2p)c], \quad (2)\end{aligned}$$

where the first expression characterizes the information rent $\delta c[1 - p^2(3 - 2p)]$, i.e., the ex ante expected cost saving from better-informed selling in the second period. This information rent implies that when the adaptive seller who sells in the first period would prefer to credibly ration, the commitment seller would also prefer to ration: $\Pi_{1A} < \Pi_{10} \Rightarrow \Pi_{11} < \Pi_{10}$. The second equality follows from Equation (1), and because $\Pi_{10} = p^2(3 - 2p) - c$, it is clear that the adaptive seller prefers $(1, A)$ over $(0, 1)$ when

$$\Pi_{1A} > \Pi_{01} = \delta\Pi_{10} \Leftrightarrow c < \frac{p^2[(3 - 2p)(1 - \delta) + \delta(-5p^2 + 9p - 3)]}{1 - \delta + \delta p^2(3 - 2p)}$$

(line A2 in Figure 2 | $\delta = 0.7$).

[10] When $r_1 = H - a$, the probability of two remaining *High* patient bidders is $p/(3 - 2p)$ (see Figure 1), and hence it is profitable to sell in the second period when $p/(3 - 2p) + [1 - p/(3 - 2p)]p > c \Leftrightarrow p(4 - 3p)/(3 - 2p) > c$ which is always satisfied when $\Pi_{01} > 0$, i.e., when it is profitable to sell at all ($\Pi_{01} < 0 \Rightarrow \Pi_{1A} < 0$ even when there is no bid shading).

3. Commitment Versus Adaptation

Neither selling regime profit-dominates the other because adaptation has both advantages and disadvantages over commitment, and the advantages and disadvantages operate in different regions of the parameter space. The advantage of adaptation is the information rent, but it is useful only when the market can bear the (adaptive) sale of more than one unit. On the other hand, the disadvantage of adaptation is the inability to credibly ration the good, which would be useful when the market can only bear one unit. The rest of this section makes this intuition precise.

The previous section demonstrates that the adaptive seller makes $\Pi_A \equiv \max\{\Pi_{1A}, \delta\Pi_{10}\}$ while the commitment seller makes $\Pi_C \equiv \max\{\Pi_{11}, \Pi_{10}\}$. Since both the always-selling commitment seller and the adaptive seller who sells in the first period face the same amount of bid shading, $\Pi_{1A} > \Pi_{11}$ because the adaptive seller gets the additional information-rent. (See the first equality in Equation (2).) Since $\delta\Pi_{10} < \Pi_{10}$, neither regime dominates, and it is immediate that the overall profit comparison Π_A versus Π_C reduces to the comparison of Π_{1A} versus Π_{10}:

Proposition 1. *When the sellers sell at all, the commitment seller is strictly better off than the adaptive seller when $\Pi_{1A} < \Pi_{10}$, and the adaptive seller is strictly better off when $\Pi_{1A} > \Pi_{10}$. In terms of the model parameters, $\Pi_{1A} < \Pi_{10} \Leftrightarrow c > (-5p^2 + 9p - 3)/(3 - 2p)$.*

The $\Pi_{1A} = \Pi_{10}$ contour is illustrated by the dashed line in Figure 2. Because the bid-shading decrement a is linear in δ, and δ is the same for the bidders and for the seller, the indifference boundary between the two selling regimes is independent of the discount factor. The Π_{1A} versus Π_{10} comparison can be interpreted as a comparison of the seller's cost to the buyer's expected valuation, i.e., to a level of expected gains from trade:

$$\frac{-5p^2 + 9p - 3}{3 - 2p} = 2p - 1 + \frac{p(1 - p)}{3 - 2p} \approx 2p - 1,$$

where p is just the expected valuation of a single buyer. Therefore, Proposition 1 can be interpreted as follows. The commitment seller is strictly better off than the adaptive seller when the expected gains from trade in the market are relatively low, i.e., when the production cost is high relative to the expected valuation of a single bidder. Conversely, the adaptive seller is strictly better off than the commitment seller when the expected gains from trade are relatively high, i.e., when production cost is high relative to the expected valuation of a bidder.

The selling behavior implied by Proposition 1 is as follows. When $\Pi_{1A} < \Pi_{10}$, the commitment seller always rations and gets Π_{10}, because $\Pi_{1A} < \Pi_{10} \Rightarrow \Pi_{11} < \Pi_{10}$ (see Equation (2)). The adaptive seller is thus worse off whether he sells in the first period and gets Π_{1A}, or does not sell and gets $\delta\Pi_{10}$. When $\Pi_{1A} > \Pi_{10}$, the adaptive seller always sells in the first period, and the commitment seller either rations or always sells. In both situations, the commitment seller is worse off: in the first by construction, and in the second because $\Pi_{1A} > \Pi_{11}$. The adaptive seller's commitment problem is highlighted in the following proposition.

Proposition 2. *There is a region of the parameter space in which the commitment seller is better off despite the adaptive seller selling in the first period and actively using the realized prices to avoid low-demand second periods. Specifically, for every* ($p > (9-\sqrt{21})/10$ *and* $\delta < 1$) *and* ($p \leq (9-\sqrt{21})/10$ *and* $\delta < (3-2p)/(5p^2-11p+6)$) *there is an interval of costs c such that* $0 < \Pi_{01} < \Pi_{1A} < \Pi_{10}$. (See the appendix for proof.)

The first inequality in $\Pi_{01} < \Pi_{1A} < \Pi_{10}$ implies that the adaptive seller would not prefer to wait. The second inequality implies that the commitment seller prefers to ration, so collects Π_{10}. The possibility of both inequalities holding at the same time follows from the fact that $\Pi_{01} = \delta\Pi_{10} < \Pi_{10}$. The resulting region, in which the adaptive seller is strictly worse off than the commitment seller despite voluntarily selling more often and collecting the information rent is shaded in Figure 2. It is clear that the region corresponds to intermediate costs.

4. Discussion

This paper analyzed seller preferences between two mutually exclusive regimes for auctioning unit-demand goods in a sequence: precommitment and adaptation. Neither regime dominates the other; the commitment regime is better at reducing the extent of strategic bid-shading behavior, but the adaptive regime resolves some demand uncertainty inherent in auction markets and can be more efficient. An adaptive seller can capture enough of the efficiency gains to prefer adaptation to commitment.

The relative profitability of the two regimes depends on the expected gains from trade: Low gains from trade favor commitment, while high gains from trade favor adaptation. The gains from trade are defined as the difference between expected demand (expected bidder valuation) and seller cost. To illustrate the result, consider a seller with particular marginal opportunity cost c and vary the demand from low to high: under both selling regimes, the seller participates in the marketplace as long as the demand is high enough for the rationing profits to be positive, i.e., as long as the overall market can bear at least one unit (with no bid shading). Just above the zero-rationing-profit level of demand, the seller effectively rations the good under both regimes to avoid bid shading, but she can ration only by waiting until the second period under the adaptive regime, while commitment allows her to credibly ration in the first period. She thus prefers the commitment regime when demand is low. As the level of demand becomes moderate, the adaptive regime starts to sell in the first period in order to capitalize on learning, while the commitment regime continues to prefer rationing. Interestingly, the commitment regime remains better off for a while despite selling less often (Proposition 2). The information rent from adaptation fully compensates for bid shading only when the demand is high. Then the information advantage of adaptation prevails, and the seller prefers adaptation to commitment.

Efficiency of the market is aligned with the preferences of the sellers for very high and very low expected gains from trade (above line A2 and below line C2 in Figure 2). For intermediate gains from trade (between lines A2 and C2), adaptive selling is more efficient because it realizes more gains from trade by resolving some of the demand uncertainty. Note that adaptive selling is not always more efficient: The adaptive seller sometimes (above line A2) prefers to forgo the learning and sell only in the second period. Then the commitment seller who offers a unit in the first period saves everyone a wait and hence is more efficient.

Several assumptions of the model are critical for the results presented here, while other assumptions can be relaxed. The discounting of second-period outcomes is necessary for the commitment seller to be sometimes better off because with $\delta = 1$ the adaptive seller could get the same net present value of profits as the credibly rationing commitment seller by waiting to sell in the second period. The capacity-constraint assumption about the seller (one unit per period) is not necessary—a seller allowed to produce two units in a period at constant marginal cost c and sell them by a third-price Vickrey auction (a natural extension of the present auction format to multiple units) would want to do so for extremely high p and extremely low c, not affecting the model predictions above line C2 in Figure 2. The distributional assumptions of the model analyzed in §2 are chosen for tractability, and they are not necessary for the qualitative results in §3 as explained in a companion paper that analyzes a model with a continuum of bidder types (Zeithammer 2007). Finally, not allowing the seller to credibly use a strategic reserve price above her own valuation of the good in the second period is necessary to preserve the demand-revelation property of the first auction. Not allowing reserve prices above seller's own valuation of the good can be motivated by the possibility of instantaneous resale of unsold units, as explained in §1. Such a motivation requires the seller's production cost c to be sunk before the auction. Otherwise, the seller could credibly set reserve prices at c, and the information rent of the adaptive seller would disappear—the seller would never want to withhold a unit from the market in the second period. To keep the playing field as level as possible, the present paper also assumes credible reserve prices away for the commitment seller (i.e., the commitment seller can also instantaneously resell unsold units). If the commitment seller could use a credible reserve price, her profits Π_{11} and Π_{10} would be boosted thanks to increased revenue, and Π_{11} would increase further due to reduced bid shading, so commitment would be preferred over adaptation more often than under the present assumptions.

While banning the use of an adaptive reserve price is necessary to preserve informativeness of first-period bidding under the present model of an auction and demand (second-price sealed-bid auction and unit demand), it is not necessary in other sequential-auction models. Caillaud and Mezzetti (2004) (hereafter CM) show that partial informativeness is preserved when the highest second-period valuation remains hidden from the seller in the first period. Assuming that the seller is also the auctioneer and bidders have the same recurring demand in every period, they therefore argue in favor of an ascending auction that does not reveal the highest (and thus the highest-remaining) valuation. This paper assumes that the auctioneer is a third party such as eBay or Agentrics, the seller only learns the price of the first object instead of all the bids, the auction is a second-price sealed-bid auction, and bidders have unit demand. Because of the unit-demand assumption, the first-auction winner drops out, and the price-determining bidder is thus exactly the highest remaining bidder, and so is fully revealed in the first period. Therefore, assuming adaptive reserve prices away is necessary to preserve informativeness in the present model, as claimed above. CM also assume that the good is perishable, so unsold units cannot be reauctioned, and the commitment problem discussed by McAfee and Vincent (1997) therefore does not operate. Potential for future exploitation (ratcheting) is thus greater in the CM's case of perishables than in the case of unit-demand durables analyzed here, where potential resale of unsold units restricts the seller's ability to use reserve prices credibly. A more complete future paper about selling unit-demand goods in auctions will build on the model proposed here by explicitly modeling the reauctioning of unsold units (not necessarily instantaneous) while allowing the adaptive seller to use arbitrary reserve prices.

Another fruitful direction for future work would be to allow for endogenous entry of bidders, compare the consumer surplus as compared to other trading mechanisms, and characterize the pros and cons of sequential auctions vis-à-vis standard posted-price selling from a marketing point of view (Shugan 2005). Finally, the present model does not account for the reference-point effect of one auction on the other, documented by Dholakia and Simonson (2005). It would be interesting to add this behavioral layer to the specification of preferences and see how the equilibrium interaction between the buyers and the seller changes.

Investigating rules of institutions such as eBay or government procurement auctions should allow researchers to be more concrete about the way commitment can be implemented. For example, both eBay and government procurement involve compulsory preannouncements of future auctions—the latter by law and the former by the fact that online auctions usually last at least a few days, so the sealed-bid endgame is preannounced (Zeithammer 2006). Preannouncements facilitate commitment selling because near-future selling decisions must be made before the outcome of the current auction is known. The natural extension of this paper to infinite-horizon settings would investigate the optimal duration of such preannouncements, trading off adaptation versus commitment. Another question posed by the present results is whether the auctioneer should fix the lead time of preannouncements or whether that decision should be left up to the sellers.

Appendix. Proofs of Lemmas and Propositions

Proof of Lemma 1. When $\lambda(H-a)=0$, bidding valuation is a dominant strategy. Suppose $\lambda(H-a)=1$. Focus on a *High* patient bidder, and suppose all other bidders are playing the proposed strategies. We show that bidding $H-a$ is the best response to every possible type of competing bidders. Consider the four different combinations of (impatient, other patient) competitors. When (H, H), the expected surplus from any bid $b' < H$ is zero because both *High* patient bidders advance to the second round and bid their valuations there. Deviation to $b' = H$ also keeps the surplus at zero by introducing the possibility of winning the first auction in a coin toss and paying H. When (H, L), no deviation to bid $b' < H$ changes the expected surplus of $\delta(1-p)(H-L)=a$, which arises from the chance that the second-period impatient bidder is *Low*. Deviation to $b' = H$ reduces the expected surplus by introducing the possibility of winning the first auction with zero surplus. When (L, L), no deviation to $b' > L$ changes the expected surplus of $(H-L)$.

Finally, when (L, H), the payoff from bidding $H-a$ is a, and there are two possible deviations:

(1) Deviating to $b' < H-a$ can reduce only the payoff because the bidder now loses for sure and the seller observes a low price of b' and may not sell again in the second period. In expectation, the bidder gets the expected surplus $\lambda(b')\delta(1-p)(H-L) \leq \lambda(H-a)\delta(1-p)(H-L) = a$.

(2) Deviating to $H-a < b' \leq H$ results in the payoff a for sure because the bidder wins the first auction and pays $H-a$.

Uniqueness of the symmetric equilibrium follows from the uniqueness of a that makes the bidder indifferent between winning and losing the first auction in the (L, H) case at the price $H-a$. □

Proof of Proposition 2. The key region of the (δ, p, c) parameter space can be characterized by a cost interval $I_{\delta p} = (\underline{c}_p, \bar{c}_{\delta p})$ such that $c \in I_{\delta p} \Rightarrow \delta\Pi_{10} < \Pi_{1A} < \Pi_{10}$. When $p > (9-\sqrt{21})/10$, $I_{\delta p}$ is characterized by inequalities $\Pi_{1A} < \Pi_{10} \Leftrightarrow c > (-5p^2+9p-3)/(3-2p) \equiv \underline{c}_p > 0$ and

$$\delta\Pi_{10} < \Pi_{1A} \Leftrightarrow c < \frac{p^2[(3-2p)(1-\delta)+\delta(-5p^2+9p-3)]}{1-\delta+\delta p^2(3-2p)} \equiv \bar{c}_{\delta p} > \underline{c}_p.$$

Both inequalities follow from Equations (1) and (2). When $p < (9-\sqrt{21})/10$, $\underline{c}_{\delta p} = 0$ because $\Pi_{1A} < \Pi_{10}$ even for $c = 0$, and $\bar{c}_{\delta p}$ is as above, which is a strictly positive quantity whenever $\delta < (3-2p)/(5p^2-11p+6)$. Since $\Pi_{1A} < \Pi_{10} \Rightarrow \Pi_{11} < \Pi_{10}$, $c \in I_{\delta p}$ implies that the commitment seller gets Π_{10}, and ends up better off than the adaptive seller. The $\Pi_{1A} < \Pi_{10} \Rightarrow \Pi_{11} < \Pi_{10}$ claim can be formally shown as follows: $\Pi_{1A} < \Pi_{10} \Leftrightarrow c > (-5p^2+9p-3)/(3-2p) \Leftrightarrow cp^2(3-2p) > p^2(-5p^2+9p-3) \Rightarrow c > p^2(-5p^2+9p-3) \Leftrightarrow \Pi_{11} < \Pi_{10}$. □

References

Besanko, D., W. L. Winston. 1990. Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* **36** 555–567.

Bruce, N., P. Desai, R. Staelin. 2006. Enabling the willing: Consumer rebates for durable goods. *Marketing Sci.* **25**(4) 350–366.

Bulow, J. 1982. Durable goods monopolists. *J. Political Econom.* **90** 314–332.

Caillaud, B., C. Mezzetti. 2004. Equilibrium reserve prices in sequential ascending auctions. *J. Econom. Theory* **117**(1) 78–95.

Coase, R. H. 1972. Durability and monopoly. *J. Law Econom.* **15** 143–149.

Desai, P., D. Purohit. 1998. Leasing and selling: Optimal marketing strategies for a durable goods firm. *Management Sci.* **44**(11) 19–34.

Dhebar, A. 1994. Durable-goods monopolists, rational consumers, and improving products. *Marketing Sci.* **13**(1) 100–120.

Dholakia, U., I. Simonson. 2005. The effect of explicit reference points on consumer choice and online bidding behavior. *Marketing Sci.* **24**(2) 206–217.

Engelbrecht-Wiggans, R. 1994. Sequential auctions of stochastically equivalent objects. *Econom. Lett.* **44** 87–90.

Fay, S. 2004. Partial-repeat-bidding in the name-your-own-price channel. *Marketing Sci.* **23**(3) 407–418.

Freixas, X., R. Guesnerie, J. Tirole. 1985. Planning under incomplete information and the ratchet effect. *Rev. Econom. Stud.* **52**(2) 173–191.

Garratt, R., T. Tröger. 2006. Speculation in standard auctions with resale. *Econometrica* **74**(3) 753–769.

Jeitschko, T. D. 1999. Equilibrium price paths in sequential auctions with stochastic supply. *Econom. Lett.* **64** 67–72.

Jofre-Bonet, M., M. Pesendorfer. 2003. Estimation of a dynamic auction game. *Econometrica* **71** 1443–1489.

Levine, D. K., W. Pesendorfer. 1995. When are agents negligible? *Amer. Econom. Rev.* **85**(5) 1160–1170.

McAfee, R. P., D. Vincent. 1997. Sequentially optimal auctions. *Games Econom. Behav.* **18** 246–276.

Milgrom, P. R., R. J. Weber. 1999. A theory of auctions and competitive bidding II. P. Klemperer, ed. *The Economic Theory of Auctions*. Edward Elgar Publishing, London, UK.

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res.* **6** 58–73.

Shugan, S. 2005. Marketing and designing transaction games. *Marketing Sci.* **24**(4) 525–530.

Skreta, V. 2006. Sequentially optimal mechanisms. *Rev. Econom. Stud.* **73**(4) 1085–1111.

Stokey, N. 1979. Intertemporal price discrimination. *Quart. J. Econom.* **93** 355–371.

Zeithammer, R. 2006. Forward-looking bidding in online auctions. *J. Marketing Res.* **43**(3) 462–476.

Zeithammer, R. 2007. On the role of commitment in sequential auctioning. Working paper, University of California–Los Angeles, CA.