

Auctions with Options for Re-auction*

Abstract

We examine the role of seller bidding and reserve prices in an infinitely repeated independent-private-value (IPV) ascending-price auction. The seller has a single object that she values at zero. At the end of any auction round, she may either sell to the highest bidder or pass-in the object and hold a new auction next period. New bidders are drawn randomly in each round. The ability to re-auction motivates a notion of reserve price as the option value of retaining the object for re-auctioning. Even in the absence of a mechanism with which to commit to a reserve price, the optimal “secret” reserve is shown to exceed zero. However, despite the infinite repetition, there may be significant value to the seller from a binding reserve price commitment: the optimal binding reserve is higher than the optimal “secret” reserve, and may be substantially so, even with very patient players. Furthermore, reserve price commitments may even be *socially* preferable at high discount factors. We also show that the optimal “phantom” bidding strategy for the seller is revenue-equivalent to a commitment to an optimal public reserve price.

JEL: D44, D82.

Key words: phantom bidding, re-auction option, reserve price, internet auctions.

Simon Grant
CentER Fellow
Department of Econometrics & Operations Research
Tilburg University
and
School of Economics
Australian National University

Atsushi Kajii
Institute of Social and Policy Planning
Tsukuba University

Flavio Menezes
School of Economics
Australian National University

Matthew J. Ryan
School of Economics
Australian National University

*While accepting responsibility for all remaining shortcomings of the paper, we would like to thank seminar participants at the Australian National University, University of New South Wales and University of Sydney 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 gratefully acknowledged.

1 Introduction

We examine the role of seller bidding and reserve prices in a repeated Independent Private Values (IPV) auction model. There is a single object for sale, and an ascending-price auction may be held in each of an infinite number of periods. At the end of any period, the seller can either sell to the highest bidder or pass-in the object for re-auctioning in the next period. Bidders are assumed to arrive randomly over time, according to a Poisson process. The ability to re-auction ensures that, with probability one, the object will eventually be sold. Hence, the seller's optimal reserve price is not based on the direct consumption value of the good to the seller (normalized to zero), but on the *option value* of retaining the object for subsequent re-sale. This repeated auction structure also captures some of the features of internet auctions for items such as used computer parts, software, etc.

For the IPV framework, it was established by Myerson (1981) and Riley and Samuelson (1981) that standard one-shot auctions (such as first- or second-price sealed-bid auctions) with a reserve price strictly above the seller's valuation of the object will maximize expected revenue within a wide class of mechanisms. In particular, the seller gains by withholding the object in some situations in which gains from trade exist, in order to extract a greater share of the surplus in trades with higher-value buyers. Clearly, such selling mechanisms generate an inefficiently low volume of trade from a Social Planner perspective: only auctions with a reserve price equal to the seller's value are efficient.

These results assume an exogenous, and commonly known, number of bidders, n , but the optimal reserve price is independent of n when all bidder valuations are drawn from the same distribution (the case considered here). Moreover, the second-price auction with reserve remains optimal even if *bidders* are uncertain of n (McAfee and McMillan (1987-a)). Indeed, provided bidders have symmetric posteriors over the number of rivals they face, these authors show that the first-price auction with reserve will also remain optimal.¹

¹ However, a "no reserve" auction may be optimal if the number of bidders is determined by endogenous entry in the presence of participation costs incurred before bidders learn their values for the object: see McAfee and McMillan (1987-b), Engelbrecht-Wiggans (1987,1993), and Levin and Smith (1994). In these models, the private gain from further entry corresponds to that of society (and the seller). Since the seller has no reason to discourage entry, the reserve price loses its only appeal. Menezes and Monteiro (2000) show that the value of a reserve is

More crucial, however, is the underlying assumption that the seller is able to commit to withhold the object from the bidders if the reserve is not met. While it is arguable that such commitment may be implementable *within* a given auction round, it is implausible to suppose that the seller can commit not to re-auction the object. By explicitly allowing for re-auctioning, we re-examine the value of a (temporary) commitment to a reserve price.

We find that binding reserves continue to be valuable to the seller, much as being the first mover in a Stackelberg duopoly game is preferable to playing the Cournot game. In fact, both the optimal “secret” reserve – the one employed in the absence of commitment where any announcement by the seller is interpreted by bidders as “cheap-talk” – and the optimal binding (or “public”) reserve are *both* higher than in the one-shot case. This is because the option value of the object exceeds its direct consumption value. The optimal “secret” reserve is thus raised to this higher option value level; while the optimal binding also increases, because the opportunity cost of not selling in the current round is reduced.

However, in practice, reserve prices are sometimes kept secret. For example, Horstmann and LaCasse (1997) argue that a common feature of many auctions, namely the seller’s refusal to sell to the high-bidder, might be interpreted as evidence of the existence of a secret reserve.² These authors explain the existence of a secret reserve price as a signalling device about the (common) value of the object being sold. Such information could not be credibly transmitted to buyers by a publicly announced reserve price because of adverse selection. Instead the seller uses costly delay as a signalling device.³

We offer an alternative rationalization for the rejection of high bids. A rejected high bid may indicate that the high bidder was in fact the seller. In the absence of a commitment technology, “phantom” bidding allows the seller to replicate the effect of a binding public reserve. Indeed, we show that seller bidding (or *shilling*) results in the *same expected revenue* as a binding public reserve when bidders learn their value for the object before making their entry decision. In particular they show that the optimal auction again involves setting a reserve price above the seller’s valuation.

² Ashenfelter (1989) reports significant percentages of high-bid rejections in wine and fine art auctions. Jones, Menezes and Vella (1996) report high-bid rejection rates for wool auctions in Australia.

³ Vincent (1995) argues that a secret reserve price might be used to increase bidder participation at the auction in a common value environment. His model, however, cannot explain why high bids are often rejected once an auction takes place.

reserve. This equivalence may at first seem surprising, since shilling gives the seller far greater strategic flexibility than a public reserve. In effect, shilling allows the seller to re-set the reserve price at any stage during the auction. We show that this added flexibility, however, adds no value for the seller. This contrasts with a view held by many that phantom bidding helps the seller. Indeed, eBay does not allow the seller to bid in its internet auctions. Similarly, the New South Wales State Government in Australia has moved to end phantom bidding by requiring bidders to register and by limiting to one the number of bids a Vendor is allowed to make.

In our repeated auction framework, the seller's optimal reserve price, and the value of commitment, are natural generalizations of the one-shot results. However, the welfare conclusions are quite different. Because the good must eventually be sold, the question of allocative efficiency asks whether it is sold too quickly or too slowly. We show that the sale is made too quickly in the absence of a reserve price commitment (unless the seller bids on her own account). That is, the "secret" reserve price is *too low*. This contrasts with the one-shot scenario, in which the "secret" reserve – which is just the seller's direct consumption valuation of the good – is efficient. With Uniformly distributed bidder valuations, we show that the optimal public reserve remains too high relative to the social optimum. However, unlike the one-shot case, it may be socially preferable to the "secret" reserve at high discount factors.

In the sequel we also consider the more general problem in which the seller chooses the duration of each auction round, as well as the reserve price. In eBay auctions, for example, the seller may choose a 3, 5, 7 or 10 day format. Unsurprisingly, if reserve prices are "secret", then the optimal length of a round of the auction is finite and strictly positive. For the "public" reserve auctions, we identify two countervailing forces at work. Shorter auctions reduce the time costs of the selling mechanism. However, longer auctions raise the probability of two or more bidders arriving, and hence increase the competition amongst bidders. In the limit, as the length of each round goes to zero, the likelihood that the object will be sold at the reserve goes to one: the auction collapses to a posted price mechanism. Simulations reveal that a zero-length auction (i.e. price posting) is in fact optimal when bidder values are Uniformly distributed.

Related Literature

The existing literature on repeated auctions has focussed on the case of long-lived bidders. For example, McAfee and Vincent (1997) consider an infinitely repeated, second-price IPV model in which the same n bidders participate in each round. As discount rates converge to unity, a version of the Coase conjecture takes hold, and the reserve price converges rapidly to zero over time. With bidder valuations distributed Uniformly on $[0, 1]$, McAfee and Vincent show that the value of commitment to a reserve price is seriously undermined by the inability to commit not to re-auction.

Burguet and Sákovics (1996) obtain a contrary result. They show that optimal reserves may be high (relative to the utility of the object to the seller) and valuable in a two-stage, first-price IPV model when buyers face participation costs, even with no discounting. They offer no justification, however, for their assumption that the seller can credibly commit not to re-auction the object after the second stage. Moreover, the seller is further assumed not to impose any reserve in the second-stage auction, despite having had access to a suitable commitment technology in the first period to implement a first-period reserve.

Haile (2000) also has a two period model where re-auctioning is assumed not to be possible after the second stage. Instead, a re-sale market opens in the second period, in which the seller commits not to participate. Since buyers are uncertain about their true valuations when they bid, this re-sale market is always active with some positive probability. Haile's main results concern the effect of the re-sale market on optimal bidding strategies in the (second-price) auction, but he also shows that high reserves may be valuable.

Therefore, when bidders are (infinitely) long-lived, and the seller can never credibly commit not to re-auction, it seems that reserve prices lose much of their value to the seller. However, the case of short-lived bidders – where each round attracts a new cohort of participants – has, to the best of our knowledge, not been studied. As we demonstrate below, with short-lived bidders, the Coasian logic vanishes and optimal reserves are typically even higher than in the one-shot case. This optimal reserve, in the absence of a suitable commitment technology, can be implemented by seller bidding, but seller bidding cannot improve on the optimal public reserve auction.

Thus, our model dispenses with the untenable assumption of a seller commitment not to re-

auction the object, but (a) implies a high value of commitment to a reserve price; and (b) provides a possible explanation for the phenomena of objects being passed-in in the absence an announced reserve, even when a costless technology exists to bind the seller the reserve.

The model is presented in the next section. Some general results are presented in section 3, and in section 4 we derive (by simulation) the optimal secret and public reserve prices when bidder valuations are Uniformly distributed, as well as the optimal length of each auction round. Longer proofs and calculations appear in the Appendix.

2 The model

An auction may be held in each of an infinite number of discrete periods. Within any given round, there is one unit of time⁴, with elapsed time described (in the obvious fashion) by the Uniform density on $[0, 1]$. We let t denote the elapsed time in the current round. No time elapses between the end of one auction round and the start of the next. Although we ignore any fixed monetary cost of conducting an auction, there is one cost to re-auctioning in our model: any bidder present at time $t = 1$ in some round is assumed no longer to be present at time $t = 0$ of the next.⁵

There are infinitely many potential buyers. Bidders arrive according to a Poisson process with rate $\lambda > 0$, so that, within any given round, the probability of n arrivals in the interval $[s, t] \subseteq [0, 1]$ is

$$e^{-\lambda(t-s)} \frac{\lambda^n (t-s)^n}{n!}.$$

In this setting, buyers arrive one at a time within any given auction period: multiple simultaneous arrivals occur with probability zero. Buyers' values are determined by independent draws from a common distribution F with strictly positive density f on $[0, 1]$. A buyer's own value is private information; but F is common knowledge.

We assume that bidders play "myopically" (e.g. that each auction round attracts a different pool of bidders). The seller, however, has regard for the future beyond the current round, and a

⁴ We shall relax this assumption, by making the duration of auction rounds endogenous, in section 3.6.

⁵ In internet auctions, for example, where automated proxy bidding is common, there may be a standing bid at $t = 1$ without the bidder being on-line. Also, the subsequent auction may potentially be conducted at a different site, with no "forwarding address" posted for bidders in the previous round.

(continuous) rate of time preference $\rho > 0$. The seller values the object at zero.

We consider cases where a publicly announced reserve may or may not be enforceable, and the seller may or may not be able to submit phantom bids. This yields the following four alternative auction mechanisms:

1. ‘Cheap-talk’ reserve price (CTRP) – announced reserves are *not* enforceable, and phantom bids are precluded.
2. Public (binding) reserve price (PRP) – announced reserves are enforced, but phantom bids are precluded.
3. Seller Phantom Bidding (with cheap-talk reserve) – announced reserves are *not* enforceable, but phantom bids are allowed.
4. Seller Phantom Bidding and public binding reserve price – publicly announced reserves are enforced, and phantom bids are allowed.

Descriptions of the operation of these alternate mechanisms, the corresponding strategy choices available to the seller, and the strategies of the bidders who arrive during the course of the auction are outlined in the subsequent subsections.

2.1 Cheap-talk-reserve-price

At time $t = 0$ the standing bid is automatically set to zero. The seller may, if she wishes, announce a price at which she purports to be willing to sell the object if the standing bid at the end of the current period has reached that amount. But this announcement in no way restricts her later choice of action.⁶ Seller (phantom) bids are precluded.

When the first bidder arrives he may post a bid to become the new standing bidder. As each subsequent bidder arrives, he engages in an open ascending auction with the standing bidder until one drops out and the other becomes the new standing bidder. Bidders may not re-enter the auction once they have dropped out.⁷

⁶ Porter (1995) reports that 12.7% of U.S. off-shore oil and gas leases auctioned between 1970 and 1979 were passed in, despite the announced reserve being exceeded.

⁷ This assumption is innocuous. See Section 3.

We assume that the bidding process between the current standing bidder and the new arrival consumes zero time. Thus, our framework captures an environment in which bilateral bidding contests consume amounts of time that are negligible relative to the length of each auction round, and also relative to the arrival rate of bidders.⁸

At the end of each period ($t = 1$), the seller may either accept the current standing bid and sell the object immediately to the high-bidder, thereby ending the repeated auction; or reject the standing bid and pass-in the object for re-auctioning. We preclude seller negotiations with bidders at the conclusion of an auction.

2.2 Public-(binding-)reserve-price

At time $t = 0$ the seller chooses the initial standing bid which is taken to be an *enforceable* reserve price. The standing bid changes only if a bidder arrives who is prepared to bid above this reserve. The auction otherwise continues as in the cheap-talk-reserve-price scenario.

2.3 Seller-phantom-bidding

At time $t = 0$ the standing bid is set to zero. The auction proceeds as in the cheap-talk-reserve-price scenario, except that at any time (*including* $t = 0$) the seller may imitate the arrival of a new buyer and bid on her own account. Since negotiation between bidders and seller are precluded by assumption, if a phantom bid stands at the end of the auction round, the object is automatically passed-in for re-auction next period.

2.4 Seller-Phantom-Bidding with public binding reserve price

This scenario combines a binding public reserve with phantom bidding by the seller. It is redundant, however, as the ability to submit phantom bids effectively undermines the enforceability of the reserve price: the seller can always trump a standing bid to force a re-auction. Hence, this scenario is equivalent to the seller-phantom-bidding (without an enforceable public reserve)

⁸ An example of an environment that fits this description is “proxy bidding” in the popular eBay on-line auction. Arriving bidders submit their maximum bids to automated “proxy bidders”. These proxy bidders maintain the bidder in the standing bidder position until the maximum bid level is exceeded. They also implement the bidding strategy at light-speed.

scenario. We shall therefore confine attention to the preceding three auction formats in the rest of the paper.

3 General results

3.1 Bidding strategies

Because bidders have independent private values, participate in at most one auction round, and the auctions have a second-price structure, optimal bidding strategies are straightforward to determine. There is a weakly dominant strategy which is to bid up to one's valuation. For example, bidders could submit maximum bids equal to their values to proxy bidders in an eBay auction.

Let us be a little more precise. The optimal bidding strategy is not affected by the presence or otherwise of the seller as a phantom bidder. Phantom bids have the same consequences for bidders as "genuine" rival bids: if a phantom bid stands at $t = 1$, the object is passed in for re-auctioning.

Suppose that a new bidder arrives and finds a standing bidder in residence. Since re-entry is precluded, it is optimal for the arriving bidder to remain in the bidding war with the standing bidder unless his (the new arrival's) valuation is exceeded.⁹ In fact, even if re-entry within the current round were allowed, there would be nothing for the arrival to gain by doing otherwise. He might drop out of the current bidding war "early", but will always find it optimal to re-enter before $t = 1$ and bid until his valuation is reached (unless it has already been exceeded in the meantime). The assumption that bidders may not re-enter avoids the unnecessary complication of these timing issues.¹⁰

Next, consider the decision of an arrival in the absence of a standing bidder. If a public reserve has been posted, this is the *de facto* standing bid, so the new bidder will either depart (if his value is no greater than this), or bid the reserve (plus epsilon) to become the standing bidder.

⁹ If the new arrival's value is no greater than the standing bid, he departs without bidding.

¹⁰ Some authors, such as Roth and Ockenfels (2001), have observed a tendency for "sniping", or last-minute bidding, in on-line auctions. However, the use of proxy bidding in eBay auctions avoids this. Other on-line auctions employ "soft" endings to the same effect. Auctions may be extended for short periods (say, 5 minutes) until no further activity is recorded; a sort of virtual "going...going..gone!". This, too, restores the usual second-price logic, and eliminates any incentive to "snipe". See Lucking-Reiley (2000).

In a cheap-talk-reserve-price auction, if no standing bidder is in residence, the effective standing bid is zero. However, if bidders anticipate that the seller has a non-zero secret reserve r , a first arrival will optimally bid the secret reserve (plus epsilon) if his value exceeds r . If he did not, he will expect to “lose” the auction, and fail to acquire the object at a price less than his valuation. If his value is no greater than r , then he has several optimal courses of action. He may choose not to bid, or to submit any bid less than r . None of our results depends upon which option the bidder chooses in this situation, as will be clear.

Finally, consider an arriving bidder who wins a bidding war with the standing bidder in a cheap-talk-reserve-price auction. If the bidder believes that the seller has a secret reserve of r , and the previous standing bidder dropped out of the bidding war before bidding reached r , then, as in the previous case, the new bidder will submit one further bid equal to r (plus epsilon). This may happen, for example, if the previous standing bidder was the first arrival, and had a value less than r (see the previous paragraph).

In summary, let r denote the reserve price anticipated by bidders: r is the public reserve for auctions in which public reserves are posted and binding; or the anticipated secret reserve in a cheap-talk-reserve-price format. A bidder with value $v \in [r, 1]$ will submit bids in $[r, v]$ just high enough to meet the reserve and remain the standing bidder. If bidding exceeds v they will drop out of the auction. Bidders with values in $[0, r)$ expect zero surplus from participation in the auction. They may choose not to bid, or else submit a bid in $[0, r)$.

3.2 The seller’s decision problem

Observe that the seller’s problem is stationary: it looks the same at the start of each auction round. Leaving aside seller bidding for the moment, the seller’s problem is to choose a reserve price, or acceptance rule, for each auction period. Because of the stationarity of the problem, we assume that the same reserve price is applied in each period.

In a cheap-talk-reserve-price auction, the reserve is secret, so we may think of the seller deciding on her acceptance rule at $t = 1$: she may accept or reject the standing bid at that time.¹¹ With

¹¹ There is also no advantage from separately announcing the reserve, as it has no signalling value in the IPV context, and is thus ignored by bidders: it is “cheap talk”.

a public reserve price, the acceptance rule is set at $t = 0$. Moreover, the reserve will affect the bidders' strategies, as described above. In each case, however, the institutional assumptions we have made impose the restriction that the acceptance rule will take the form of a cut-off (reserve) price, r , such that the standing bid at $t = 1$ is accepted iff it is at least r . This restriction is innocuous.

If seller bidding is allowed, the seller has a much more complex strategy space. In addition to setting a (secret) reserve, she also chooses a bid function that may depend on the elapsed time in the current round, the reserve, and the level of the last genuine bid submitted. By stationarity, we may assume that the seller employs the same bidding strategy in any round, but the range of potential bid functions is still very large. Fortunately, there exists an optimal seller bidding strategy with a very simple form. We derive this strategy in subsection 3.5.

3.3 The value of a public reserve

Let us first compare, from the seller's point of view, the cheap-talk-reserve-price format to an auction with a public-(binding-)reserve-price. Is there value to the seller from committing to a (public) reserve price?

For a cheap-talk-reserve-price (CTRP) auction, the seller's strategy involves choosing an acceptance price (or reserve), r , being the lowest price at which she is prepared to sell the object in any round. This choice will depend on the bidders' strategies, which in turn depend on their expectation of the reserve price. We impose the usual equilibrium conditions that players maximize their expected payoffs at all points in the game, and beliefs about rival strategies are correct.

Let $v_0(r)$ be the seller's discounted expected payoff when she chooses a reserve price of r , and all bidders correctly anticipate r and play a best response. Equilibrium in the CTRP game therefore requires that

$$r = v_0(r) \tag{1}$$

Equation (1) says that the seller's reserve is equal to the value of re-starting the auction game, given the bidders' optimal responses to this reserve. If bidding fails to reach v_0 , it is better to pass in the object; and accepting a bid above v_0 is preferable to re-auctioning the object.

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

The formal proof appears in the Appendix, in which an explicit expression for $v_0(r)$ is derived, but the basic idea is straightforward. First, continuity of v_0 guarantees the existence of a fixed point. Second, it is obvious that $v_0(0) > 0$ and $v_0(1) = 0$. That is, the seller expects a non-zero surplus in a “no reserve” auction; and setting a reserve of $r = 1$ ensures that the probability of achieving a sale in any given round is zero.

For the public-(binding-)reserve-price (PRP) scenario, the seller’s strategy again consists of choosing an acceptance price (reserve) r . However, r must now be stated publicly and committed to at the start of the auction round. The optimal reserve price therefore solves

$$\max_{r \in [0,1]} v_0(r) \tag{2}$$

Again, continuity of v_0 ensures the existence of a solution to (2).

The seller can clearly do no worse in the PRP auction, since the value of the CTRP auction is $v_0(r)$ evaluated at an r that satisfies (1). The seller would rather choose the *best anticipated* reserve, than choose a reserve that is *anticipated* and best given this *fixed bidder expectation*. Hence we have:

Proposition 2 *From the seller’s point of view, the public-reserve-price scenario is preferable to the cheap-talk-reserve-price scenario.*

In section 4 we shall quantify the value a reserve price commitment for the case of bidders with Uniformly distributed values.

In general, the optimal PRP (“public”) reserve will be higher than the equilibrium CTRP (“secret”) reserve price. Raising a reserve is more beneficial if done publicly, since it induces higher bids from high-value arrivals. Conversely, reductions in reserve are best kept secret, to raise the chances of a sale without encouraging bid reductions from sole high-value bidders. This creates a tendency for reserves to be higher when public and binding, than when they are secret. More precisely:

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

$$v_0(r^*) = r^* - \frac{[1 - F(r^*)]}{f(r^*)} \quad (3)$$

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

Re-arranging (3) gives

$$r^* = v_0(r^*) + \frac{[1 - F(r^*)]}{f(r^*)}$$

This matches the familiar expression for the optimal reserve in a one-shot IPV auction – see, for example, Riley and Samuelson (1981, Proposition 3) – except that the seller’s valuation of the object, $v_0(r^*)$, is now endogenous: it is the “option value” of retaining the object for re-auctioning. Since $v_0(r^*) > 0$, the optimal public reserve is *higher* than in the one-shot case, since the seller’s opportunity cost of not trading in any given round is greater than the consumption value of the object. With short-lived bidders, the Coasian logic of McAfee and Vincent (1997) is reversed.

Recall that in a one-shot auction, the optimal public reserve is independent of the number n of bidders. However, r^* will, in general, vary with λ , the average number of arrivals per round, since λ affects the option value $v_0(r^*)$.¹² However, this dependence vanishes as the seller becomes increasingly myopic, and r^* converges to the usual one-shot reserve price.¹³

3.4 The surplus maximizing reserve price

In a one-shot auction, any non-zero reserve price is surplus reducing, since there is some chance that no trade occurs. Hence, privately optimal reserve price commitments are necessarily too high from a social efficiency point of view. With a repeated auction, this is no longer the case. Unless the reserve price is set at unity – which is precluded by Propositions 1 and 3 – the object will

¹² For the Uniform case, see Figure 1.

¹³ Using equation (22) from the Appendix, we may define

$$V(r) \equiv \lim_{\rho \rightarrow \infty} e^\rho v_0(r) = \frac{\lambda [1 - F(r)] r + \lambda^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda[F(z) - F(r)]} dz}{e^{\lambda[1 - F(r)]}}$$

to be the (undiscounted) expected revenue from a *one-shot* auction with public reserve r . Simple algebra gives

$$V'(r) = 0 \iff r = \frac{1 - F(r)}{f(r)}$$

Therefore, the optimal public reserve in a one-shot auction with random bidder arrivals is independent of λ .

sell with probability one. The efficiency issue concerns whether the sale occurs too quickly or too slowly.

To be more precise, 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 period, the Planner may therefore choose to allocate the object to any one of the arrivals during that period, or else wait one more round. We assume that the Social Planner observes the values of all arrivals during the current round, but does not know the values of future arrivals.¹⁴ Which allocation rule maximizes the discounted expected value of total surplus?

As usual, stationarity implies that the Planner's rule will consist of a cut-off, r , such that the good is allocated to the highest value arrival during the current period if and only if the highest value exceeds r . Note, therefore, that the Social Planner's cut-off is a *valuation*, while 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 cut-off. In this sense, they are directly comparable for the purposes of determining the *allocative efficiency* of the auction.

Let $v_S(r)$ denotes the Planner's value function. The Social Planner problem is of the "optimal stopping" variety. In each period, a maximum surplus is drawn at random according to some density h on $[0, 1]$, with associated cumulative distribution function H .¹⁵ The Planner's problem is to determine a cut-off level of surplus at which to stop the process. In general:

$$\frac{1}{\delta}v_S(r) = [1 - H(r)]\mathbb{E}[z | z \geq r] + H(r)v_S(r) \quad (4)$$

(where $\delta = e^{-\rho}$). Thus, using the fact that

$$\frac{d}{dr}\mathbb{E}[z | z \geq r] = \frac{-h(r)}{1 - H(r)}\{r - \mathbb{E}[z | z \geq r]\},$$

we have

$$\frac{1}{\delta}v'_S(r) = -h(r)r + H(r)v'_S(r) + h(r)v_S(r).$$

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

¹⁵ In the following calculations we treat these as continuous distributions for simplicity. Of course, in the auction example they are not. In particular, there is an atom of size $e^{-\lambda}$ at zero. However, we shall shortly verify that the same results obtain in this case.

Therefore, $v'_S(r) = 0$ iff $v_S(r) = r$. The Planner should allocate the object whenever the highest value arrival exceeds the continuation value, and this cut-off should also maximize the value function.

In the present situation, one may show that¹⁶

$$v_S(r) = \frac{\delta \left[1 - r e^{-\lambda[1-F(r)]} - \int_r^1 e^{-\lambda[1-F(z)]} dz \right]}{1 - \delta e^{-\lambda[1-F(r)]}} \quad (5)$$

If we let $\hat{N}(r)$ denote the numerator of (5) and $\hat{D}(r)$ the denominator, one easily shows that

$$\frac{\hat{N}'(r)}{\hat{D}'(r)} = r$$

so

$$\frac{\hat{N}'(r)}{\hat{D}'(r)} = \frac{\hat{N}(r)}{\hat{D}(r)} \quad \text{iff} \quad \frac{\hat{N}(r)}{\hat{D}(r)} = r$$

That is, \hat{r} maximizes $v_S(r)$ if and only if \hat{r} solves

$$v_S(r) = r \quad (6)$$

Hence, the presence of atoms in the distribution of maximum surplus does not affect the basic optimal stopping logic.

We may immediately conclude:

Proposition 4 *The CTRP auction sells the object too quickly from a social efficiency point of view. That is, if r^{**} and \hat{r} solve (1) and (6) respectively, then $r^{**} < \hat{r}$.*

Proof. It is clear that $v_0(r^{**}) < v_S(r^{**})$, since the seller must share the (strictly positive) expected surplus with the buyer, while the Social Planner does not. Therefore, if $\hat{r} \leq r^{**}$ we have

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

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

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

¹⁶ Details are in the Appendix.

high bid earns r for both the seller and the Social Planner, while passing-in earns the former $v_0(r)$ and the latter $v_S(r)$. Since the Social Planner anticipates the whole expected surplus from future trade, while the seller only anticipates a fraction – that is $v_0(r) < v_S(r)$ – the Social Planner is more patient. The CTRP auction sells too quickly.

It is not so clear whether the optimal *public* reserve is too high or too low. Let $R(r)$ denote the expected revenue generated in an auction with (binding) reserve r , and let $S(r)$ be the expected total surplus from such an auction. We also define $\Pi_1(r)$ to be the probability that no bidder with value at least r arrives during the auction. Then, using the expressions in the Appendix, one easily shows that:

$$v'_S(r) \propto [1 - \delta\Pi_1(r)] \delta S'(r) + \delta^2\Pi'_1(r)S(r)$$

$$v'_0(r) \propto [1 - \delta\Pi_1(r)] \delta R'(r) + \delta^2\Pi'_1(r)R(r).$$

The second term in each expression reflects the “option” value of not selling in the current round as a result of raising the reserve. The first term reflects the impact of an increased r on current period returns (surplus or revenue respectively).

Since $\Pi'_1(r) \geq 0$, and $S(r) \geq R(r)$, we see that

$$\delta^2\Pi'_1(r)S(r) \geq \delta^2\Pi'_1(r)R(r) \geq 0.$$

This reflects the extra patience of the Social Planner due to her higher “option” value.

On the other hand $S'(r) \leq 0$, since raising r lowers the chances of an allocation being made in any period, without affecting the expected value of the highest-value arrival. Raising r therefore reduces the expected realized surplus within any given round, and is a cost to the Social Planner. For the seller, however, there are gains as well as losses from raising r . First, there is the reduced probability of sale, as for the Social Planner. However, since the seller does not always get all the surplus from a sale, the opportunity cost is smaller than for the Social Planner. Second, raising r increases the seller’s portion of surplus when exactly one bidder arrives with a value above the reserve.¹⁷

¹⁷ By direct calculation:

$$R'(r) = S'(r) + \lambda[1 - F(r)][1 - (1 - r)\lambda f(r)] e^{-\lambda[1 - F(r)]}.$$

The overall effect is ambiguous: the optimal public reserve price may either be too high or too low from a surplus-maximization point of view. In section 4, we show that the public reserve is too high when bidder values are Uniformly distributed.

3.5 Public reserves versus seller bids

A seller bid submitted at time t has the same effect on bidder behavior as imposing a public reserve at t equal to the phantom bid. The ability to make phantom bids is therefore equivalent to the ability to set a public reserve at $t = 0$ and alter it continuously throughout the auction. Hence, the seller-phantom-bidding format is the best of the three, and it follows that the seller is happy to forego access to a reserve price commitment technology provided she can bid in the auction.

Of course, many auctions, including on-line eBay auctions, have rules against seller bidding. The extent to which these rules are enforceable, especially in the context of virtual auctions, is highly debatable. Surprisingly, however, in our framework these regulations are entirely redundant, as the seller is *completely indifferent* between the seller-phantom-bidding and public-reserve-price formats. The additional strategic freedom afforded by seller bidding – essentially, the freedom to use a “flexible reserve” – adds no additional value.

Graham, Marshall and Richard (1990) showed that in a one-shot, second price, IPV auction, if the seller has the opportunity to bid once all genuine bidding activity has ceased, her optimal bid is independent of the standing bid. We extend this result to our framework. The incentive to phantom bid at t is shown to be independent of both the current standing bid and t , even though additional bidding activity may occur in $(t, 1]$ and there is the possibility of re-auctioning the object.

Before stating this result formally, let us first define a suitable value function for a seller-phantom-bidding auction. Once again, by the stationarity of the problem, the value function will not depend on the period, but it will depend on the time t that has elapsed in the current period. Therefore, if t units of time have elapsed in the current round, and there is a standing bid b , we let $\hat{v}_t(r, b)$ denote the seller’s expected continuation value when she nominates a public reserve price

of r at t and uses an optimal continuation strategy thereafter.

We shall make use of the following “regularity assumption”:

Assumption 5 *The value function $\hat{v}_t(r, b)$ is differentiable in r for any $(t, b) \in [0, 1] \times [0, 1)$ and any $r \in (b, 1)$. Moreover,*

$$\frac{\partial \hat{v}_t(\hat{r}, b)}{\partial r} = 0$$

is a sufficient condition for $\hat{r} \in (b, 1)$ to maximize $\hat{v}_t(\cdot, b)$ on $(b, 1]$.

In particular, Assumption 5 entails that $\hat{v}_t(r, b)$ is continuous in r , so the *ex ante* value of the seller-phantom-bidding auction is

$$\max_{r \in [0, 1]} \hat{v}_0(r, 0) \tag{7}$$

Under this regularity condition we obtain the following equivalence result.

Proposition 6 *Under Assumption 5, the seller-phantom-bidding and public-reserve-price auction formats generate the same ex ante expected revenue for the seller. That is:*

$$\max_{r \in [0, 1]} \hat{v}_0(r, 0) = \max_{r \in [0, 1]} v_0(r) \tag{8}$$

3.6 Optimal auction length

Until now, we have normalized the length of auctions to one unit of time. However, on-line auctions often allow sellers to choose their duration, at least within some bounds. It is therefore of interest to endogenize the length of auction rounds.

Let T denote this parameter. The seller’s value function becomes¹⁸

$$v_0(r, T) = \frac{\lambda T [1 - F(r)] r + (\lambda T)^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda T [F(z) - F(r)]} dz}{e^{\{\rho + \lambda [1 - F(r)]\} T} - 1} \tag{9}$$

It is clear that all of the preceding results apply (*pari passu*) for any given T . The question we address here is what value of T will the seller choose under each of the cheap-talk-reserve-price and public-reserve-price formats?

¹⁸ Equation (9) follows by straightforward adaptation of the calculation of $v_0(r)$ in the Appendix.

By familiar logic, a seller who uses a public reserve will choose (r, T) to maximize $v_0(r, T)$.

With a secret reserve, for any T , r must satisfy

$$r = v_0(r, T) \tag{10}$$

Assuming that this has a unique solution $r(T)$ for each T , T will be chosen to maximize $v_0(r(T), T)$.

We shall allow T to be chosen from \mathbb{R}_+ . That is, duration is infinitely divisible, and $T = 0$ is a possible choice. With regard to the latter assumption, what does it mean to run an auction lasting zero units of time? Observe that the value function (9) is not even defined when $T = 0$.

Formally, we may define $v_0(r, 0)$ as follows

$$v_0(r, 0) \equiv \lim_{T \rightarrow 0} v_0(r, T) \tag{11}$$

However, an alternative, more intuitive, definition is the following. Since the probability of two or more simultaneous arrivals is zero, instantaneous re-auctioning ($T = 0$) means that the object is offered for sale to the first person willing to pay the reserve price, r . The object will therefore be sold to the first arriving bidder whose value exceeds r , and the sale price will be exactly r . A zero-length auction is simply a *posted price mechanism*. If it is a zero-length CTRP auction, the posted price r is “cheap talk”, while the reserve in a zero-length PRP auction is a posted price to which the seller can commit in any bilateral bargaining process with a potential buyer.

Thus:

$$v_0(r, 0) = \int_0^\infty e^{-\rho t} r \mathcal{P}_r(t) dt$$

where $\mathcal{P}_r(t)$ is the density associated with the arrival time of the first bidder with a valuation that exceeds r . Since bidders with values in excess of r arrive according to a Poisson process with mean $\lambda[1 - F(r)]$, it is well known that arrival times will have the exponential density

$$\mathcal{P}_r(t) = \lambda [1 - F(r)] e^{-\lambda t [1 - F(r)]}.$$

Therefore:

$$\begin{aligned}
v_0(r, 0) &= r\lambda [1 - F(r)] \int_0^\infty e^{-\rho t - \lambda[1 - F(r)]t} dt \\
&= r\lambda [1 - F(r)] \left[\frac{e^{-\rho t - \lambda[1 - F(r)]t}}{-(\rho + \lambda[1 - F(r)])} \right]_{t=0}^\infty \\
&= \frac{r\lambda [1 - F(r)]}{\rho + \lambda [1 - F(r)]} \tag{12}
\end{aligned}$$

Re-writing (12), we have:

$$\rho v_0(r, 0) = \lambda [1 - F(r)] [r - v_0(r, 0)] \tag{13}$$

This expression is a Bellman equation familiar from the literature on optimal search.¹⁹ The imputed instantaneous rate of income accumulation, $\rho v_0(r, 0)$, is equal to the instantaneous arrival rate of bidders with value above r , $\lambda [1 - F(r)]$, multiplied by the net benefit from such an arrival, $r - v_0(r, 0)$.

Let us first observe that definitions (11) and (12) are equivalent. Using definition (12) for $v_0(r, 0)$, we may write

$$\begin{aligned}
\frac{v_0(r, T)}{v_0(r, 0)} &= \\
&\frac{\lambda [1 - F(r)] r + T\lambda^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda T[F(z) - F(r)]} dz}{r\lambda [1 - F(r)]} \times \left[\frac{\{\rho + \lambda [1 - F(r)]\} T}{e^{\{\rho + \lambda [1 - F(r)]\} T} - 1} \right]
\end{aligned}$$

Since the second term goes to 1 as $T \rightarrow 0$ (by L'Hôpital's rule), we see that

$$\begin{aligned}
\lim_{T \rightarrow 0} \frac{v_0(r, T)}{v_0(r, 0)} &= \lim_{T \rightarrow 0} \frac{\lambda [1 - F(r)] r + T\lambda^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda T[F(z) - F(r)]} dz}{r\lambda [1 - F(r)]} \\
&= 1 + \lim_{T \rightarrow 0} \frac{T\lambda^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda T[F(z) - F(r)]} dz}{r\lambda [1 - F(r)]} \\
&= 1.
\end{aligned}$$

¹⁹ See, for example, Mortensen (1986).

Now that we have defined the value function $v_0(r, T)$ on $[0, 1] \times \mathbb{R}_+$ we may address the question of the optimal auction duration. For the case of a cheap-talk-reserve-price auction it is clear that $T = 0$ cannot be optimal:²⁰

Proposition 7 *For a cheap-talk-reserve-price auction, some $T > 0$ will be chosen.*

Proof. When $T = 0$, the equilibrium condition (10) is

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

which is satisfied if and only if $r = 0$. Thus

$$v_0(r(0), 0) = 0.$$

Since we already know that $v_0(r(1), 1) > 0$, the result follows. \square

When the seller cannot commit to her posted price, and she faces a “search cost” of finding another buyer (the expected delay until the next arrival), the only credible price announcement is zero. Absence of commitment deprives the seller of any effective market power. To acquire any surplus from trade, 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 non-zero auction length to overcome her price commitment problem. Of course, longer rounds do increase buyer competition and hence seller surplus as before, but they also delay trade. It is not clear *a priori* which T will provide an optimal balance between these two effects. However, the case in which F is the Uniform distribution shows that zero-length auctions (price posting) may be optimal with a public reserve (see section 4).

4 The case of Uniformly distributed values

Our objective in the present section is to identify the optimal secret and public reserve prices when F is the Uniform distribution. We may then determine the value (to the seller) from a commitment

²⁰ Readers familiar with the literature on search will recognise the logic underlying Proposition 7 as a version of the so-called *Diamond paradox* (Diamond (1971)).

to a public reserve. We also compute the social costs of the alternative auction formats, and the optimal auction length.

For the case of Uniformly distributed bidder valuations, $f(z) = 1$, $F(z) = z$. Hence, as we show in the Appendix, $v_0(r)$ may be expressed as:

$$v_0(r) = \frac{e^{\lambda(1-r)} \left(1 - \frac{2}{\lambda}\right) + \left(1 - 2r + \frac{2}{\lambda}\right)}{e^{\rho + \lambda(1-r)} - 1} \quad (14)$$

Straightforward calculations reveal that

$$\begin{aligned} v_0'(r) &\geq 0 \\ \Leftrightarrow 2\delta e^{-\lambda(1-r)} &\geq [2\delta - \lambda(1 + \delta)] + 2\lambda r \end{aligned} \quad (15)$$

(where $\delta \equiv \exp(-\rho)$).

4.1 Optimal public reserve

Consider the public-reserve-price scenario.

Proposition 8 *There is a unique optimal reserve in the public-reserve-price auction for any $(\lambda, \delta) \in \mathbb{R}_{++} \times (0, 1)$. This optimal reserve lies in $(0.5, 1)$, converging to 0.5 as $\lambda \rightarrow 0$ or $\delta \rightarrow 0$.*

Proof. Existence of an optimal public reserve, and the fact that any such reserve price is strictly less than 1, are guaranteed by Proposition 3. Uniqueness of the optimal public reserve follows from the facts that the left-hand side of (15) is convex in r , while the right-hand side is linear. The lower bound on the optimal reserve comes from equation (3) in Proposition 3, which implies

$$r^* > 1 - r^*$$

and hence $r^* > 0.5$.

Letting $r = 0.5$ in (15) gives

$$\frac{2\delta}{e^{\lambda/2}} \geq (2 - \lambda)\delta \quad (16)$$

Both sides of this expression converge to 2δ as $\lambda \rightarrow 0$, or to 0 as $\delta \rightarrow 0$. This proves that $r = 0.5$ is optimal in each of these limiting cases, since we have already established the existence of a unique interior maximizer of $v_0(r)$. \square

In a one-period version of this auction game, it is well-known that $r = 0.5$ is the optimal (enforceable) reserve for any $\lambda > 0$ (i.e. for any number of bidders). Hence, if $\delta \rightarrow 0$ the optimal reserve must converge to $r = 0.5$. A similar result obtains in a model in which the same bidders vie for the object in each round – see McAfee and Vincent (1997, p.250). However, in their model, the optimal reserve in period 1 converges to 0.5 *from below* as $\delta \rightarrow 0$ (*ibid.*, Figure 1(b)). When facing the same set of bidders each period, the inability to commit *not* to re-auction the object places downward Coasian pressure on the reserve.

If $\delta > 0$ and re-auction is possible, the seller has an incentive to raise her reserve above $r = 0.5$, since she receives a positive value even if the object is passed in. However, as the expected number of bidders in any given round goes to zero, the value of this re-auction option becomes negligible, so the optimal reserve again converges on $r = 0.5$.

Conversely:

Proposition 9 *For any $\delta \in (0, 1)$, the optimal public reserve converges to $r = (1 + \delta)/2$ as $\lambda \rightarrow \infty$; while for any $\lambda > 0$, the optimal public reserve converges to $r = 1$ as $\delta \rightarrow 1$.*

Proof. The left-hand side of (15) goes to zero as $\lambda \rightarrow \infty$. The right-hand side will only go to zero if $r \rightarrow (1 + \delta)/2$.²¹

As $\delta \rightarrow 1$, (15) converges to the condition

$$e^{\lambda r} \geq e^{\lambda} [1 - \lambda(1 - r)].$$

The left-hand side of this expression is strictly increasing and strictly convex, rising to e^{λ} as $r \rightarrow 1$. The right-hand side is linear and strictly increasing in r , also rising to e^{λ} as $r \rightarrow 1$. When $r = 1$ both sides have slope λe^{λ} . Therefore, the two sides are equal if and only if $r = 1$. \square

Contrast this result with McAfee and Vincent (1997, p.251 and Figure 1(d)). In their model, with values drawn from the Uniform distribution on $[0, 1]$, the optimal first period reserve also rises with the number of bidders, n , but converges to the static solution, $r = 0.5$. In our model, a higher arrival rate λ augments the incentive to pass-in the object and search for higher value

²¹ This puts an upper bound on the optimal reserve price (for given δ), since it is intuitive that r is increasing in λ . Indeed, if $r > (1 + \delta)/2$, then the right-hand side of (15) strictly exceeds 2δ (for any $\lambda > 0$), while the left-hand side is no greater than 2δ for any $r \leq 1$ and $\lambda > 0$.

bidders. This reduces the pressure on the seller to sell in any given round, and pushes the reserve further and further above its optimal one-shot value. In the McAfee and Vincent model, with the same bidders each period, increasing the number of bidders raises the reserve for a quite different reason. Higher n raises the pressure on bidders to bid early, allowing the seller to raise her reserve. In the limit, the competitive pressure on bidders completely overcomes the Coasian dynamic, and the optimal one-shot reserve is achieved.

In contrast to our Proposition 9, McAfee and Vincent (1997, p.251 and Figure 1(b)) obtain that the optimal reserve declines with increases in δ in their model.

4.2 Optimal secret reserve

Let us now turn to properties of the cheap-talk-reserve-price auction. The first observation is the following:

Proposition 10 *There exists a unique equilibrium reserve in the cheap-talk-reserve-price auction for any $(\lambda, \delta) \in \mathbb{R}_{++} \times (0, 1)$.*

Proof. Use $v_0 = r$ in (14) and re-arrange the resulting expression to get

$$e^\lambda [z + 2\delta - \delta\lambda] = \delta e^z [2 + \lambda - z] \quad (17)$$

where $z = \lambda r$. The left-hand side of (17) is linear and strictly increasing in z , while the right-hand side is strictly increasing and strictly convex. Hence there will be at most two solutions to (17) in $z \in [0, \lambda]$. A necessary and sufficient condition for two solutions is that the LHS \geq RHS at $z = 0$ and $z = \lambda$. The latter is always the case, but the former is so if and only if

$$e^\lambda \geq \frac{2 + \lambda}{2 - \lambda} \quad (18)$$

In fact, one may verify that (18) fails for any $\lambda > 0$. Hence, the equilibrium secret reserve is unique. \square

Proposition 3 implies that the optimal secret reserve will be strictly smaller than the optimal public reserve. Furthermore, it is immediate from (17) that:

Proposition 11 *For any $\lambda > 0$, the equilibrium secret reserve $r \rightarrow 0$ as $\delta \rightarrow 0$.*

This is just convergence to the optimal secret reserve in the one-shot auction.

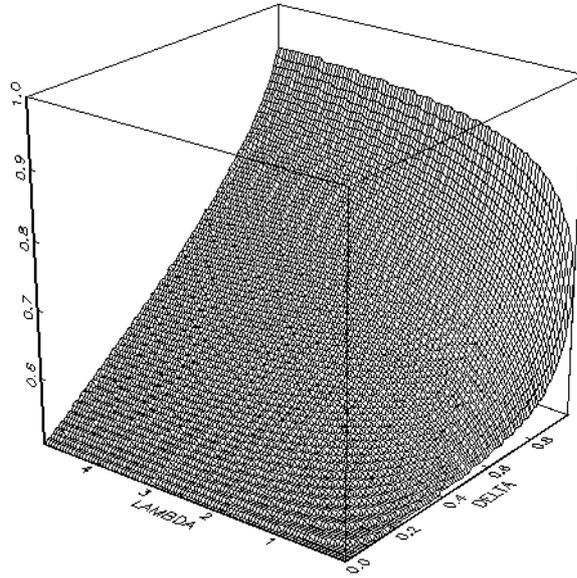


Figure 1: Optimal public reserve price

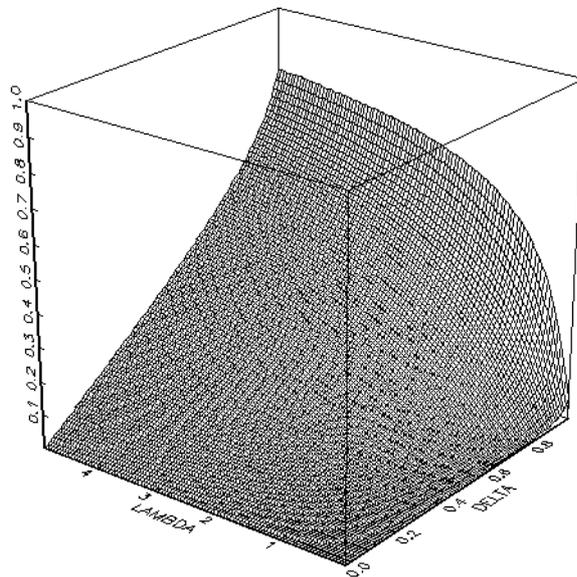


Figure 2: Equilibrium secret reserve price

4.3 Numerical results

4.3.1 The case $T = 1$

Although explicit solutions for the optimal secret and public reserves are elusive, we may easily obtain them numerically. Figures 1 and 2 give the optimal reserves for a wide range of (δ, λ) values.²²

The figures illustrate the behavior of the reserves for limiting values of the parameters. In relation to Proposition 8, we observe that convergence of the optimal public reserve to 0.5 as $\lambda \rightarrow 0$ is comparatively slow. Figure 1 uses λ values that are bounded below by $\lambda = 0.1$. One verifies Proposition 8 by taking λ values down to $\lambda = 0.005$, as in Figure 3.

Similarly, Proposition 9 is also obscured by the scale Figure 1. Indeed, it is clear that Propositions 8 and 9 imply “extreme” behavior of the surface near $(\delta = 1, \lambda = 0)$, with convergence to a non-smooth contour.

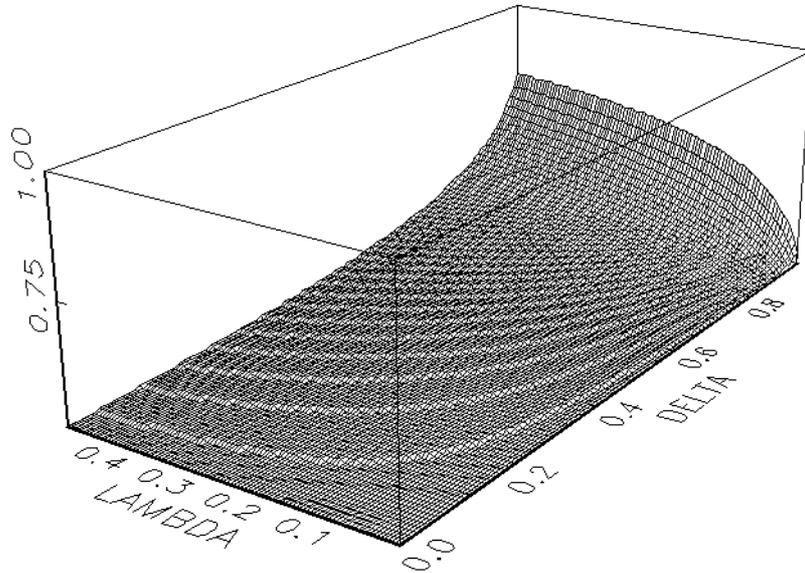


Figure 3: Optimal public reserve for low λ values

Figure 4 plots $r^{PRP} - r^{CTRP}$, where r^{PRP} denotes the optimal public reserve, and r^{CTRP}

²² Recall that $\delta = e^{-\rho}$.

is the equilibrium secret reserve price. This is the increment in the reserve as a result of public announcements binding the seller.

Figure 5 illustrates the value of this commitment to a public reserve. It plots the gain from commitment, as a percentage of the secret reserve value. That is:

$$100 \left[\frac{v_0(r^{PRP}) - v_0(r^{CTRP})}{v_0(r^{CTRP})} \right].$$

We observe that this value falls sharply as the average number of arrivals per period, λ , rises. The percentage gain is below 10% (for any δ) before λ reaches 5. Gains also tend to diminish with rises in δ , though much more slowly. In this sense, the inability to commit not to re-auction the object does devalue the commitment to a reserve price, but we do not see the dramatic decline in value observed by McAfee and Vincent (1997) in the case of long-lived bidders.

Since public reserves are higher than secret reserves, on average it will take longer to sell an object using the public-reserve-price format. The probability of concluding a sale in any given period is $1 - \Pi_1(r)$, which is equal to $1 - e^{-\lambda(1-r)}$ in the Uniform case. Let $p(r, \lambda)$ denote this quantity. It is decreasing in r for any λ , as one would expect. Figure 6 indicates the increment in this probability from using a cheap-talk-reserve-price auction: $p(r^{CTRP}, \lambda) - p(r^{PRP}, \lambda)$.²³

The expected number of periods until a sale is achieved is $p(r, \lambda)^{-1}$. Figure 7 indicates the expected extra delay (in numbers of periods) from using the public-reserve-price format rather than a cheap-talk-reserve-price auction: $p(r^{PRP}, \lambda)^{-1} - p(r^{CTRP}, \lambda)^{-1}$.

Finally, let us compare the optimal public reserve to the socially optimal reserve price. When F is Uniform on $[0, 1]$, the Social Planner's value function (5) becomes

$$v_S(r) = \frac{\delta \{1 - r e^{-\lambda(1-r)} - e^{-\lambda(1-r)} (r - \lambda^{-1})\}}{1 - \delta e^{-\lambda(1-r)}} \quad (19)$$

Figure 8 reveals that the optimal public reserve is too high. It graphs the difference between the optimal public reserve and the maximizer of (19).

²³ This difference depends on δ , as well as λ , since r^{CTRP} and r^{PRP} depend on both parameters.

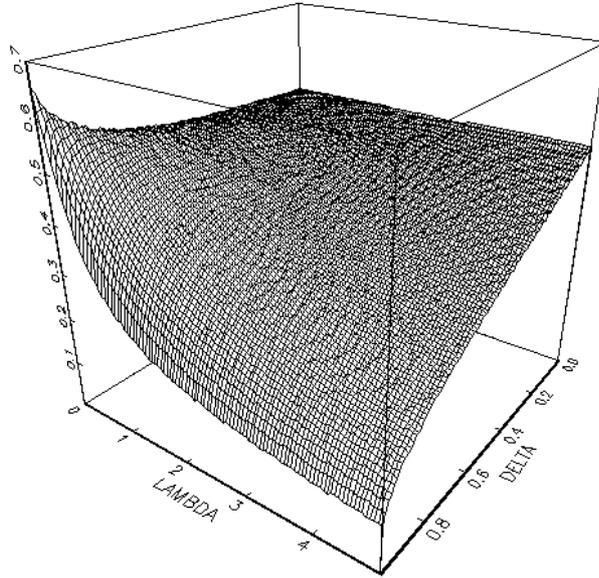


Figure 4: Public reserve less secret reserve

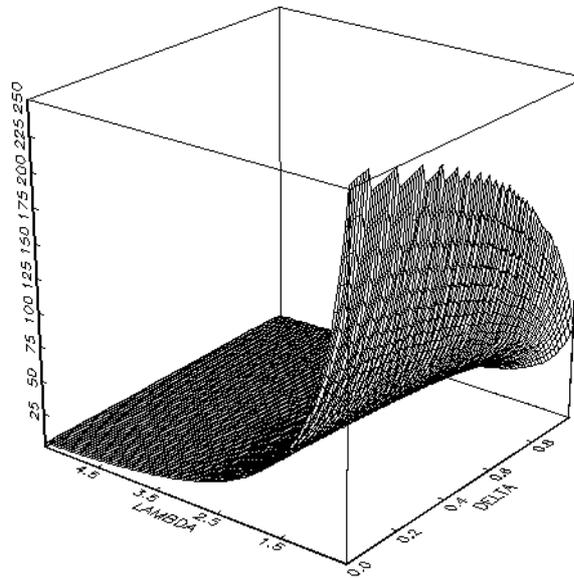


Figure 5: Value of commitment

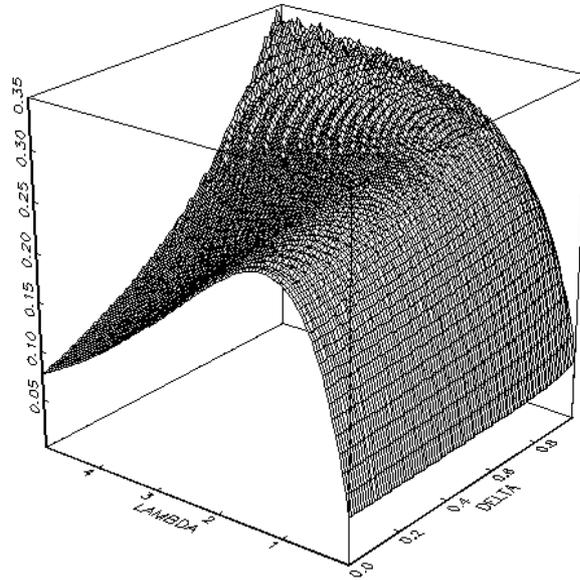


Figure 6: Extra per period sale probability from a secret reserve

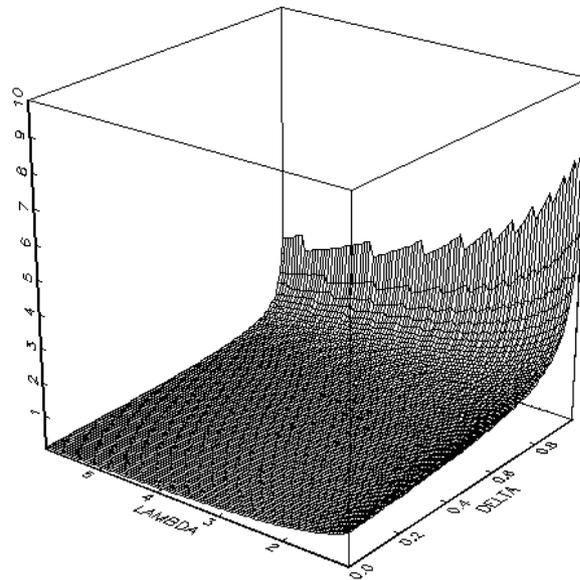


Figure 7: Expected extra delay from a public reserve

Since we know (Proposition 4) that the equilibrium secret reserve price is too low, it is of interest to determine which auction format comes closest to achieving the socially optimal expected surplus. Are reserve price commitments socially desirable? Figure 9 provides the answer: it plots the difference between $v_S(r^{CTRP})$ and $v_S(r^{PRP})$ (as a percentage of the latter). One observes that, for sufficiently high discount factors, the public reserve is relatively more efficient.

4.3.2 Endogenous T

For the special case in which F is Uniform, Figure 10 plots the value function $v_0(r, T)$ when $\rho = 0.1$ and $\lambda = 3$, using (9).

It is easy to observe that the optimal choice of (r, T) will have $T = 0$.²⁴ In fact, if we plot

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

as a function of T (again assuming that values are Uniformly distributed, $\rho = 0.1$ and $\lambda = 3$) we obtain Figure 11. Furthermore, straightforward calculations reveal that $v_0(r, 0)$ is maximized when

$$r = (1 + \theta) - \sqrt{\theta(1 + \theta)} \quad (21)$$

²⁴ We obtain the same conclusion for all other parameter values that we have tried. In fact, one may show that

$$\left. \frac{dv_0(r^*(T), T)}{dT} \right|_{T \downarrow 0} < 0 \quad (20)$$

where $r^*(T)$ denotes the optimal public reserve given T . To do so, we first calculate

$$\begin{aligned} \frac{\partial v_0(r, T)}{\partial r} &= 0 \\ \Leftrightarrow e^{-[\rho + \lambda(1-r)]T} &= e^{-\rho T} + \frac{\lambda T}{2} [2r - 1 - e^{-\rho T}] \end{aligned}$$

Now substitute this into the expression resulting from the calculation

$$\frac{\partial v_0(r, T)}{\partial T} \geq 0,$$

divide through by T^2 and simplify to obtain

$$\begin{aligned} -[\rho + \lambda(1-r)] \left[\lambda + \lambda(1-2r)e^{-\lambda T(1-r)} + \frac{2(e^{-\lambda T(1-r)} - 1)}{T} \right] &\geq \\ \left[\frac{e^{\rho T} - 1}{T} - \frac{\lambda}{2} e^{-\rho T} + \lambda r \right] \left[\lambda(1-r)(\lambda T - 2) - \frac{2(e^{-\lambda T(1-r)} - 1)}{T} \right] & \end{aligned}$$

Taking limits as $T \rightarrow 0$ we get

$$-[\rho + \lambda(1-r)] 2\lambda \geq 0$$

Since $\rho + \lambda(1-r) > 0$, (20) follows.

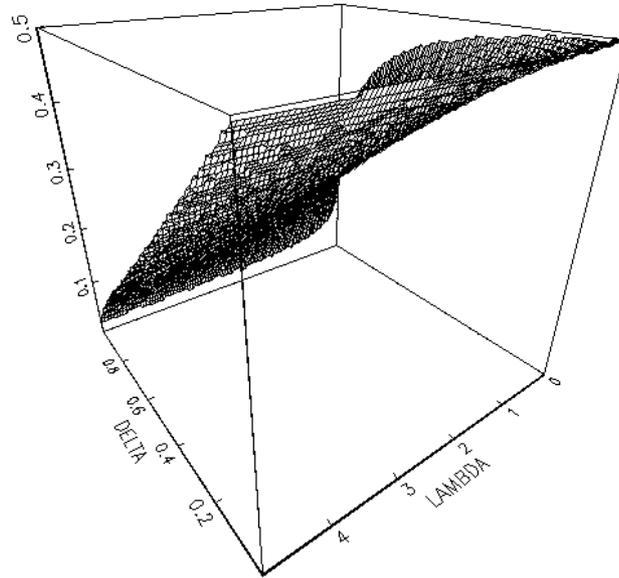


Figure 8: Optimal public reserve less socially optimal reserve.

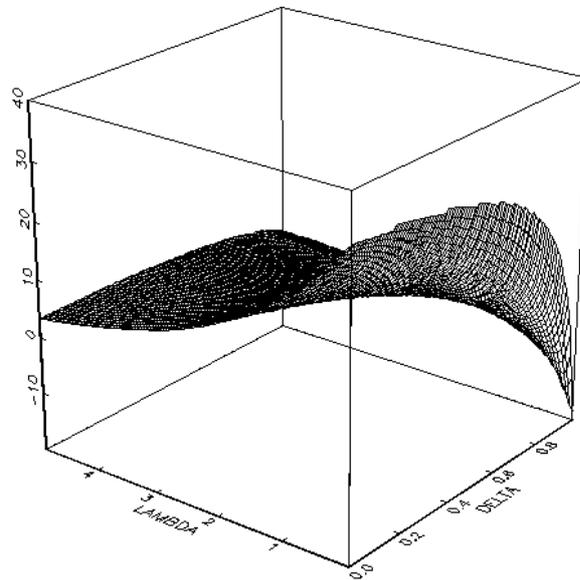


Figure 9: Social cost of commitment: $100 [v_S(r^{CTRP}) - v_S(r^{PRP})] / v_S(r^{PRP})$

(where $\theta = \rho/\lambda$). Hence, with Uniformly distributed valuations, the optimal public-reserve-price auction has $T = 0$ and reserve price (21). We therefore have a complete characterization of the optimal auction (within the scenarios considered here) for the Uniform case.

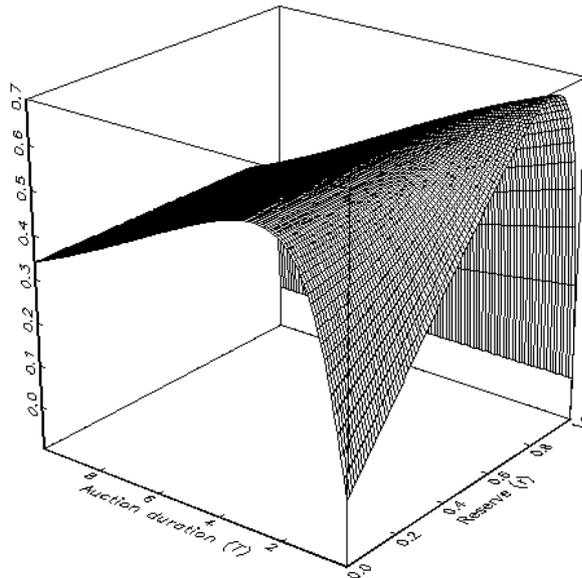


Figure 10: $v_0(r, T)$ when F Uniform, $\rho = 0.1$ and $\lambda = 3$

The optimality of $T = 0$ in a public-reserve-price-auction should be interpreted with care. It relies on our assumption that there is no delay *between* auction rounds, and no other fixed (monetary) cost to running another auction. In reality, at least one, and probably both, of these assumptions will be violated. For example, eBay charges sellers a small fee for each auction they run. Time or other costs of running many auctions will provide pressure to increase T .

On the other hand, the optimality of price-posting (in the presence of commitment) does not undermine the relevance of the model for the analysis of internet auctions. On eBay, for example, sellers have the option of posting a “Buy It Now” price. In the absence of a standing bid (above the reserve), if an arriving bidder offers the “Buy It Now” price, it is automatically accepted and the auction is cancelled.

Figure 12 illustrates $v_0(r(T), T)$, the value of a cheap-talk-reserve-price auction as a function of T . This has a maximum at an auction length around 2.

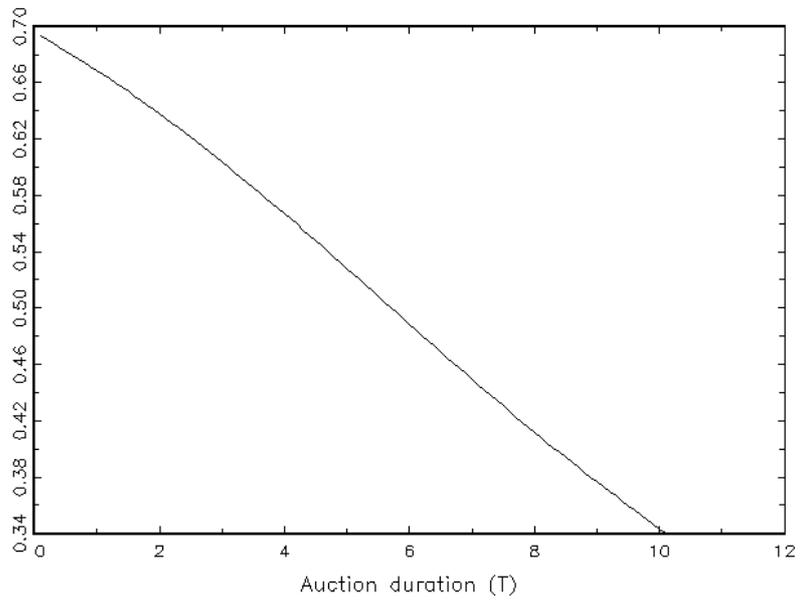


Figure 11: Auction value with a public reserve

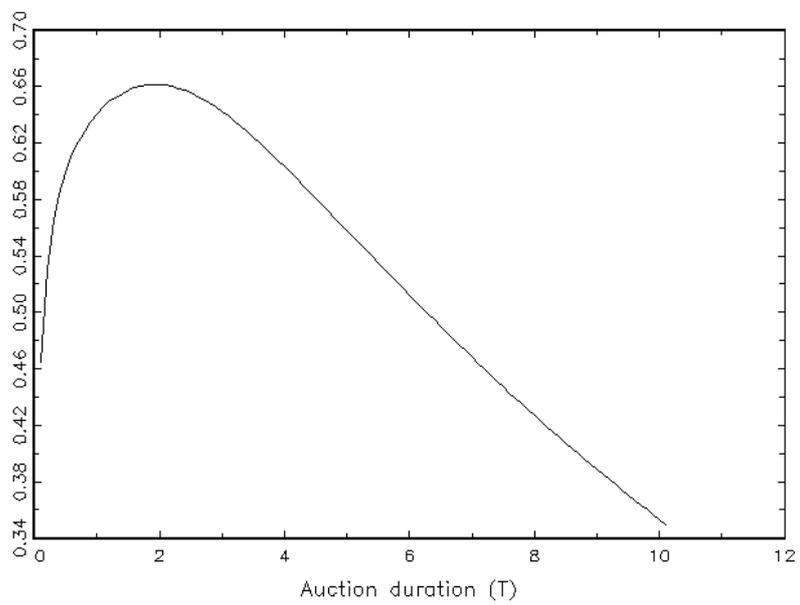


Figure 12: Auction value with a secret reserve

5 Concluding Remarks

The model presented here offers a natural generalization of the one-shot auction that allows for re-auctioning. As $\delta \rightarrow 0$ the familiar one-shot results are recovered. However, the generalized framework is useful in a number of respects.

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

Second, the potential to re-auction allows for a more realistic analysis of reserve prices. This potential generates a non-zero option value of retaining the object, and hence raises both “secret” and “public” reserves. However, with our short-lived bidders, a substantial value from a reserve price commitment remains. We are also able to confirm that seller bidding offers no advantage over a reserve price commitment in our setting.

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

Finally, by modelling the bidder arrival process and endogenizing auction duration, our framework also nests price posting as a limiting case (when $T \rightarrow 0$). One may then observe the relative value of price versus time commitments for revenue generation. In the absence of price commitments, the logic of the *Diamond paradox* implies value in a commitment to $T > 0$, as this will increase buyer competition. However, with Uniformly distributed bidder valuations and reserve price commitments, price posting is actually optimal for the seller.

Appendix

Proof of Proposition 1. We may express $v_0(r)$ as

$$\frac{e^{-\rho}R(r)}{1 - e^{-\rho}\Pi_1(r)},$$

where $\Pi_1(r)$ is the probability that, in any given period, the object is passed in, and $R(r)$ is the expected revenue from any given auction round. Thus, $\Pi_1(r)$ is equal to the probability that

no bidder arrives with a value in $[r, 1]$. The process of arrival of bidders with values in $[r, 1]$ is easily shown to be Poisson with mean $\lambda [1 - F(r)]$, so $\Pi_1(r) = \exp(-\lambda [1 - F(r)])$. The expected revenue per round is:

$$R(r) = \Pi_2(r)r + \int_r^1 \left(\sum_{n=2}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} g_2^n(z) \right) z dz,$$

where $\Pi_2(r) = \lambda [1 - F(r)] \exp(-\lambda [1 - F(r)])$ is the probability that the auction yields r – which is just the probability that exactly one bidder arrives with a valuation in $[r, 1]$. The function g_2^n is the density of the second order statistic from n random draws from the distribution F . In particular:

$$g_2^n(z) = n(n-1)f(z)F^{n-2}(z)[1 - F(z)].$$

Therefore, $\int_r^1 (\sum_{n=2}^{\infty} \exp(-\lambda) \lambda^n g_2^n(z)/n!) z dz$ is the expected value of the second-highest arrival with a value above r (the random variable defaulting to zero along sample paths for which there do not exist two such arrivals). The bracketed term in this expression may be simplified as follows:

$$\begin{aligned} \sum_{n=2}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} g_2^n(z) &= f(z) [1 - F(z)] \sum_{n=2}^{\infty} e^{-\lambda} \frac{\lambda^n}{(n-2)!} F^{n-2}(z) \\ &= f(z) [1 - F(z)] \lambda^2 \sum_{m=0}^{\infty} e^{-\lambda} \frac{\lambda^m}{m!} F^m(z) \\ &= f(z) [1 - F(z)] \lambda^2 e^{-\lambda [1 - F(z)]} \sum_{m=0}^{\infty} e^{-\lambda F(z)} \frac{[\lambda F(z)]^m}{m!} \\ &= \lambda^2 f(z) [1 - F(z)] e^{-\lambda [1 - F(z)]} \end{aligned}$$

Hence:

$$v_0(r) = \frac{\lambda [1 - F(r)] r + \lambda^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda [F(z) - F(r)]} dz}{e^{\lambda [1 - F(r)]} - 1} \quad (22)$$

Notice that $v_0(0) > 0$, $v_0(1) = 0$ and $v_0(r)$ is a continuous function. Hence from the intermediate value theorem it follows there exists an $\hat{r} \in (0, 1)$ for which $\hat{r} - v_0(\hat{r}) = 0$ as required. \square

Proof of Proposition 3.²⁵ Let

$$N(r) = \lambda [1 - F(r)] r + \lambda^2 \int_r^1 z f(z) [1 - F(z)] e^{\lambda [F(z) - F(r)]} dz$$

²⁵ We thank Rhema Vaithianathan for correcting an error in a previous “proof.”

and

$$D(r) = e^{\rho + \lambda[1-F(r)]} - 1.$$

Recall that $v_0(r) = N(r)/D(r)$ (see equation (22)). An optimal public reserve must therefore satisfy the necessary condition

$$\frac{N(r)}{D(r)} = \frac{N'(r)}{D'(r)}.$$

By direct calculation:

$$\frac{N'(r)}{D'(r)} = \frac{N(r) + r - \frac{[1-F(r)]}{f(r)}}{D(r) + 1}$$

Therefore,

$$\frac{N(r^*)}{D(r^*)} = \frac{N'(r^*)}{D'(r^*)}$$

if and only if (3) holds.

>From (1), (2) and (3): $r^{**} = v_0(r^{**}) \leq v_0(r^*) \leq r^*$. 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 1 rules out $r^{**} = 1$. Therefore, $r^{**} < r^* < 1$. Finally, $r^* > r^{**}$ and Proposition 1 imply $r^* > 0$. \square

Derivation of (5). Observe that

$$v_S(r) = \frac{\delta S(r)}{1 - \delta \Pi_1(r)}$$

where

$$\Pi_1(r) = e^{-\lambda[1-F(r)]}$$

is the probability that there is no arrival with a value above r during a period of length 1, and $S(r)$ is the expected total surplus generated in the current round (given that the object has not yet been allocated). The latter is

$$S(r) = \int_r^1 \left[\sum_{n=1}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} g_1^n(z) \right] z dz$$

where $g_1^n(z) = n f(z) F^{n-1}(z)$ is the pdf of the highest of n random draws from distribution F .²⁶

²⁶ We observe that $S(r) \geq R(r)$, so $v_S(r) \geq v_0(r)$, as one would expect. The difference represents the buyers' expected surplus from an auction with reserve r .

To evaluate $v_S(r)$, we first note that the square-bracketed term in the definition of $S(r)$ is

$$\sum_{n=1}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} n f(z) F^{n-1}(z) = \lambda f(z) e^{-\lambda[1-F(z)]}$$

Therefore:

$$\begin{aligned} S(r) &= \int_r^1 \lambda f(z) e^{-\lambda[1-F(z)]} z dz \\ &= 1 - r e^{-\lambda[1-F(r)]} - \int_r^1 e^{-\lambda[1-F(z)]} dz \end{aligned}$$

Proof of Proposition 6. Define

$$h(r | b, t) = \frac{\partial \hat{v}_t(r, b)}{\partial r}$$

to be the marginal incentive to raise the phantom bid from r to $r + dr$ at time t when there is a (genuine) standing bid of $b \leq r$. Then

$$\begin{aligned} h(r | b, t) &= \\ &= \left[\frac{1 - F(r)}{1 - F(b)} \right] h(r | r, t) + \left[\frac{F(r) - F(b)}{1 - F(b)} \right] \int_t^1 \mathcal{P}_r(t' | t) h(r | r, t') dt' \end{aligned}$$

where $\mathcal{P}_r(t + x | t)$ is the density describing the random variable x , being the elapsed time until the arrival of the first bidder with a value in excess of r , given no such arrivals by t . The arrival process of bidders with values above r is Poisson with mean $\lambda[1 - F(r)]$, so:

$$\mathcal{P}_r(t + x | t) = \lambda[1 - F(r)] e^{-\lambda[1-F(r)]x}.$$

Therefore:

$$\begin{aligned} h(r | b, t) &= \\ &= \left[\frac{1 - F(r)}{1 - F(b)} \right] h(r | r, t) + \frac{\lambda[1 - F(r)] [F(r) - F(b)]}{[1 - F(b)]} \int_t^1 e^{-\lambda[1-F(r)](t'-t)} h(r | r, t') dt' \end{aligned}$$

Defining $g_t(r) = h(r | r, t)$, we have $h(r | b, t) = 0$ iff

$$g_t(r) + \lambda[F(r) - F(b)] \int_t^1 e^{-\lambda[1-F(r)](t'-t)} g_{t'}(r) dt' = 0 \quad (23)$$

Let r^* denote the optimal public reserve in a public-reserve-price auction. We claim that an optimal seller bidding strategy is to counter-bid r^* at any t at which there is a standing bid $b < r^*$.

To verify our claim, it suffices to show that $g_t(r^*) = 0$ for all t when the seller adopts this strategy. It then follows that the optimal seller bidding strategy is equivalent to conducting an optimal public-reserve-price auction.²⁷

However, it is natural to suspect that $g_t(r^*)$ might strictly decrease with t . If there is less time remaining in the current round, the seller may wish to bid less aggressively. Nevertheless, one may show that in fact $g_t(r^*)$ is independent of t . In particular, given a standing bid of r^* at $t = 1$, the optimal counter-bid solves

$$\max_{r \geq r^*} \left[\frac{1 - F(r)}{1 - F(r^*)} \right] r + \left[\frac{F(r) - F(r^*)}{1 - F(r^*)} \right] v_0(r^*)$$

since the seller's strategy ensures a continuation value of $v_0(r^*)$ from the start of the next round. The first-order condition for this problem is

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

which is satisfied iff $r = r^*$ (recall equation (3)). Therefore, the optimal public reserve remains the optimal counter-bid at $t = 1$. So $g_t(r^*)$ does not decline strictly with t .

More generally, $g_t(r^*)$ may be decomposed as follows:

$$g_t(r^*) = \Gamma_t(r^*) - e^{-\{\rho + \lambda[1 - F(r^*)]\}(1-t)} \frac{f(r^*)}{[1 - F(r^*)]} v_0(r^*).$$

The first term, $\Gamma_t(r^*)$, represents the change in the expected revenue from the current round, while the second reflects the change in the probability that the object will not be sold in the current round. In the event of no sale, the seller's continuation strategy ensures a continuation value of $v_0(r^*)$. To determine $\Gamma_t(r^*)$, note that revenue is unaffected provided at least one new bidder arrives in $(t, 1]$ with a value above r^* . The impact on revenue is only non-zero if there are no such arrivals. Thus:

$$\Gamma_t(r^*) = e^{-\{\rho + \lambda[1 - F(r^*)]\}(1-t)} \left\{ \frac{f(r^*)r^*}{[1 - F(r^*)]} - 1 \right\}.$$

²⁷ In particular, the calculations here are based on the assumption that the seller makes a *single counter-bid* for each (b, t) , rather than engaging in a bidding war with the standing bidder. The independence of the optimal counter-bid with respect to the standing bid b implies that nothing is lost by excluding this alternative form of seller bidding strategy.

Substituting this into the expression for $g_t(r^*)$ and implies that $g_t(r^*) = 0$ iff

$$v_0(r^*) = r^* - \frac{[1 - F(r^*)]}{f(r^*)}.$$

Hence, the condition $g_t(r^*) = 0$ is independent of t , and is satisfied when r^* is the optimal public reserve (recall equation (3)).

Hence, an optimal strategy for the seller is to submit an optimal phantom bid at time $t = 0$, and play no further part in that round. This implies (8) and completes the proof. \square

Derivation of Equation 14. For the case of Uniformly distributed bidder valuations, $f(z) = 1$, $F(z) = z$. Hence, (22) becomes:

$$v_0(r) = \frac{\lambda(1-r)r + \lambda^2 \int_r^1 z(1-z)e^{\lambda(z-r)} dz}{e^{\rho+\lambda(1-r)} - 1}$$

The integral in the numerator may be evaluated using integration by parts:

$$\begin{aligned} \lambda^2 \int_r^1 z(1-z)e^{\lambda(z-r)} dz &= \lambda \int_r^1 z(1-z)\lambda e^{\lambda(z-r)} dz \\ &= \lambda \left[z(1-z)e^{\lambda(z-r)} \right]_r^1 - \int_r^1 (1-2z)\lambda e^{\lambda(z-r)} dz \\ &= -\lambda r(1-r) - \left[(1-2z)e^{\lambda(z-r)} \right]_r^1 - 2 \int_r^1 e^{\lambda(z-r)} dz \\ &= -\lambda r(1-r) + e^{\lambda(1-r)} \left[1 - \frac{2}{\lambda} \right] + \left(1 - 2r + \frac{2}{\lambda} \right) \end{aligned}$$

Therefore

$$v_0(r) = \frac{e^{\lambda(1-r)} \left(1 - \frac{2}{\lambda} \right) + \left(1 - 2r + \frac{2}{\lambda} \right)}{e^{\rho+\lambda(1-r)} - 1}$$

References

- Burguet, R. and J. Sákovic, 1996, "Reserve Prices without Commitment," **Games and Economic Behavior** **15**, 149-164.
- Durrett, R., 1991, **Probability: Theory and Examples**. Brooks/Cole Publishing, California.
- Diamond, P., 1971, "A Model of Price Adjustment," **Journal of Economic Theory** **3**, 156-168.
- Engelbrecht-Wiggans, R., 1987, "On Optimal Reservation Prices in Auctions," **Management Science** **33(6)**, 763-770.

- Engelbrecht-Wiggans, R., 1993, "Optimal Auctions Revisited, " **Games and Economic Behavior** **5**, 227-239.
- Graham, D. A., R. Marshall and J.F. Richard, "Phantom Bidding Against Heterogeneous Bidders," **Economics Letters** **32**, 1990, 13-17.
- Haile, P. A., 2000, "Auctions with Private Uncertainty and Resale Opportunities," forthcoming in **Journal of Economic Theory**.
- Harstad, R. M., J. H. Kagel and D. Levin, 1990, "Equilibrium Bid Functions for Auctions with an Uncertain Number of Bidders," **Economics Letters** **33**, 35-40.
- Holt, C. A., Jr., 1979, "Uncertainty and the Bidding for Incentive Contracts," **American Economic Review** **69**(4), 697-705.
- Horstman, I. and C. LaCasse, 1997, "Secret Reserve Prices in a Bidding Model with a Resale Option," **American Economic Review** **87**, 663-684.
- Jones, C., F. M. Menezes and F. Vella, 1996, "Auction Price Anomalies: Evidence from Wool Auctions in Australia." Australian National University Working Paper No. 303.
- Levin, D. and J. L. Smith, 1994, "Equilibrium in Auctions with Entry, " **American Economic Review** **84**(3), 585-599.
- Lucking-Reiley, D., 2000, "Auctions on the Internet: What's Being Auctioned, and How?" **Journal of Industrial Economics** **48**(3), 227-252.
- McAfee, R. P. and J. McMillan, 1987-a, "Auction with Entry," **Economics Letters** **23**, 343-347.
- McAfee, R. P. and J. McMillan, 1987-b, "Auctions 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, 2000, "Auctions with Endogenous Participation," **Review of Economic Design** **5**, 71-89.
- Milgrom, P. R. and R. J. Weber, 1982, "A Theory of Auctions and Competitive Bidding," **Econometrica** **50**, 1089-1122.
- Mortensen, D.T., 1986, "Job Search and Labor Market Analysis," in O. Ashenfelter and R. Layard (eds), **The Handbook of Labor Economics, Volume II**. Elsevier Science.
- Myerson, R. B., 1981, "Optimal Auction Design," **Mathematics of Operations Research** **6**(1), 58-73.
- Porter, R. H., 1995, "The Role of Information in Offshore Oil and Gas Lease Auctions." **Econometrica** **63**(1), 1-28.
- Riley, J.G. and W. F. Samuelson, 1981, "Optimal Auctions," **American Economic Review** **71**(3), 381-392.
- Roth, A., E. and A. Ockenfels, 2001, "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet." Unpublished working paper, Harvard University.
- Vincent, D., 1995, "Why Reserve Prices May Be Kept Secret," **Journal of Economic Theory** **65**(2), 575-584.