doi: 10.1111/j.1365-2966.2006.0022.x

Auctions with options to re-auction

Simon Grant,* Atsushi Kajii,† Flavio Menezes‡ and Matthew J. Ryan¶

A seller has one unit to sell using an English auction mechanism similar to internet auction markets, such as eBay. Bidders appear according to a random arrival process. The seller chooses a reserve price and duration for each auction. If the reserve is not met, the seller passes in the object and conducts another auction with a new, randomly chosen, set of bidders. We distinguish reserves that embody an institutional commitment not to sell below that price, from those that do not. In each case, we find the optimal reserve price and the optimal auction duration. Without price commitment, the equilibrium reserve is too low for allocative efficiency, whereas the optimal reserve with commitment is shown to be too high when the distribution of bidder valuations exhibits an increasing hazard rate. It might even be socially preferable to allow reserve price commitments. With respect to duration, a version of the Diamond paradox afflicts sellers who cannot commit to price; auctions facilitate valuable duration commitments that increase buyer competition and raise expected revenue. With commitment, price posting (equivalent to a zero-length auction) is the dominant selling mechanism.

Key words auctions, price posting, re-auction, reserve price

JEL classification D44, D82

1 Introduction

Auction theory delivers strong and elegant results about optimal reserve prices in standard auctions. In the independent private values (IPV) context, and under standard regularity

*Department of Economics, Rice University, Houston, Texas, USA and School of Economics, Australian National University, Canberra, Australia

†Institute of Economic Research, Kyoto University, Kyoto, Japan. Email: kajii@kier.kyoto-u.ac.jp

‡Australian Centre of Regulatory Economics, Australian National University, Canberra, Australia and Charles River Asssociates International, Canberra, Australia

¶Department of Economics, University of Auckland, Auckland, New Zealand

We accept responsibility for all remaining shortcomings of the paper. We would like to thank Eric Maskin, Paulo Monteiro, Tomoyuki Nakajima and an anonymous referee, as well as seminar audiences at the Australian National University, Universities of Auckland, Heidelberg, Melbourne, New South Wales, Sydney and Tasmania, ESAM 2002 and the 2004 Zeuthen Workshop (University of Copenhagen) for their comments. Particular thanks are owed to Guillaume Rocheteau, who carefully read earlier drafts and offered many valuable suggestions. The financial support of ARC (Grant no. A000000055) is also gratefully acknowledged. Ryan would like to thank the University of Copenhagen for its invitation to present this paper at the Zeuthen Workshop, and its hospitality during his stay in Denmark. Finally, the viewpoints and opinions expressed in this paper are the views of the authors and are not necessarily those of Charles River Associates (Asia Pacific) (CRA). CRA respects the rights of individuals to express opinions but assumes no responsibility for any errors or omissions contained therein.

assumptions, the seller's optimal reserve price is the solution (in r) to the equation

$$r = v_0 + \frac{1 - F(r)}{f(r)}, \tag{1}$$

where v_0 is the seller's valuation of the object, and F (respectively, f) is the cumulative distribution function (respectively, density function) according to which bidder values are drawn.[1] This solution has several noteworthy features, including the following: (i) the optimal reserve is independent of the number of bidders in the auction; and (ii) it exceeds v_0 and therefore creates an allocative inefficiency.

However, both of these features are artefacts of the institutional assumptions made about reserve prices in theoretical models. In such models, a reserve price is a commitment on the part of the seller to consume the object if bidding fails to reach the reserve. In practice, reserve prices have no such legal force. A seller is free to re-auction the object, or to negotiate with the highest bidder *ex post*. The legal force of a reserve price constrains the agency of the auctioneer, depriving the auctioneer of the authority to commit the seller to a contract below the reserve, or commits the seller to sell to the highest bidder in the event that bidding exceeds the reserve.[2]

Although these characteristics of reserve prices in real-world auctions are well known, and the commitment problem is often mentioned in discussions of auction theory (Krishna 2002; Milgrom 1987), there has been little attempt to study reserve prices under more realistic institutional assumptions.

In the absence of price commitment, it is clear that $r = v_0$ is the only incentive-compatible reserve price in a one-shot auction, and it results in an efficient allocation. This suggests that price commitments are socially undesirable. But this view is overly simplistic. It rests on the assumption that the seller's only outside option is to consume the object. In practice, the seller's outside option is a complex negotiation process with the high bidder, with a further option to re-auction if negotiations break down. The value of this outside option game will not, in general, coincide with v_0.

In the present paper, we re-evaluate optimal reserve prices when re-auctioning is permitted. One important motivation for our set-up is the phenomenal growth in transactions through virtual auctions, such as eBay. In these online markets, the costs of re-auctioning are negligible. The dominant eBay site also allows sellers to choose a minimum bid, which shares important similarities with theoretical reserve prices. Bids below this minimum cannot be submitted, so the seller never learns of the existence of such bidders. This effectively removes the option of *ex post* negotiation.[3] Accordingly, our model also excludes this option.[4]

[1] See, for example, Riley and Samuelson (1981, proposition 3).

[2] In theoretical models, we might be agnostic about whether reserves embody the latter commitment to sell above the reserve, because the presence or absence of this commitment is irrelevant. The seller's optimal reserve price is the same in both cases, and it is always rational for the seller to sell above the optimally chosen reserve.

[3] As with many auction sites, eBay gives bidders the opportunity to contact sellers directly via email. Although it is forbidden by eBay 's rules, a bidder could use this channel to make known their willingness to bid below the minimum. However, such a bidder risks annoying the seller and suffering adverse consequences through the formal and informal reputation systems that are used to maintain the integrity of the virtual marketplace.

[4] The possibility of *ex post* negotiation is the subject of ongoing research. For some initial results, see Menezes and Ryan (2005).

A seller holds a sequence of (ascending price) auctions, with bidders arriving randomly over time. For example, potential buyers might log on to a virtual auction site according to some random process. We allow that there might or might not be a mechanism for making reserve price commitments. In the presence of a commitment technology (such as eBay's minimum bid facility), any reserve price announced at the start of an auction is binding for that auction: if the announced reserve is not met, the seller must either consume the object or conduct another auction. However, the seller is free to set a different reserve price in any subsequent auction.[5] Auctions with reserve price commitments are called *public reserve* auctions in this paper.

If no commitment technology exists, the seller is free to accept or reject the standing bid at the end of an auction. If the seller rejects the standing bid, he or she might consume the object or hold another auction. Because any announcement would be "cheap talk" anyway, we assume that the seller makes no reserve price announcement, and we call these *secret reserve* auctions. The rationale for this terminology is twofold. First, it would be inappropriate to call them "no reserve" auctions, because, in legal parlance, a "no reserve" auction is one in which the seller is obliged to accept the highest bid. Second, even though *ex ante* reserve price commitment is impossible, the seller will still have an *ex post* reservation price: a minimum standing bid that he or she is prepared to accept at the conclusion of the auction. This reservation price is what we refer to as the seller's secret reserve; "secret", because it cannot be credibly announced (although we shall assume that bidders can infer its value).

In a *public reserve* environment with re-auctioning, we obtain a simple modification to (1): v_0 is replaced by the seller's option value of re-auctioning. However, this simple change has significant consequences. First, because the option value depends on the anticipated arrival rate of new bidders, the optimal reserve price depends on the "thickness" of the market, in contrast to (i). Second, the allocative efficiency properties of the optimal public reserve are not as clear-cut as (ii) suggests. Against the tendency of the seller to hike up the reserve to appropriate surplus, must be balanced the seller's relative impatience: because the social planner anticipates the entire surplus on any sale, he or she is prepared to wait longer for a higher value bidder. Reserve price commitments might, in some circumstances, make auctions relatively more efficient than in the absence of such commitments.

Our model also endogenizes the duration of auctions. For example, in Yahoo! Japan, which is by far the dominant internet auction site in Japan, sellers may choose among 2-day to 7-day formats. With a suitably homogeneous bidder arrival process, letting the duration go to zero replicates a price-posting mechanism.[6] If price commitments are disallowed, a version of the Diamond paradox takes effect: price posting yields zero expected return to the seller (Diamond 1971).In this case, sellers will always choose auctions of strictly positive duration. In other words, auctions allow "time commitments" (a commitment to wait a

[5] In our model, like the standard one-shot auction, it does not matter whether or not the seller is obliged to sell above the stated reserve.

[6] For example, in Yahoo! Japan, the item is sold as soon as a bid meets the *Buy It Now* price at any time. We note, however, that there is no explicit Buy It Now price in our model. It is known that if the duration T is positive, there are important strategic differences between these 'Buy It Now' rules. See Milgrom (2004) and analysis of Reynolds and Wooders (2004).

fixed length of time for more buyers to arrive) that are valuable for sellers who cannot credibly commit to price. However, for sellers who can commit to price, we show that the cost of waiting always outweighs the benefits of increased buyer competition: price posting is optimal.

The paper most closely related to the present work is Wang (1993).[7] He too considers a repeated IPV auction structure, with bidders who arrive randomly over time and participate in one auction only. Wang's main focus is the comparison of auctions and price posting, although he also looks at the socially optimal public reserve. Our results advance Wang's analysis in three respects.

First, Wang assumes a Poisson arrival process. We allow for more general arrival processes. The arguments in Wang's paper rely heavily on algebraic manipulations specific to his Poisson model, so the underlying logic is further obscured. We have attempted to provide arguments that illustrate more clearly the intuition for results.

Second, Wang assumes that the seller's waiting penalty consists of a linear storage or display cost, plus a fixed re-auctioning cost. Because exponential discounting is more standard in dynamic economic models, we have adopted this formulation of waiting costs.

Finally, Wang does not analyze the secret reserve (no commitment) case, so these results are new (as is Proposition 6). One of the major motivations for the present paper is to study the private versus social costs and benefits of price commitments.

To conclude this section, we shall mention a few limitations that follow from our assumption of stationary and memoryless bidder behavior. This of course simplifies the problem and enables us to obtain very clean results. But in an internet auction, similar items do tend to be auctioned repeatedly. Therefore, the losers will participate in the next auction and, hence, bidders' behavior will not be time invariant. Some bidders might be interested in obtaining multiple units by participating in many auctions, suggesting that the assumption of independent private values is very strong in this context. These are all interesting but challenging questions, and we shall leave them for future research.

2 The model

A seller has a single object to sell, which he or she values at zero. The seller might hold a potentially infinite sequence of auctions. Any given auction round lasts T units of time,[8] with elapsed time in the current round denoted by t, an element of the interval $[0, T]$, No time elapses between the end of one auction round and the start of the next. In each round, an English auction, modeled as a "button auction" (Milgrom and Weber 1982), is used.[9] Any bidder present at time T in the current round is assumed no longer to be

[7] Nagareda (2003) studies the role of a secret reserve in a static auction, where bidders bid simultaneously and a secret reserve is also chosen simultaneously by a seller whose value is also private information. Our model on the other hand is a dynamic ascending bid auction, and the seller's value is common knowledge.

[8] We shall relax this assumption, by making the T endogenous, in Section 3.5.

[9] Because the bidders arrive at different times, we require a slight modification to the mechanics of the button auction. If the current standing bidder faces no competition, either from a rival bidder, or from the seller's reserve price, he temporarily releases the "button" to stop the price from rising further, but does not exit the auction. If

present at time $t = 0$ of the next. This assumption greatly simplifies the analysis, ensuring a recursive structure to the seller's problem, but is also consistent with stylized features of internet auctions. Automated proxy bidding is common in such markets, so the standing bidder at $t = T$ need not be online. Also, in the event of the object being passed-in, the subsequent auction might be conducted at a different site, with no forwarding address posted for bidders in the previous round.

There are infinitely many potential buyers. Bidders arrive according to a stochastic process to be described below. Once buyers arrive, their values are determined by independent draws from a common cumulative distribution function, F, with strictly positive density, f, on $[0, 1]$. A buyer's own value is private information; but F is common knowledge. We shall assume throughout this paper:

Assumption 1 *The draws determining individuals' valuations are independent of the arrival process.*

Assumption 2 *The probability of n arrivals in any auction round lasting T units of time is given by $p(n, T)$, and is independent of the history of arrivals in previous rounds. The expected number of arrivals for any $T < \infty$ is finite; that is, $\sum_{n=0}^{\infty} np(n, T) < \infty$.*

The first part of Assumption 2 makes the seller's problem stationary. For example, we could fix some stochastic process on $[0, T]$ that is used to determine the number of arrivals. The same process is used in each round. Players arrive according to the stochastic process described above and at the end of each round they vanish and a new auction starts, with bidders arriving according to the same process. Therefore, the distribution of the number of arrivals during a round is independent and identical across rounds. Alternatively, one might imagine a homogeneous Poisson arrival process on $[0, \infty)$ that is chopped into segments of length T. The homogeneity property implies that this is equivalent to restarting the process from time zero each T units of time.

The second part of Assumption 2 is a mild regularity assumption, and is satisfied, for example, by a Poisson process with constant arrival rate $\lambda > 0$.

We assume that bidders play 'myopically' (i.e. that each auction round attracts a different pool of bidders) and, therefore, we ignore buyers' time preferences.[10] The seller, however, has regard for the future beyond the current round, and a continuous rate of time preference $\rho > 0$. The seller values the object at zero.

a new bidder arrives, the original bidder depresses the button again. Once the bidder's valuation is reached, he permanently exits the auction.

[10] The latter assumption is not entirely innocuous. Internet auctions might last for a week or more, and bid retraction is typically forbidden (see http://pages.ebay.com/help/buy/bid-retract.html for eBay's rules on bid retraction). Bidding, therefore, has an opportunity cost that can vary over time. We ignore such costs for analytical convenience. However, they are not without theoretical interest, as they could explain the prevalence of "sniping" in online auctions (Roth and Ockenfels, 2002), even within an IPV set-up. The figure a buyer submits to a proxy bidder at time $t < T$ might understate his willingness to pay if he remains the standing bidder at the conclusion of the auction. A "sniper" with a lower valuation at time T can, therefore, win the auction if the standing bidder is not at his computer when the auction ends.

The seller is under no obligation to accept any bid[11]: the seller always has the option of passing in the object for re-auctioning. However, the seller might or might not be able to commit to reject bids below an announced reserve price. When commitment is not possible, we have a secret reserve auction: no reserve price is announced. The secret reserve is the lowest price the seller will voluntarily accept at $t = T$, given his or her option to re-auction. We use *public reserve* to refer to a reserve price announced at $t = 0$ that is binding: the seller must pass-in the object if bidding fails to reach the public reserve by time T. However, if the object is passed in, the seller is free to set a new public reserve price in the next auction.

3 Results

3.1 Bidding strategies

The optimal bidding strategies are straightforward to determine. It is a weakly dominant strategy to bid up to one's valuation. In a public-reserve auction, with announced reserve r, a bidder with value $v \in [r, 1]$ will submit bids in $[r, v]$ just sufficient to meet the reserve and remain the standing bidder. If bidding exceeds v the bidder drops out of the auction. Bidders with values in $[0, r)$ expect zero surplus from participation in the auction. We assume, without loss of generality, that they choose not to bid.

In a secret-reserve auction, we consider a game where the secret reserve (which we also denote by r) is a choice variable of the seller. The secret reserve is not announced, and so bidders form some conjecture about r, and play a best response against it. In particular, they expect that they do not win the object unless they bid at least r. So for each bidder, the best response to r is the bidding strategy as described above when all the bidders infer r precisely, in particular in equilibrium.[12]

3.2 The option value of re-auctioning

Observe that the seller's problem is stationary: it looks the same at the start of each auction round and can be analyzed using dynamic programming methods. The seller's problem is to choose a reserve price, or acceptance rule, for each auction round. Because of the stationarity of the problem, and because the revenue in each round is bounded, an optimal strategy exists, and it can be written as a stationary policy function. Therefore, we shall focus on strategies in which the same reserve price is applied in each period.

In a secret-reserve scenario, we can think of the seller deciding on his or her acceptance rule at $t = T$: she may accept or reject the standing bid at that time. With a public reserve

[11] As discussed in the introduction, this assumption is innocuous. We could just as well assume that the seller is obliged to sell to the high bidder if bidding exceeds an announced reserve (public or otherwise).

[12] This means that we are implicitly assuming that seller's parameters ρ and the consumption value of the object (zero) are public information. If bidders do not know ρ or the seller's consumption value precisely, then the bidders' and the seller's optimal strategies in a repeated secret-reserve auction game become very complex to determine. This is an interesting avenue for further research.

price, the acceptance rule is set at $t = 0$, and the seller is committed to it. Moreover, the public reserve announcement will affect the bidders' strategies, as described above. In each case, however, the acceptance rule will take the form of a cut-off (reserve) price, r, such that the standing bid at T is accepted if and only if it is at least r.

We can now compute the seller's value function, given T and r. In the following calculations, we assume that r is anticipated by bidders: it is either announced in a public reserve context, or conjectured by bidders in a secret-reserve context. Let $\Pi(r, T)$ be the probability that an auction round with reserve r ends with the object being passed in; and let $\mathcal{R}(r, T)$ be the expected price obtained in a T -length auction with reserve price r, conditional on a sale being achieved. Assumption 1 allows us to write:

$$\Pi(r, T) = \sum_{n=0}^{\infty} p(n, T) F(r)^n$$

and

$$\mathcal{R}(r, T) = \left[\frac{1}{1 - \Pi(r, T)}\right] \sum_{n=1}^{\infty} p(n, T)[1 - F(r)^n]\mathcal{R}_n(r),$$

where $\mathcal{R}_n(r)$ is the expected revenue from a one-shot auction with n bidders and reserve r, conditional on sale. Familiar manipulations show that:[13]

$$\mathcal{R}_n(r) = \left[\frac{1}{1 - F(r)^n}\right] \int_r^1 J(z)\, dF(z)^n$$

where

$$J(z) = z - \frac{[1 - F(z)]}{f(z)}.$$

Let $v_0(r, T)$ denote the discounted expected revenue from holding repeated T-length auctions, with reserve price r in each round, until a sale is achieved. This value function is defined by the recursive formula

$$v_0(r, T) = e^{-\rho T} \{[1 - \Pi(r, T)]\, \mathcal{R}(r, T) + \Pi(r, T) v_0(r, T)\}.$$

This can be solved to obtain:

$$v_0(r, T) = D(r, T)\mathcal{R}(r, T), \tag{2}$$

where

$$D(r, T) = \frac{1 - \Pi(r, T)}{e^{\rho T} - \Pi(r, T)}$$

represents a discount factor reflecting the expected delay until sale. The following lemma, whose proof is in the appendix, will be important in what follows:

[13] See, for example, Menezes and Monteiro (2004).

Lemma 1 *The value function $v_0(r, T)$ is continuous in r.*

3.3 The seller's optimal reserve price

Recall that the secret reserve is set at the end of the auction, given existing bids, and, therefore, the seller is never interested in randomizing the acceptance price because it does not affect bidders' behavior. Therefore, writing r for the secret reserve, it is then straightforward to show that, in a secret reserve auction in which bidders form rational expectations, the seller's secret reserve is the solution to

$$r = v_0(r, T). \tag{3}$$

Equation (3) says that the seller's acceptance price at $t = T$ is equal to the value of re-starting the auction game. If bidding fails to reach $v_0(r, T)$, it is better to pass in the object; and accepting a bid above $v_0(r, T)$ is preferable to re-auctioning. Of course, because $v_0(r, T)$ depends on r, the solution to (3) might be complex to determine. However:

Proposition 2 *There exists an equilibrium secret reserve, and all solutions to (3) are interior to* $[0, 1]$.

PROOF: An equilibrium secret reserve, if one exists, must lie in the unit interval, and $v_0(r, T) \in [0, 1]$ for any such r. By Lemma 1 and Brouwer's fixed point theorem, it follows that (3) has a solution in $[0, 1]$. Moreover, it is clear that $v_0(0, T) > 0$ and $v_0(1, T) < 1$ so any solution is interior to $[0, 1]$. □

In a public reserve auction, the optimal reserve price is the solution to

$$\max_{r \in [0,1]} v_0(r, T). \tag{4}$$

Again, continuity of $v_0(r, T)$ in r ensures the existence of a solution.

The optimal public reserve will be higher than the equilibrium secret reserve price under the following regularity assumption.

Assumption 3 *The function*

$$J(r) = r - \frac{[1 - F(r)]}{f(r)}$$

is non-decreasing in r.

More precisely:

Proposition 3 *Let r^* denote an optimal public reserve and r^{**} an equilibrium secret reserve price. Suppose Assumption 3 holds. Then*

$$v_0(r^*, T) = r^* - \frac{[1 - F(r^*)]}{f(r^*)} \tag{5}$$

and, hence, $r^* > r^{**}$. Furthermore, $r^* \in (0, 1)$.

PROOF: Let $\hat{v}(T)$ denote $v_0(r, T)$ evaluated at an optimal public reserve price. Then

$$e^{\rho T}\hat{v}(T) = \max_r\{[1 - \Pi(r, T)]\mathcal{R}(r, T) + \Pi(r, T)\hat{v}(T)\}$$
$$= \max_r \left\{\sum_{n=1}^{\infty} p(n, T)\left[\int_r^1 J(z)\, dF(z)^n + F(r)^n\hat{v}(T)\right]\right\} + p(0, T)\hat{v}(T) \tag{6}$$

But

$$\max_r \int_r^1 J(z)dF(z)^n + F(r)^n\hat{v}(T)$$

is the problem of choosing an optimal public reserve in a one-shot (first-price) auction with n bidders and seller valuation $\hat{v}(T)$. For *any* n, the first order condition for this problem is

$$J(r) = \hat{v}(T) \tag{7}$$

and Assumption 3 ensures that the problem is concave. Hence, the solution to (7) solves problem (6) also. Then (5) follows from (7).

From (3), (4) and (5):

$$r^{**} = v_0(r^{**}, T) \leq v_0(r^*, T) \leq r^*.$$

It follows that $r^* \geq r^{**}$, with equality if and only if $1 - F(r^*) = 0$. But the latter is equivalent to $r^* = 1$, and Proposition 2 rules out $r^{**} = 1$. Therefore, $r^{**} < r^* < 1$. Finally, $r^* > r^{**}$ and Proposition 2 imply $r^* > 0$. □

Therefore, the optimal public reserve satisfies

$$r = v_0(r, T) + \frac{[1 - F(r)]}{f(r)}. \tag{8}$$

This matches (1), except that the seller's consumption value of the object is replaced by the 'option value' of retaining the object for re-auctioning. Because $v_0(r, T) > 0$, the optimal public reserve is higher than in the one-shot case, as a result of the seller's opportunity cost of not trading in any given round being greater than the consumption value of the object.

Once again, (8) might be complex to solve because of the dependence of $v_0(r, T)$ on r. Moreover, because $v_0(r, T)$ also depends on the details of the bidder arrival process, the equilibrium secret reserve and optimal public reserve prices will both depend on the "thickness" of the market, unlike in the one-shot case. For example, if bidders arrive according to a Poisson process with arrival rate λ, the secret and public reserves will both be (increasing) functions of λ.

3.4 The efficient reserve price

In a one-shot auction, any (credible) reserve price above the seller's valuation is surplus reducing, because there is some chance that the object does not end up in the hands of

the person who values it most highly. Hence, reserve price commitments are allocatively inefficient, whereas the equilibrium 'secret reserve' is socially optimal.

With a repeated auction, this is no longer the case. Let us consider a social planner who must respect all the exogenous constraints of the auction mechanism. The social planner discounts at the same rate as the seller. At the end of any round, the planner might allocate the object to any one of the arrivals during that round, or else wait another T units of time to collect a new sample of buyers. The planner can only allocate to a buyer who arrived in the current round. We assume that the planner observes the values of arriving buyers, but cannot foresee the values of future arrivals.[14] The planner's objective is to maximize discounted expected total surplus.

As usual, stationarity implies that the planner's optimal allocation rule consists of a cut-off r, such that the good is allocated to the highest value arrival during the current period if and only if this highest value exceeds r. Note, therefore, that the social planner's cut-off is a valuation, whereas the seller's is a bid level. However, in each case, the good is allocated in the current round if and only if a bidder arrives whose valuation exceeds the relevant cut-off, and it is allocated to the highest value arrival in that round. Therefore, the seller's reserve price is directly comparable with the planner's value cut-off for the purposes of assessing allocative efficiency.

Let $v_S(r, T)$ denote the planner's value function, when r is the cut-off level for maximum buyer value. It is straightforward to see that

$$v_S(r, T) = D(r, T)\mathcal{S}(r, T), \tag{9}$$

where

$$\mathcal{S}(r, T) = \left[\frac{1}{1 - \Pi(r, T)}\right] \sum_{n=1}^{\infty} p(n, T) \int_r^1 z\, dF(z)^n$$

is the expected total surplus conditional on allocating the object. The planner will choose r to maximize $v_S(r, T)$, and it is also clearly optimal to allocate the object if and only if the highest value arrival in the current round exceeds the Planner's option value of dismissing the current buyers and re-starting the process. Therefore:

$$\hat{r} \in \arg\max_{r \in [0,1]} v_S(r, T) \Longrightarrow v_S(\hat{r}, T) = \hat{r}. \tag{10}$$

We immediately conclude:

Proposition 4 *The secret-reserve auction sells the object too quickly from a social efficiency point of view. That is, if r^{**} and $\hat{r}$ solve (3) and (10) respectively, then $r^{**} < \hat{r}$.*

PROOF: It is clear that $v_0(r^{**}, T) < v_S(r^{**}, T)$, because the seller must share the (strictly positive) expected surplus with the buyer, whereas the social planner does not. Therefore,

[14] This seems to us the natural benchmark against which to assess the allocative efficiency of the auction mechanism. Alternatively, one might justify the lack of social planner omniscience by supposing that bidders do not receive their values for the object until they arrive. This value might depend, for example, on actions taken prior to arrival.

if $\hat{r} \leq r^{**}$ we have

$$v_S(\hat{r}, T) = \hat{r} \leq r^{**} = v_0(r^{**}, T) < v_S(r^{**}, T).$$

But this contradicts the fact that $\hat{r}$ maximizes $v_S(r, T)$. □

The intuition behind Proposition 4 is easy to see. Consider an auction round that ends with a high bid equal to r, which in turn is equal to the high bidder's valuation. Accepting the current high bid earns r for both the seller and the social planner, whereas passing in earns the former $v_0(r, T)$ and the latter $v_S(r, T)$. Because $v_0(r, T) < v_S(r, T)$, the social planner is more patient. Unlike the one-shot case, removing the seller's access to reserve price commitments does not produce an efficient allocation mechanism.

Is the optimal public reserve too high or too low from an efficiency point of view? The answer depends on two opposing tendencies: the seller's relative impatience versus his or her use of reserve price commitment to extract surplus from the buyer. However, the seller's optimal public reserve is too high in the following special case:

Proposition 5 *If*

$$\frac{1 - F(z)}{f(z)} \tag{11}$$

is decreasing, the socially optimal r is less than the privately optimal public reserve.

PROOF: Comparing (2) and (9):

$$v_S(r, T) - v_0(r, T) = D(r, T)[\mathcal{S}(r, T) - \mathcal{R}(r, T)]. \tag{12}$$

First, we claim that

$$D(r, T) = \frac{1 - \Pi(r, T)}{e^{\rho T} - \Pi(r, T)}$$

is strictly decreasing in r for any $T > 0$. It is obvious that $\Pi(r, T)$ (the probability that the maximum value arrival in any given T-length round has value less than r) is increasing in r. Because $\varepsilon^{\rho T} > 1$ for any $T > 0$, the claim follows.

Next, observe that

$$\mathcal{S}(r, T) - \mathcal{R}(r, T) = \int_r^1 \left[\frac{1 - F(z)}{f(z)}\right] \mu(z; r, T)\, dz, \tag{13}$$

where

$$\mu(z; r, T)\, dz = \frac{\sum_{n=1}^{\infty} p(n, T) dF(z)^n}{1 - \Pi(r, T)} = \frac{\sum_{n=1}^{\infty} p(n, T) n f(z) F(z)^{n-1}}{1 - \Pi(r, T)} dz.$$

The function $\mu(z; r, T)$ is the density function for the maximum bidder value in a T-length auction, conditional on this maximum exceeding r. Therefore, (13) is the expected value of (11) over $[r, 1]$ with respect to the conditional density $\mu(z; r, T)$. If (11) is decreasing, it is clear that this expected value must decrease with increases in r.

Therefore, we have established that (12) is decreasing in r. Furthermore, if $J(z)$ is increasing (which is implied by (11) decreasing) there is a unique interior maximum of $v_0(r, T)$ in r, and this is equal to the global maximum. From this fact and that (12) is decreasing in r, it follows that the socially optimal reserve is less than the privately optimal one.[15] □

Because there appears to be some confusion on this point in the published literature,[16] we observe here that Assumption 3 is inconsistent with (11) being everywhere non-decreasing. The following result is proved in the appendix.

Proposition 6 *There does not exist a distribution such that both (11) and $J(z)$ are well-defined, finite and non-decreasing on* (0,1).

Combining Propositions 4 and 5 establishes that when (11) is monotonically decreasing, the efficient reserve price lies between the seller's optimal secret reserve and public reserve. Therefore, it is of interest to determine which auction format comes closest to achieving the socially optimal expected surplus. In particular, are reserve price commitments ever socially desirable?

For the special case in which $F(z) = z$ (values are uniformly distributed on $[0, 1]$), $T = 1$ and arrivals follow a Poisson process with arrival rate λ, Figure 1 provides an affirmative answer. Note that, in this case,

$$\frac{1 - F(z)}{f(z)} = 1 - z$$

and $J(z) = 2z - 1$, so Assumption 3 is satisfied and (11) is decreasing. Figure 1 plots the difference between the social planner's value function evaluated at the seller's optimal secret reserve (v_S^{sr}) and at the seller's optimal public reserve (v_S^{pr}), expressed as a percentage of the latter. This quantity is represented as a function of λ and the per period discount factor $\delta = e^{-\rho}$. For a sufficiently patient seller operating in a sufficiently thin market, reserve price commitments improve efficiency. [17]

3.5 Optimal auction length

Until now, we have fixed the length of auctions at T units of time. However, online auctions often allow sellers to choose their duration, at least within some bounds. Therefore, it is of interest to endogenize T. We shall allow $T = 0$ by defining $v_0(r, 0) := \lim_{T \downarrow 0} v_0(r, T)$.

[15] Note that (13) is also decreasing in T as follows. If $T' > T$, then

$$G(x; r, T') = \int_r^x \mu(z; r, T')\, dz$$

first-order stochastically dominates $G(x; r, T)$. Because (11) is decreasing, (13) must be decreasing in T. This allows us to generalize an observation in Wang (1993, theorem 3). If waiting costs are linear in T and there is a fixed re-auctioning cost, $\mathcal{S}(r, T) - \mathcal{R}(r, T)$ is the difference between the planner's and the seller's value functions. Therefore, the socially optimal value for T (the optimal expected 'sample size' when sampling buyers) is lower than the privately optimal T

[16] For example, theorem 3 in Wang (1996) proves a result under the (untenable) assumption that both $J(z)$ and (11) are increasing.

[17] For the algebraic details, see Grant *et al.* (2002).

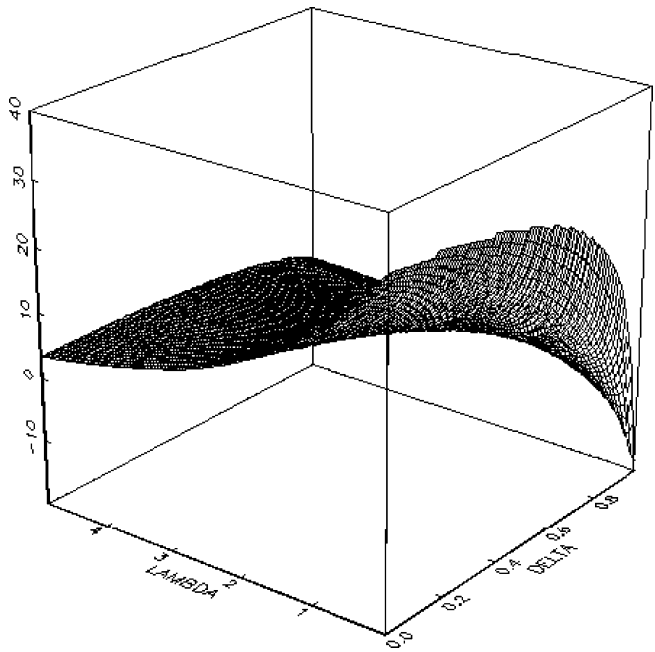

Figure 1 Social value of commitment: $100[v_S^{sr} - v_S^{pr}]/v_S^{pr}$.

We now ask the question: *Which value of $T \in [0, \infty)$ should the seller choose to maximize her expected discounted revenue?* To provide an answer, we make the following additional assumption on the behavior of the arrival process as T vanishes.

Assumption 4 *For each n, the function $p(n, T)$ is continuous in T. Moreover, the following two regularity conditions are satisfied:*

$$\lim_{T \downarrow 0} \frac{T}{1 - p(0, T)} > 0 \tag{14}$$

and

$$\lim_{k \to \infty} \frac{\sum_{n=2}^{\infty} p(n, T_k)}{p(1, T_k)} = 0 \tag{15}$$

for any decreasing sequence $\{T_k\}_{k=1}^{\infty}$ with $\lim_{k \to \infty} T_k = 0$.

Condition (14) ensures that, at any reserve price, the expected elapsed time until a sale is achieved remains bounded away from zero as T vanishes. Condition (15) ensures that, as $T \downarrow 0$, the number of active bidders converges to unity conditional on a sale being achieved. In particular, this assumption rules out multiple simultaneous arrivals. A sufficient condition for (15) is "uniform orderliness" (Synder 1975).[18] Assumption 4

[18] The uniform orderliness condition in our model can be stated as follows. Write $p(n, s, \zeta)$ for the probability of n arrivals in the interval $[s, s + \zeta]$. Then for any $\varepsilon > 0$ there exists a $\zeta > 0$ such that $\sum_{n=2}^{\infty} p(n, s, \zeta) \leq \varepsilon\, p(1, s, \zeta)$, for all $s \in (0, \infty)$.

delivers a version of the Diamond paradox (Diamond 1971): seller surplus vanishes as $T \downarrow 0$. The intuition is straightforward. Condition (15) implies that the expected sale price converges to r, whereas the expected delay remains bounded away from zero because of (14). Therefore, the option value is strictly less than the reserve price unless $r = 0$.

Proposition 7 *Under Assumptions 1–4, $T = 0$ is never optimal in a secret-reserve auction.*

PROOF: Let $r > 0$ be given. We shall show that

$$\lim_{T \downarrow 0} v_0(r, T) < r. \tag{16}$$

Recall that $1 - \Pi(r, T)$ is the probability of sale in any T-length round with reserve r. Because $1 - \Pi(r, T) \leq 1 - p(0, T)$, (14) implies

$$\lim_{T \downarrow 0} [1 - \Pi(r, T)] = 0.$$

If there is some $\hat{T} > 0$ such that $1 - \Pi(r, T) = 0$ for all $T < \hat{T}$, then (16) follows immediately.

Otherwise, there must exist an infinite sequence $\{T_k\}_{k=1}^{\infty}$ such that $T_k \downarrow 0$ as $k \to \infty$ and $1 - \Pi(r, T_k) > 0$ for all k. We first show that

$$\lim_{k \to \infty} D(r, T_k) < 1. \tag{17}$$

Write

$$D(r, T_k) = \frac{\left[\frac{1-\Pi(r, T_k)}{T_k}\right]}{\left(\frac{e^{\rho T_k}-1}{T_k}\right) + \left[\frac{1-\Pi(r, T_k)}{T_k}\right]}$$

and observe that

$$\lim_{k \to \infty} \frac{e^{\rho T_k} - 1}{T_k} = \rho > 0.$$

Using (14):

$$\lim_{k \to \infty} \frac{1 - \Pi(r, T_k)}{T_k} \leq \lim_{k \to \infty} \frac{1 - \Pi(0, T_k)}{T_k} = \lim_{k \to \infty} \frac{1 - p(0, T_k)}{T_k} < \infty$$

and (17) follows.

We next show that

$$\lim_{k \to \infty} \mathcal{R}(r, T_k) \leq r. \tag{18}$$

Recall that $\mathcal{R}(r, T)$ is the price obtained in a T-length round with reserve r, conditional on making a sale in that round. Therefore:

$$\begin{aligned}\mathcal{R}(r, T_k) &\leq \frac{p(1, T_k)[1 - F(r)]}{1 - \Pi(r, T_k)} r + \frac{\sum_{n=2}^{\infty} p(n, T_k)}{1 - \Pi(r, T_k)} \\ &\leq r + \frac{\sum_{n=2}^{\infty} p(n, T_k)}{p(1, T_k)[1 - F(r)]}.\end{aligned}$$

By (15), the second term vanishes, and so (18) follows.

Combining (17) and (18) gives

$$\lim_{k \to \infty} v_0(r, T_k) < r$$

so (16) is proved.

Because $v_0(r, 0) \geq 0$ for all r, we deduce from (16) that $v_0(r, 0) = r$ if and only if $r = 0$. This completes the proof. □

When sellers cannot commit to his or her posted price, and he or she faces a 'search cost' of finding another buyer (the expected delay until the next arrival), the only incentive compatible reserve price is zero, and bidders know this. Absence of commitment deprives the seller of any effective market power. To acquire surplus from trade, he or she must hold an auction long enough that there is some non-zero probability of buyer competition.

However, for a public-reserve auction, matters are not so clear. The seller does not need to use a commitment to delay the transaction to put pressure on buyers, as he or she can use price commitments instead. Of course, longer rounds still increase buyer competition and, hence, seller surplus, but they also delay trade. It is not clear a priori which T will provide an optimal balance between these two effects.

Nevertheless, one can show that $T = 0$ is optimal under the following additional assumption:

Assumption 5 *The bidder arrival process on* $[0, \infty)$ *has the property that the numbers of arrivals in disjoint intervals are independent, and the distribution of the number of arrivals in* $(s, t]$ *depends only on* $t - s$.

Assumption 5 is a homogeneity (or stationarity) property. Arrival processes satisfying Assumption 5 might be stopped and re-started at arbitrary points in time, while remaining identical to the process run uninterrupted over $[0, \infty)$. For such a process, letting $T \downarrow 0$ in a public-reserve auction is equivalent to price-posting.

In fact, the only arrival process satisfying Assumptions 1–5 is the Poisson (Brown 1984, theorem 3). However, our proof of Proposition 8 invokes Assumptions 1–5 rather than the algebraic form of the Poisson. This allows the underlying logic to be better understood.

Proposition 8 *Under Assumptions 1–5,* $T = 0$ *is optimal in a public-reserve-price auction.*

A detailed proof can be found in the appendix, but we shall here sketch the basic idea. First, consider the arrival process on $[0, \infty)$. Suppose that the seller posts a public reserve price of r, and bidders arrive and bid as in our usual ascending auction game. Hence, the

standing bid evolves according to a well-defined stochastic process that depends on the value of r. Let V_r denote this stochastic process (see the proof of the next lemma in the appendix for a formal definition). Consider a seller who must choose when to stop this process and accept the current standing bid. Under our assumptions, this 'optimal stopping' problem has a simple solution:

Lemma 9 *There exists a p (which might depend on r) such that it is optimal to stop as soon as* $V_r \geq p$.

In other words, given r, the seller stops the process as soon as the cut-off price is met. A proof of this result can also be found in the appendix.[19]

Of course, the seller should also choose r in an optimal fashion. We claim that if $(\hat{r},\ \hat{p})$ are jointly optimal, then so is $(r = \hat{p},\ \hat{p})$. To see why, consider the possibility that $\hat{r} < \hat{p}$.[20] Then the object sells to the second arrival whose value weakly exceeds $\hat{p}$, at a price equal to the minimum of the values of the first two such arrivals. If instead, the seller had chosen $r = \hat{p}$, the object would have sold to the first arrival whose value weakly exceeds $\hat{p}$, at a price equal to $\hat{p}$. But if $\hat{p}$ is optimally chosen (given $\hat{r}$), as we assume, then the seller receives no less expected discounted revenue under the latter scenario than under the former. In particular, it is necessary that $\hat{p}$ satisfy the following condition: given $V_{\hat{r}} = \hat{p}$, it must be weakly preferable to accept $\hat{p}$ immediately than to wait for next increment in $V_{\hat{r}}$ and accept at that point.

It follows that there exists an optimal pair $(r,\ p)$ with the property that $r = p$, which amounts to a price-posting mechanism. In our repeated auction scenario, the seller does not choose $(r,\ p)$, but $(r,\ T)$. The latter corresponds to a fixed-time stopping rule: stop the process at the first integer multiple of T for which $V_r \geq r$. The optimal $(r,\ p)$ rule must weakly dominate the best choice of $(r,\ T)$. However, because we have just seen that price-posting (i.e. $r = p$) is optimal when the seller can choose $(r,\ p)$, and because it isequivalent to the $(r,\ T)$ rule with $T = 0$, Proposition 8 follows.

To obtain some intuition for Proposition 8, observe that, in our model, a decision to re-auction dismisses the standing bidder (if any) and excludes him from receiving the object. The ability to commit to a reserve price therefore places all the bargaining power on the side of the seller: it is a credible take-it-or-leave-it offer to the buyer. There is no advantage to the seller from waiting for the random arrival process to place competitive pressure on the standing bidder, as the seller can place whatever pressure he she likes through his or her reserve price. Waiting merely adds delay costs.

Propositions 7 and 8, taken together, provide some useful perspective on the price mechanism. Posted price selling will be favored by sellers who can credibly commit to prices above their option value of retaining the object. In the absence of a suitable commitment technology, auctions become useful. They allow sellers to credibly commit to a costly waiting period before transacting. Waiting gives an opportunity for buyer competition to push the price above the seller's option value.

[19] In a similar vein, McAfee and McMillan (1988, theorems 9 and 10) prove that an optimal procurement mechanism for a monopsonist who pays a linear communication cost is to approach buyers sequentially with a fixed-price offer until it is accepted. (See also Riley and Zeckhauser 1983).

[20] It is obvious that $\hat{r} > \hat{p}$ can be ruled out.

Consider a Poisson arrival process with arrival rate λ. This satisfies all of our assumptions, so $T = 0$ is optimal. Moreover:[21]

$$v_0(r,\ 0) = D(r, 0)r,$$

where

$$D(r, 0) = \frac{\lambda[1 - F(r)]}{\rho + \lambda[1 - F(r)]}.$$

One can think of $D(r, 0)$ as a demand curve, so $v_0(r,\ 0)$ is analogous to the profit function of a monopolist with zero costs. For example, if we further assume that bidder values are uniform on $[0, 1]$, straightforward calculations reveal that $v_0(r, 0)$ is maximized when

$$r\ =\ (1 + \theta)\ -\ \sqrt{\theta(1 + \theta)}$$

(where $\theta = \rho/\lambda$).

Sellers on eBay have long had the option of declaring a *Buy It Now* price. In the absence of a standing bid above the reserve, an arriving bidder who offers to pay the Buy It Now price is immediately awarded the object and the auction cancelled. A seller on eBay could, therefore, implement a posted-price mechanism by choosing her minimum bid amount equal to the Buy It Now price.[22] More recently, and in response to popular demand, eBay has introduced a standard fixed price sale format.

4 Concluding remarks

The model presented here offers a natural generalization of the one-shot auction that allows for re-auctioning.[23] This generalized framework is useful in several of respects.

First and foremost, it will be rare that sellers can credibly commit not to re-auction the object. In particular, internet auction sites allow rapid re-auctioning at negligible cost.

Second, the potential to re-auction allows for a more realistic analysis of reserve prices. The option value of retaining the object for re-auctioning exceeds the consumption value of the object to the seller, so the optimal 'secret' and 'public' reserves both increase. These optimal reserve prices also depend on the thickness of the auction markets.

Third, re-auctioning significantly alters the welfare properties of different auction mechanisms. Secret reserves are too low for allocative efficiency; and reserve price commitments might be socially preferable.

Finally, in the absence of price commitments, the logic of the Diamond paradox implies that auctions facilitate a duration commitment ($T > 0$) that sellers will use to increase buyer

[21] See Grant *et al.* (2002) for details. A similar expression is derived in Wang (1995), who shows that $v_0(r, 0)$ is concave in r under Assumption 3.

[22] In the theoretical literature, the term "posted price" usually implies, as we assume here, an institutional commitment to reject lower offers (see e.g. Wang. 1993 and 1995). In practice, however, posted prices do not embody such commitments, and are more correctly interpreted as ceiling prices (Chen and Rosenthal 1996). This is why the Buy It Now facility must be combined with a minimum bid to convert it to a posted price mechanism in our sense.

[23] As $\rho \to \infty$ the familiar one-shot results are recovered.

competition. Furthermore, if bidders arrive according to a strongly stationary process, our framework nests price-posting as a limiting case ($T \to 0$). When sellers have access to price commitments, price-posting is preferable to auctioning under strong stationarity of the arrival process.

Appendix

Proof of Lemma 1: We first show that $\Pi(r, T)$ is continuous in r. Fix any r and $\varepsilon > 0$. By the continuity of F, there is $\delta > 0$ such that $|r - r'| < \delta$ implies that $|[F(r)] - [F(r')]| < \varepsilon$. Now

$$[F(r)]^n - [F(r')]^n = \\ ([F(r)] - [F(r')])(F(r)^{n-1} + F(r)^{n-2}F(r') + \cdots + F(r')^{n-1}),$$

and $0 \leq (F(r)^{n-1} + F(r)^{n-2}F(r') + \cdots + F(r')^{n-1}) \leq n$ because $F(r)$ is a cumulative distribution. So $|r - r'| < \delta$ implies $|[F(r)]^n - [F(r')]^n| < \varepsilon n$ for all n. Now

$$\begin{aligned}\Pi(r, T) - \Pi(r', T) &= \sum_{n=1}^{\infty} p(n, T)F(r)^n - \sum_{n=1}^{\infty} p(n, T)F(r')^n \\ &= \lim_{k\to\infty} \sum_{n=1}^{k} p(n, T)F(r)^n - \lim_{k\to\infty} \sum_{n=1}^{k} p(n, T)F(r')^n \\ &= \lim_{k} \sum_{n=1}^{k} p(n, T)(F(r)^n - F(r')^n).\end{aligned}$$

So $|r - r'| < \delta$ implies that

$$\begin{aligned}\left|\lim_{k} \sum_{n=1}^{k} p(n, T)(F(r)^n - F(r')^n)\right| &\leq \lim_{k} \sum_{n=1}^{k} p(n, T)|(F(r)^n - F(r')^n)| \\ &< \lim_{k} \sum_{n=1}^{k} p(n, T)n\varepsilon \\ &= \varepsilon \lim_{k} \sum_{n=1}^{k} p(n, T)n,\end{aligned}$$

and by assumption $\lim_k \sum_{n=1}^{k} p(n, T)n$ is a finite constant. This shows that $\Pi(r, T)$ is continuous in r.

Therefore, $v_0(r, T)$ is continuous in r, because

$$\int_r^1 J(z)dF(z)^n$$

is continuous and $e^{\rho T} - \Pi(r, T) > 0$ for any $r \in [0, 1]$ and any $T > 0$.

This completes the proof. □

Proof of Proposition 6: Suppose that

$$h(z) = \frac{1 - F(z)}{f(z)}$$

and $J(z) = z - h(z)$ are well-defined, finite and non-decreasing on $(0, 1)$. Because

$$\frac{d}{dz} \ln\left(\frac{1}{1 - F(z)}\right) = \frac{1}{h(z)}$$

we see that

$$\ln\left(\frac{1}{1 - F(x)}\right) = \int_0^x \frac{1}{h(z)}\, dz.$$

Therefore:

$$F(x) = 1 - \exp\left[-\int_0^x \frac{1}{h(z)}\, dz\right]. \tag{19}$$

In order that $F(1) = 1$, we must have

$$\int_0^1 \frac{1}{h(z)}\, dz = \infty.$$

Because h is non-decreasing, this in turn requires $h(0) = 0$. This fact, and the assumption that $J(z)$ is non-decreasing imply $h(z) \leq z$ on $(0, 1)$. Therefore,

$$\int_0^x \frac{1}{h(z)}\, dz \geq \int_0^x \frac{1}{z}\, dz = \infty$$

for any $x > 0$. But then $F(x) = 1$ for all $x > 0$ from (19). This means that the distribution has an atom at $z = 0$, which is inconsistent with existence of a density function. □

Proof of Lemma 9: Consider an arrival process on $[0, \infty)$ satisfying Assumptions 1–5. We first introduce some notation, which will also be used in the proof of Proposition 8 later, and then define V_r formally.

Denote by $V^{(k)}$ the stochastic process describing the kth highest value, with sample path $V^{(k)}(t)$, $t \geq 0$: $V^{(k)}(t)$ is the kth highest value arrival up to time t. Because arrivals are discrete, a sample path of process $V^{(k)}$ is a non-decreasing step function of t. For $z \in [0, 1)$, denote by $\tau_z^{(k)}$ the first time at which $V^{(k)} \geq z$. Therefore, by definition, $V^{(k)}(\tau_z^{(k)}) \geq z$ with probability one. In fact, because arrivals are discrete and the distribution of bidder values is continuous, $V^{(k)}(\tau_z^{(k)}) > z$ holds with probability one.

Consider an English auction mechanism with reserve r. By definition V_r is the stochastic process describing the high bid value as follows: at time t, if $\tau_r^{(1)} > t$ (no bidder with value above the reserve has arrived by t) then $V_r(t) = 0$; $V_r(t) = r$ if $\tau_r^{(2)} > t \geq \tau_r^{(1)}$ (exactly one bidder with value above the reserve has arrived by t); and $V_r(t) = V^{(2)}(t)$ if $t \geq \tau_r^{(2)}$.

Now we shall give a proof. For the purposes of argument, suppose that the seller observes not only the value of V_r, but also the arrival of bidders. Therefore, even if a bidder arrives whose value is less than V_r and, hence, does not bid, the seller is assumed to observe his arrival.[24] We shall show that the optimal stopping rule is as stated in the Lemma under these assumptions. *A fortiori*, the Lemma holds when the seller observes only the value of V_r.

Consider the situation where $V_r = x \geq r$ is observed, and so there is exactly one standing (unobserved) bidder whose value exceeds x. Let $v(x)$ denote the expected discounted continuation value of the optimally stopped process. Consider the cumulative distribution function F_x conditional on $V_r = x$ for the value of V_r right after the next arrival. With probability $F(x)$, the next arrival is no larger than x. In this case no bid is made and V_r remains x, so $F_x(z) = 0$ if $z < x$, and $F_x(x) = F(x)$. For $z > x$, V_r will be above z after the next arrival if and only if both standing bid and the next arrival exceeds z. The next draw exceeds z with probability $1 - F(z)$ and

[24] Note that this is a version of a standard search problem for a seller facing unknown demand (e.g. job search). In our version of the problem, all prices previously found remain available to the seller.

the standing bidder's value is also over z with probability $\frac{1-F(z)}{1-F(x)}$ and these are independent. Therefore:

$$F_x(z) = \begin{cases} 1 - \frac{[1-F(z)]^2}{[1-F(x)]} & \text{if } z \geq x \\ 0 & \text{otherwise.} \end{cases}$$

When $z > x$, write $f_x(z) = \frac{d}{dz} F_x(z)$ for the associated density function.

Because V_r is a Markov process, the optimal stopping rule is independent of time. If $V_r(t) = x$ and the process is stopped at t, the seller receives x. If the seller waits until the next arrival, he or she will receive the expected value of $v(z)$ with respect to F_x at that time. By Assumptions 1–5, the expected waiting cost is a constant $\beta \in (0, 1)$ and v solves the Bellman equation:

$$v(x) = \max\left\{x, \beta\left[F(x)v(x) + \int_x^1 f_x(z)v(z)\,dz\right]\right\}. \tag{20}$$

We claim that there is a p such that $v(x) > x$ if $x < p$, and $v(x) = x$ if $x \geq p$. This value function can be realized by stopping when $V_r \geq p$, so the Lemma is proved once we verify our claim.

Let T be the operator defined as follows: for any $w \in C([0, 1])$,[25]

$$Tw(x) = \max\left\{x, \beta\left[F(x)w(x) + \int_x^1 f_x(z)w(z)\,dz\}\right]\right\}.$$

Note that T is monotone: $Tw \geq Tw'$ if $w \geq w'$. Furthermore, for any constant a, $T(w + a)(x) \leq T(w)(x) + \beta a$ with $\beta < 1$, so Blackwell's sufficient conditions for a contraction are satisfied (Stokey and Lucas 1989, theorem 3.3). Hence, T is a contraction mapping and it has a unique fixed point, which can be obtained as $\lim_{n\to\infty} T^n(w_0)$, where w_0 is an arbitrary continuous function from [0, 1] to itself. Any fixed point of T satisfies the Bellman equation (20).

Set $w_0(x) = x$ and investigate $\lim_{n\to\infty} T^n(w_0)$. Assume for now that the equation

$$x - \beta\left[F(x)x + \int_x^1 f_x(z)z\,dz\right] = 0 \tag{21}$$

has a unique solution, and let $x = p$ denote this solution. It is clear there is at least one solution, because the left-hand side is negative at $x = 0$ and positive at $x = 1$. We will verify the uniqueness part of the assumption shortly.

If $x \geq p$, then $x - \beta[F(x)x + \int_x^1 f_x(z)z\,dz] \geq 0$ and, hence, $Tw_0(x) = x$. Iterating, we obtain that $\lim_{n\to\infty} T^n(w_0)(x) = x$ in this case. If $x < p$, then $x - \beta[F(x)x + \int_x^1 f_x(z)z\,dz] < 0$ and so $Tw_0(x) = \beta[F(x)x + \int_x^1 f_x(z)z\,dz] > x$. By the monotonicity of T, we have

$$\lim_{n\to\infty} T^n(w_0)(x) \geq Tw_0(x) = \beta\left[F(x)x + \int_x^1 f_x(z)z\,dz\right] > x$$

when $x < p$. To sum up, the limit function $\lim_{n\to\infty} T^n w_0$, which is equal to the value function v, has the desired property.

Finally, let's return to the uniqueness assumption on the solution to (21). Note that (21) can be written

$$\int_x^1 f_x(z)z\,dz = [\beta^{-1} - F(x)]x. \tag{22}$$

A sufficient condition for uniqueness is that the derivative of the left-hand side of (21) is positive at any solution. That is,

$$1 - \beta\left[-f_x(x)x + f(x)x + F(x) + \int_x^1 \frac{\partial f_x(z)}{\partial x} z\,dz\right] > 0,$$

[25] $C([0, 1])$ is the Banach space of continuous functions $f : [0, 1] \to \mathbb{R}$ equipped with the sup norm.

or equivalently

$$-f_x(x)x + f(x)x + F(x) + \int_x^1 \frac{\partial f_x(z)}{\partial x} z\,dz < \beta^{-1},$$

whenever x satisfies (22). Note that

$$f_x(z) = \frac{\partial}{\partial z} F_x(z) = 2\frac{1-F(z)}{1-F(x)} f(z) \tag{23}$$

and, in particular, $f_x(x) = 2f(x)$. Also:

$$\frac{\partial f_x(z)}{\partial x} = \frac{f(x)}{[1-F(x)]^2} 2[1-F(z)]f(z) = \frac{f(x)}{1-F(x)} f_x(z). \tag{24}$$

Using (22)–(24):

$$\begin{aligned} &-f_x(x)x + f(x)x + F(x) + \int_x^1 \frac{\partial f_x(z)}{\partial x} z\,dz = -f(x)x + F(x) + \frac{f(x)}{1-F(x)}[\beta^{-1} - F(x)]x \\ &< -f(x)x + F(x) + f(x)x \\ &< \beta^{-1}, \end{aligned}$$

where both inequalities follow from $\beta < 1$. This proves uniqueness. □

Proof of Proposition 8: Let $\triangle\tau_z = \tau_z^{(2)} - \tau_z^{(1)}$ be the elapsed time between the first and the second arrivals with value no smaller than z. Note that $\triangle\tau_z$ is independent of $\tau_z^{(1)}$ under our assumptions.

Denote by $P_z^{(k)}$ the cumulative probability distribution of the random variable $\tau_z^{(k)}$; that is, $P_z^{(k)}(t)$ is the probability that $V^{(k)}$ goes above z at or before t. Denote by $P_z^{\triangle}$ the cumulative distribution function of $\triangle\tau_z$. Because $\triangle\tau_z$ is independent of $\tau_z^{(1)}$, we have,

$$\begin{aligned} \int_0^\infty e^{-\rho t_2} P_z^{(2)}(dt_2) &= \int_0^\infty \int_0^\infty e^{-\rho t_1} e^{-\rho\triangle t} P_z^{\triangle}(d\triangle t) P_z^{(1)}(dt_1) \\ &= \int_0^\infty e^{-\rho\triangle t} P_z^{\triangle}(d\triangle t) \int_0^\infty e^{-\rho t_1} P_z^{(1)}(dt_1). \end{aligned} \tag{25}$$

Finally, let $Q_z^{(k)}(\cdot \mid t)$ be the probability distribution of $V^{(k)}$ at time t, conditional on $\tau_z^{(k)} = t$; that is, conditional on $V^{(k)}$ jumping above r at time t. Note that $Q_z^{(k)}(\cdot \mid t)$ is in fact independent of t, because the conditional probability only depends on the fact that the kth highest value has just exceeded z for the first time, and not when this occurs.

With these notational preparations, the expected discounted revenue for a seller using a posted price mechanism, with price p, can be written:

$$p \int_0^\infty e^{-\rho t} P_p^{(1)}(dt). \tag{26}$$

To understand this expression, note that the item is sold for a price of p , at the time when $V^{(1)}$ first exceeds p; that is, at time $\tau_p^{(1)}$.

By Lemma 9, there exists a value $s \in [0, 1]$ such that an optimal stopping rule for the seller is to stop as soon as V_r weakly exceeds s. We claim that for any reserve r and stopping price s, there exists a price-posting rule with expected revenue at least as high. It suffices to establish this claim under the additional assumption that s is chosen optimally with respect to r. Furthermore, because $r = s$ is equivalent to price posting, and $s < r$ is clearly suboptimal, we need only consider the case $s > r$.

Therefore, suppose $s > r$ and s is optimal given r. The expected discounted value of stopping process V_r as soon as it weakly exceeds s can be found as follows. Because $s > r$, the process is stopped as soon as $V^{(2)} \geq s$; that

is, at time $\tau_s^{(2)}$. Conditional on $\tau_s^{(2)} = t$, the expected value of the stopped process is

$$\int_0^\infty v\, Q_s^{(2)}(dv|t),$$

which is independent of t, as pointed out above. Therefore, the expected discounted value of the stopping rule is given by:

$$\begin{aligned}
&\int_0^\infty e^{-\rho t}\left[\int_0^\infty v\, Q_s^{(2)}(dv|t)\right] P_s^{(2)}(dt) \\
&= \left[\int_0^\infty v\, Q_s^{(2)}(dv|t)\right]\int_0^\infty e^{-\rho t} P_s^{(2)}(dt) \\
&= \left\{\int_0^\infty \left[\int_0^\infty v\, Q_s^{(2)}(dv|t)\right] e^{-\rho\triangle t} P_s^{\triangle}(d\triangle t)\right\}\int_0^\infty e^{-\rho t_1} P_s^{(1)}(dt_1),
\end{aligned} \tag{27}$$

where the last equality follows from (25).

The term

$$\int_0^\infty \left[\int_0^\infty v\, Q_s^{(2)}(dv|t)\right] e^{-\rho\triangle t} P_s^{\triangle}(d\triangle t) \tag{28}$$

is the expected discounted continuation value of the selling mechanism, evaluated at time $\tau_s^{(1)}$. If s is optimal, this value must be no greater than s. To see why, suppose $V^{(2)}(\tau_s^{(2)}) = s$, so the seller has s in hand at time $\tau_s^{(2)}$. The optimality condition implies that s is at least as great as the continuation value from waiting for the next increment in V_r (which is equivalent to waiting for the next arrival with a value above s) and taking V_r at that point. Under Assumptions 1–5, the latter continuation value is equal to (28).

Therefore, we have:

$$\int_0^\infty \int_0^\infty v\, Q_s^{(2)}(dv|t) e^{-\rho\triangle t} P_s^{\triangle}(d\triangle t)\int_0^\infty e^{-\rho t_1} P_s^{(1)}(dt_1) \le s\int_0^\infty e^{-\rho t_1} P_s^{(1)}(dt_1).$$

Recalling (26), the right-hand side is the expected discounted revenue from posting price $p = s$. This proves our claim: for any optimal choice of r and s, there is a price posting mechanism that generates the same expected discounted revenue for the seller.

Of course, in our repeated auction game, the seller does not choose r and s, but r and T. That is, the seller uses a stopping rule of the form: stop the process at the first multiple of T at which $V_r \ge r$. This 'fixed time' stopping rule is, *a fortiori*, no better than the stopping rule based on an optimal choice of s, because the latter is optimal amongst all possible stopping rules. However, because price posting can be achieved by setting $T = 0$, it follows that we can replicate an optimal stopping rule using 'fixed time' stopping and, hence, that $T = 0$ is optimal in the public-reserve auction mechanism. □

References

Brown, T. C. (1984), "Poisson approximations and the definition of the Poisson process," *American Mathematical Monthly* **91,** 116–23.

Chen, Y., and R. W. Rosenthal (1996), "Asking prices as commitment devices," *International Economic Review* **37,** 129–55.

Diamond, P. (1971), "A model of price adjustment," *Journal of Economic Theory* **3,** 156–68.

Graham, D. A., R. Marshall, and J. F. Richard (1990), "Phantom bidding against heterogeneous bidders," *Economics Letters* **32,** 13–17.

Grant, S., A. Kajii, F. Menezes, and M. J. Ryan (2002), "Auctions with options to re-auction," Discussion Paper 55, Tilburg University, Center for Economic Research. [online; cited 25 January 2005]. Tilberg University, Tilberg, the Netherlands]. Available: http://greywww.kub.nl:2080/greyfiles/center/2002/doc/55.pdf

Krishna, V. (2002), *Auction Theory*, San Diego: Academic Press.

McAfee, R. P., and J. McMillan (1987), "Auction with a stochastic number of bidders," *Journal of Economic Theory* **43,** 1–19.

McAfee, R. P., and D. Vincent (1997), "Sequentially optimal auctions," *Games and Economic Behavior* **18,** 246–76.

Menezes, F. M., and P. K. Monteiro (2004), *An Introduction to Auction Theory*, Oxford and New York: Oxford University Press.

Menezes, F. M., and M. J. Ryan (2005), "Reserve price commitments in auctions," *Economics Letters* **87,** 35–39.

Milgrom, P. R. (1987), "Auction theory," in T. Bewley, ed., *Advances in Economic Theory: Fifth World Congress of the Econometric Society*, Econometric Society Monographs, no. 12, 1–32, New York and Melbourne: Cambridge University Press.

Milgrom, P. R. (2004), *Putting Auction Theory to Work*, Cambridge, New York and Melbourne: Cambridge University Press.

Milgrom, P. R., and R. J. Weber (1982), "A theory of auctions and competitive bidding," *Econometrica* **50,** 1089–122.

Nagareda, T. (2003), Announced reserve prices, secret reserve prices, and winner's curse [online; cited 1 February 2006]. Available: http://www.e.u-tokyo.ac.jp/cirje/research/workshops/micro/micropaper02/gs/nagareda.pdf

Reynolds, S. S., and J. C. Wooders (2004), "Auctions with a buy price," Eller College Working Paper, No. 1010-05, University of Arizona, Tuscon, AZ, USA.

Riley, J. G., and W. F. Samuelson (1981), "Optimal auctions," *American Economic Review* **71,** 381–92.

Riley, J. G., and R. Zeckhauser (1983), "Optimal selling strategies: When to haggle, when to hold firm," *Quarterly Journal of Economics* **98,** 267–89.

Roth, A. E., and A. Ockenfels (2002), "Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the internet," *American Economic Review* **92,** 1093–103.

Stokey, N. L., and R. E. Lucas Jr (1989), *Recursive Methods in Economic Dynamics*, Cambridge, MA: Harvard University Press.

Synder, D. (1975), *Random Point Processes*, New York: John Wiley & Sons.

Wang, R. (1993), "Auctions versus posted-price selling," *American Economic Review* **83,** 838–51.

Wang, R. (1995), "Bargaining versus posted-price selling," *European Economic Review* **39,** 1747–64.

Copyright of International Journal of Economic Theory is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.