

The case for last-second bidding

Bryan Lim*

Department of Finance
University of Melbourne

November 17, 2010

Abstract

Several recent empirical studies document the frequent occurrence of “sniping” – the act of submitting last-second bids – in online auctions. At first glance, sniping would appear to be superfluous, given the second-price structure of online auctions. This conclusion, however, neglects the sequential arrival of new auctions and the implicit option value of losing the current auction. The option is the expected surplus from bidding on subsequent auctions. By waiting until the last moment to bid, a bidder maximizes her information about the value of the surplus. Since a bid is binding, last-second bidding precludes her from being locked into a potentially suboptimal bid submitted earlier. To illustrate the dynamics of the model, I provide a quantitative solution for the option value.

JEL Classification: C73, D44, D83

Keywords: sniping; Internet auctions; last-second bidding; sequential auctions

*blim@unimelb.edu.au

One of the more striking characteristics of online auctions is the preponderance of bidding activity in the closing minutes of an auction. Anyone with even a casual familiarity with eBay, the dominant brand in the space, has surely witnessed an auction with little to no activity for the first six days, 23 hours, and 55 minutes, with a sudden flurry of frenzied bidding in the subsequent and final five minutes. The phenomenon is so widespread it has both a name, “sniping”, and secondary websites devoted strictly to facilitating it (e.g., esnipe.com, auctionsniper.com, just-snipe.com). Sniping has been formally documented by Wilcox (2000), Roth and Ockenfels (2002), Bajari and Hortaçsu (2003), Schindler (2003), and Gonzalez, Hasker and Sickles (2009), among others. Roth and Ockenfels (2002), for example, find that in a sample of 240 auctions, 89 (37%) have bids placed in the last minute and 29 (12%) in the last ten seconds.

On the face of it, last-second bidding is puzzling to game theorists. Internet auction sites like eBay typically run second-price auctions, with the winning bidder paying the second-highest bid plus some nominal increment. As originally conceived by Vickrey (1961), the *raison d’être* of second-price auctions is for bidders to submit their private valuations. The optimal strategy is independent of the actions of the other players. One can verify this result is robust to making the second-highest bids public, as is the case in online auctions. (See, for example, Bajari and Hortaçsu, 2004.)

The Vickrey result, however, requires certain assumptions that may not apply to real-world auctions. Specifically, standard second-price auction theory typically assumes that each bidder has an *independent private valuation* of the asset in question, a characterization which may apply to auctions for homogeneous goods like, say, consumer electronics but may not be relevant for heterogeneous goods like art, for which one’s valuation may be conditional on what others’ valuations are. There is considerably more incentive not to reveal one’s private information in an art auction than in a consumer electronics auction, given the free rider problem in the former. Indeed, when bidders have *interdependent* or *common valuations*, last-second bidding may be justified, as argued by Schindler (2003), Bajari and Hortaçsu (2003), and Hossain (2008).¹

Nevertheless, the rationale for last-second bidding when bidders have independent private valuations remains comparatively unclear. One might speculate that in bidding early, a player announces her presence in the marketplace, thereby altering

¹Hossain (2008), for example, conjectures that some bidders do not know how much they value an auction but only know whether they value it more than the current price. Such bidders learn their valuation through the bidding process, which creates incentive for better informed agents to withhold their bids until the last second.

the strategies of the other players. In the context of the standard Vickrey model, of course, this argument is not relevant. The design of second-price auctions is such that bidding one's private valuation is the strictly dominant strategy. When the second-highest bid is public during the auction period, last-second bidding is certainly optimal if one believes other bidders play retaliatory strategies. If it were the case that another bidder's strategy is always to outbid the prevailing high bid, last-second bidding would be the best response, but this scenario simply substitutes the original question for another: why would a bidder would play such a win-at-all-costs strategy?

To date, there is only a limited literature which has sought to explain the occurrence of sniping in auctions for goods without common valuations. Roth and Ockenfels (2002) offer a pair of possibilities: one, that last-second bidding is the best response to inexperienced bidders who play as if the format were a standard first-price English auction; two, that last-second bidding represents tacit collusion by experienced bidders, who keep the final price low in the process. The latter possibility is explored more rigorously in Ockenfels and Roth (2006), in which the authors note that very late bids have a positive probability of not being successfully submitted, and this opens a way for bidders to implicitly collude and avoid bidding wars. Ely and Hossain (2009) test the efficacy of early versus late bidding – squatting versus sniping, in their terminology – in an experimental setting. Finding that sniping leads to statistically but not necessarily economically significant increases in surplus, they hypothesize that squatting discourages competition by announcing the bidder as an interested party in the given auction, while sniping attenuates the possibility of an escalating bidding war. Barbara and Bracht (2006) argue sniping is rational since early bidders can retract bids prior to the auction's close.

What to this point has been overlooked is the sequential and random arrival of new auctions. Internet auction sites typically trade in homogeneous or near-homogeneous goods. While true homogeneity in online auctions is rare given the high proportion of used items, near-homogeneity is ubiquitous. A bidder comes to eBay looking for a television or a bike or an MP3 player. She may or may not be flexible in terms of the degree of heterogeneity she is willing to accept, but it seems reasonable enough to assume that the average bidder does not come to an auction site to bid on one specific auction, upon whose conclusion she will exit the market, win or lose. More likely, a bidder comes to an auction site to search for multiple auctions, any one of which would suit her need. She may periodically revisit the site to update her search, identifying newly opened auctions which fit her criteria. Moreover, she has some deadline by which she wants to have won an auction.

In this paper, I argue that insofar as the preceding characterization is correct, last-second bidding is the optimal strategy for a sufficiently patient bidder. The intuition is based on two simple observations. First, a patient bidder knows that if she loses the current auction, she will have the opportunity to bid on a subsequent one. With sequential auctions, there is an option value to losing an auction. Second, by waiting until the last moment to bid on an auction, a bidder will maximize her information about the value of that option. The value is conditional on how many subsequent auctions she might bid and the surplus she would extract if she were to bid. As future auctions get bid up, the option value decreases; as additional auctions open, the option value increases. If she had bid early, her preferred bid just prior to the auction's close may be below what she has already submitted. Since she can only revise her bid upward, she will be locked into a higher bid than she would like. By waiting until the last second to bid, she learns more about the value of the option to lose the current auction. Her bid conditional on this updated information set is consequently at least as good as a bid conditional on any earlier information. When monitoring the auction market is costly or uncertain, the bidder may contract the bidding out to a sniper service, knowing that such a bid can be cancelled if she revisits the market and wishes to revise her bid.

Though the option value is in practice subjective, I provide an analytic solution. According to this solution, the bidder's preferred bid is increasing with each subsequent auction. On early auctions, she prefers to bid relatively small amounts in an attempt to maximize the surplus, knowing that if she loses, she can bid on later-closing auctions. As her deadline approaches, she is willing to bid more as her future opportunities to win an auction diminish. In expectation, she will bid relatively low on the current auction and progressively higher on each future auction until she wins or her deadline expires. It follows that the more patient the bidder, the less she will bid on the current auction. In the limit, as her deadline extends to infinity, she will consistently bid the lowest possible amount that has some positive probability of being accepted. Thus, so-called "low ball" bidding is justified for sufficiently patient bidders.

The contribution of this paper is providing a rationale for last-second bidding which does not rely on common valuations, behavioral motives, or any particular institutional features² of an auction. The random arrival of new auctions implies that last-second bidding is optimal for patient bidders. The optimality is expressed with respect to a strategy not to win a particular auction but rather to maximize one's

²For example, provisions for bidders to retract bids or for sellers to offer extendable auction closing times.

expected surplus conditional on winning any auction. While the narrative is couched in terms of Internet auctions, the underlying intuition is much more generalizable. One could, for example, drop the specification that auctions are second-price and apply the key results of this paper to a housing search model, wherein a bidder may postpone placing an offer on a particular house for as long as possible while searching for substitutes on which to potentially bid. Indeed, the intuition may apply to any model of auctions involving search. I stress here that this paper is not intended to invalidate existing theories of last-second bidding – for example, the behavioral mechanism proposed by Ely and Hossain (2009) is particularly persuasive. Rather, my intent is to demonstrate that whatever the true motivation of sniping, it is in fact an optimal strategy in common circumstances.

The primary innovation is the timing component of bidding in sequential auctions. The extant literature on sequential auctions typically models the situation in which the number of auctions – and often the number of bidders – is known rather than random. Engelbrecht-Wiggans (1993) examines sequential auctions for stochastically-equivalent objects. Arora, Xu, Padman and Vogt (2004) model optimal bids in sequential closed-bid second price auctions. Wang (2006) argues that an early bid raises opponents' conditional expectation of the next auction's price, which in turn raises their bids on the current auction and lowers the bidder's surplus if she wins. With random arrival of auctions, however, bidders must consider not only how much to bid but also when to bid. In that regard, perhaps most comparable to my paper is Said (2008), who in a considerably more generalized model explores the option value of sequential auctions with random arrivals.

This paper is arranged as follows. Section 1 presents a simplified example, to guide the underlying intuition. The generalized model is given in Section 2. A solution for the option value is given in Section 3. Section 4 discusses the determination of the bidder's beliefs regarding winning bids, and Section 5 outlines possible extensions. Section 6 concludes.

1 A simple example

A bidder is looking to purchase a bike for up to \$100 within the next six days. She has identified an auction for a bike she is interested in, and this auction ends in three days. She would prefer to identify as many auctions as possible: for any auction she lost, she could potentially bid on the auction closing next. As it stands, however, she has found only one suitable auction. Should new auctions arrive, they are open for bidding for three days.

The bidder could bid some amount $x \leq 100$ immediately, but she could also wait three days. She assumes that the value of the bike to other bidders is independent of her value and that whether or not she bids on the auction does not affect the amounts other interested parties are willing to bid on that or any other similar auction. Under this assumption, there is no downside to waiting. Whether she bids today or in 3 days affects neither the probability she wins the auction with a bid of x nor the price at which she would pay if she won. By waiting three days, she could monitor the auction site for new listings, and one of two outcomes will occur: either new auctions arrive or they do not. Moreover, given the duration new auctions are open, in three days she will know exactly how many bikes she can bid on before her deadline.

Suppose, however, that she bids x today. In three days, she revisits the site and calculates her preferred bid x^* just before the auction ends conditional on whether or not new auctions for bikes she is interested in have opened. Comparing x^* to x , and bearing in mind that bids can only be revised upward, there are three possibilities. The first is $x^* > x$. She can update her bid to x^* , so no advantage was lost by bidding early. The second is $x^* = x$. In this case, the timing of her bid proved inconsequential. The last is $x^* < x$. Her preferred bid is less than x , implying she is worse off having bid three days earlier. She would be locked into a bid higher than she would prefer. Note that the preceding analysis does not require any specific calculation of x or x^* . It is based only on the fact that one of the three possibilities must occur. Given the requirement that bids be revised only upward, bidding just prior to the earliest auction's closing time weakly dominates bidding at any time earlier.

2 The model

There are a series of second-price auctions, each for a unit of a homogeneous good. An auction is uniquely identified by its closing time a . a refers interchangeably to both a specific auction and its closing time. Time t is discrete. At t , the time remaining in an open auction is $a - t$.

At the start of each period, a new auction opens with probability λ . The specification implies that sellers are not strategically timing when they are listing their auctions and that two auctions never close at the same time. A newly listed auction closes after a fixed duration D , so an auction opening at t will close at $t + D$.

Bidders may submit bids on a at any time $t < a$. Bidding is anonymous and conducted remotely, so the number of bidders interested in a is unknown. An auction

arriving at t is open for bidding at t with a minimum bid of 0. The winning bidder is revealed at time a , and she pays the second highest bid.³

At the start of each period, all prior bids on an auction are posted publicly except for the current high bid. The highest publicly posted bid on a at t , denoted $b(a, t)$, consequently represents the lower bound of the current high bid.⁴ I will refer to $b(a, t)$ as the *minimum bid* on a at t : If the bidder does not have the high bid on a , a bid submitted on a must be larger than $b(a, t)$. If she has the current high bid on a , she may only submit a bid higher than her existing bid. That is, bids may be revised upward but not downward.

The state $s(t)$ is characterized by the set of auctions $a \in \mathbf{A}(t)$ that are open for bidding at t as well as the corresponding set of minimum bids $b(a, t) \in \mathbf{B}(t)$ for those auctions. $s(t)$ can be thought of as the results of a search for auctions of this type of good at t .

2.1 Bidder Strategy at t

I will focus on the optimal strategy of a bidder i at time t , taking the behavior of all other bidders as given. By not endogenizing the behavior of her competitors, I allow for a wide range of possible bidding strategies amongst them.⁵

Bidder i has a private valuation V drawn from the interval $[0, 1]$. She has a fixed bidding horizon T : she will exit the market upon winning an auction or the time equaling T , whichever comes first. T measures i 's patience, with low realizations implying a need to obtain the item sooner than later. The distributions of T and V are independent, meaning it is not necessarily the case that if she is impatient, she also has a higher valuation. See Appendix A for a listing of all variables.

She is interested in *eligible* auctions: those which close on or before her deadline. At t , there are $N \equiv N(t)$ eligible auctions for bidder i . Formally, $\mathbf{a} = [a_1 a_2 \dots a_N]$, $\mathbf{a} \subset \mathbf{A}(t)$ represents the set of auctions for which $a_n \leq T$, and $b_n \equiv b(a_n, t)$, $\mathbf{b} = [b_1 b_2 \dots b_N]$ represents the corresponding set of minimum bids. Auctions are indexed sequentially, with $a_n < a_{n+1}$ for all n . Figure 1 presents a sample timeline.

³On Internet auctions sites like eBay, the posted bid typically represents the second-highest bid plus some nominal increment $\epsilon > 0$. For analytical clarity, I implicitly assume $\epsilon = 0$, though altering the model such that $\epsilon > 0$ does not materially affect the analysis.

⁴Given that all but the current high bid in each auction are publicly posted, a bidder knows if she has the current high bid.

⁵Though one might argue that the game should be modeled symmetrically by solving for a multiplayer equilibrium, that approach is not necessarily appropriate for auctions. Solving for such an equilibrium typically requires all bidders to play "rationally", which, though mathematically more elegant, may not necessarily be more realistic. Indeed, it may be the case that in an auction, there are players who play off-equilibrium strategies that are difficult to support analytically.




Figure 1: Sample timeline with $t = 0$, $D = 4$ and $T = 6$. There are two auctions currently open for bidding, the first closing in 2 periods and the second closing in 4. Given the bidder's deadline T , additional eligible auctions could open up at periods 1 and 2. An auction opening up after period 2 will close after T .

She assigns probability $f_n(y_n)$ to the maximum bid y_n among all other bidders on a_n at its close, where $f_n(y_n) > 0$ for all $y_n \in [0, 1]$.⁶ The corresponding cumulative distribution function $F_n(y_n)$ is equivalent to the probability she wins a_n with a bid of y_n . The exact specification of $f_n(y_n)$ is left purposefully ambiguous since it represents a subjective belief of the bidder. The only restriction I place is that the distribution of y_n is independent of bids on earlier-closing auctions.

Assumption 1. *A bid on an auction has no effect on the maximum bid among her competitors on that auction or any auction closing subsequently:*

$$f_n(y_n|x_m) = f_n(y_n) \forall n \geq m$$

Put differently, Assumption 1 states that the bidder cannot affect the likelihood that a bid of x wins an auction by bidding some x' on that auction or any earlier auction. In the absence of this restriction, one must consider potentially extremely complex conditional distributions specifying how much and when bidder i 's competitors would choose to bid as functions of how much and when i chose to bid. Later in this section, I will show that given Assumption 1, it is in fact optimal for a bidder to behave accordingly and not to condition her bid for a_n on others' bids for a_n or for any earlier auction.⁷

Assumption 2. *The bidder only submits bids on a_1 , the current auction.*

Intuitively, the sequential nature of auctions suggests that she approach the bidding problem sequentially and focus on a_1 : henceforth, the *current* auction. If she does not win a_1 , then she will focus on a_2 . If she does not win a_2 , then she

⁶Consequently, in the event of a single bid placed on an auction, the minimum bid equals 0. If no other bids are submitted, the winning bidder wins the item at 0.

⁷Note that the assumption does necessarily not rule out a bidder $j \neq i$ from playing a retaliatory strategy whereby j immediately outbids i whenever i submits a bid. Assumption 1 allows for j to play such a strategy insofar as it does not affect what j 's ultimate maximum bid on a_n will be.

will focus on a_3 , and so on. Later, I will argue that this is in fact optimal, but for now the specification will be taken as given. This restriction keeps the strategy space compact by precluding the bidder from playing strategies involving bids on later-closing auctions.

Given Assumptions 1 and 2, i 's strategy at t will involve

1. Determining whether to submit a bid on a_1 immediately or not.
2. If submitting a bid on a_1 , determining the optimal amount x_1^* to bid.⁸

I will approach each of these questions separately.

2.1.1 Optimal Amount to Bid

Assume, for the moment, that the bidder wants to submit a bid on a_1 at t . I denote the option value of losing a_1 as $\pi_2 \equiv \pi_2(s(t))$. π_2 represents the expected surplus from bidding optimally on the set of eligible auctions ending after a_1 , conditional on the state at t . Crucially, for a bidder with a sufficiently high realization of T , the set of eligible auctions includes those which are not currently open but may arrive in the future.

Conditional on submitting the bid at t , her objective function is

$$\max_{x_1} \int_{b_1}^{x_1} (V - y_1) f_1(y_1) dy_1 + [1 - F_1(x_1) + F_1(b_1(t))] \pi_2 \quad (1)$$

The first term is the expected surplus from bidding x_1 and winning the auction, while the second term is the expected surplus from losing and potentially bidding on subsequent eligible auctions. Thus, Equation (1) is the expected surplus of bidding on all eligible auctions, starting with a_1 . The resulting first order condition is

$$f_1(x_1^*)(V - x_1^*) - f_1(x_1^*)\pi_2 = 0$$

Assuming $f_n(y_n) > 0$ for all $y_n \in [0, 1]$, I have

$$x_1^* = V - \pi_2 \quad (2)$$

x_1^* is her *preferred* bid at t . By bidding $x_1^* = V - \pi_2$, she ensures that if she wins, she is at least as well off as she would be if she chose not to bid on a_1 at all. If

⁸This definition of strategy may seem too simple, given the relatively complicated auction environment (stochastic minimum bids on existing auctions and random arrivals of new auctions). One must keep in mind, however, that the bidder is repeatedly conducting this strategy at each new period.

she bids $V - \pi_2$ and wins, the maximum she pays is $V - \pi_2$, implying her minimum surplus would be $V - (V - \pi_2) = \pi_2$. Since a bid submitted at t must be greater than the prevailing minimum bid, Equation (2) implies that she would only submit her bid if $V - \pi_2 > b_1$. If $V - \pi_2 \leq b_1$, the preferred bid is lower than the minimum she is allowed to submit. Accordingly, the bidder would refrain from bidding on a_1 , since bidding some $x_1 > b_1 \geq V - \pi_2$ would yield lower expected surplus than not bidding at all.

Remark 1. The preferred bid x_1^* is independent of the current minimum bid b_1 .

Insofar as a bidder $j \neq i$ is similar to i , this result implies that no such bidder j has an incentive to bid up a_1 beyond $x_{1,j}^*$, thus providing justification for Assumption 1. If b_1 is high, it will preclude her from submitting her bid, but it will not affect the amount that she would bid if she could.

Remark 1 echoes the main result in Section 3 of Vickrey (1961). One may observe that Equation (2) is simply the objective function of a bidder in a single second-price auction in which she receives π_2 if she loses. The characteristics of the optimal bid x_1^* follow accordingly.

2.1.2 Optimal Time to Bid

Differentiation of Equation (1) yields the preferred bid at t , but one must consider whether submitting a bid at t is in fact optimal. $x_1^* > b_1$ represents a once-and-for-all bid on a_1 conditional on the state at t . That is, if the bidder were to commit herself to bidding on a_1 immediately and not revising the bid in the future regardless of whether circumstances change (though maintaining the option to bid on subsequent auctions should she lose), her optimal bid would be x_1^* . This is a highly contrived scenario, of course, but the conceit begs the question of whether there is a superior strategy to submitting x_1^* at t . Should the bidder wait until the time is closer to a_1 ?

Before addressing this question, I first place some structure on the option value π_2 . In a practical sense, just as $f_n(y_n)$ is subjective, so too is π_2 . Each bidder has her own procedure for incorporating the current state into her beliefs about the distributions of y_n , which would then be used to estimate π_2 . Nevertheless, I conjecture some common characteristics of π_2 . Insofar as my conjecture is accurate, I can then derive the optimal timing strategy.

Conjecture 1. π_2 , the option value of losing the current auction, satisfies the following conditions:

1. From t to $t + 1$, π_2 increases when an eligible auction opens and decreases otherwise.

2. π_2 is decreasing in the minimum bid on any eligible auction closing after a_1 .

The first condition captures the bidder's preference for more eligible auctions on which to bid. Implicitly, being more patient (higher T) is beneficial insofar as it allows the bidder potentially to bid on more auctions. The second condition conveys the declining option value of losing a_1 as bids on later-closing auctions increase. The higher b_2 is, for example, the less expected surplus from bidding on a_2 , implying the option value of losing a_1 decreases.

At this point, I refrain from further specifying π_2 . It is sufficient to interpret it as a subjective calculation by the bidder, conditioning on the current state. The exact manner in which she calculates π_2 is less important than the assumption that she does calculate it and that it matches the generic qualities listed in Conjecture 1. In Section 3, I explicitly solve for π_2 under a set of assumptions, though the intent there is more to corroborate mathematically the current narrative rather than to provide a literal quantitative bid strategy.

Having proposed some dynamics for π_2 , I can now solve for the optimal time to bid. Recall that x_1^* is calculated on the assumption that the bidder wants to submit at bid a t . But she must determine whether there is some point $\tilde{t} > t$ in the future such that bidding optimally at \tilde{t} will be preferable to submitting x_1^* at t . Assuming $x_1^* > b_1$, she must consider two broad contingencies. The first is that a_1 closes immediately: $t = a_1 - 1$. In this case, the bidder cannot postpone the decision of whether or not to bid on a_1 . If she doesn't bid, she expects a surplus of π_2 from bidding on the remaining auctions. If she does bid, she knows bidding x_1^* will yield at least the same surplus as not bidding. Hence, she should submit x_1^* if $V - \pi_2 > b_1$ and not bid otherwise. This is a *last-second bid*.

The second and more interesting case is that a_1 does not close immediately: $t < a_1 - 1$. In the interval between t and a_1 , the option value π_2 will change as new information is revealed. It may change because new auctions open. It may change because later-closing auctions get bid up. The interval of time between t and $a_1 - 1$ reveals information to the bidder and, in doing so, alters her calculation of the value of losing auction a_1 . Let $t_1 \equiv a_1 - 1$ represent the last period before the current auction closes, and let $\tilde{\pi}_2 \equiv \pi_2(s(t_1))$ represent the option value at t_1 . The corresponding optimal bid \tilde{x}_1^* is given by

$$\tilde{x}_1^* = V - \tilde{\pi}_2$$

The bidder knows ex ante that at t_1 , she will want to bid \tilde{x}_1^* no matter whether she had bid some amount previously or not. The question is whether she could affect

$f_n(\cdot)$ or $s(t_1)$ (and therefore $\tilde{\pi}_2$) with an earlier bid on a_1 .

Proposition 1. *If π_2 satisfies Conjecture 1, the weakly dominant strategy is to bid on a_1 when $t = a_1 - 1$ and $V - \pi_2 > b_1$.*

Proof. For Proposition 1 to hold, the bidder cannot advantageously affect $s(t_1)$ or $f_n(\cdot)$ through strategic early bidding. According to Conjecture 1, a more advantageous $s(t_1)$ would have either more eligible auctions or lower bids on later-closing auctions (or both). Neither is affected by an early bid on a_1 . The number of eligible auctions at t_1 is governed by the arrival rate λ and is assumed independent of bidding. Similarly, by Assumption 1, the distribution of bids on any auction is independent of her bidding on a_1 . The state at t_1 is inevitable, implying so too is the corresponding option value $\tilde{\pi}_2$ and preferred bid \tilde{x}_1^* . If she submitted some bid x_1 prior to t_1 , should this bid turn out to be more than \tilde{x}_1^* , she will be locked into a sub-optimal bid. \square

If \mathbf{S}_1 represents the set of all possible states at t_1 , the only undominated bids within the continuum of possible bids placed at $t < a_1 - 1$ are those less than $\min_{s \in \mathbf{S}_1} V - \pi_2(s)$. Though an undominated strategy would be to bid $\min_{s \in \mathbf{S}_1} V - \pi_2(s)$ any earlier period, this repeated bidding is superfluous and would not generate any additional surplus than simply bidding at the last second. Proposition 1 implies that as a practical strategy to implement, last-second bidding is optimal.

One can frame Proposition 1 in terms of the example in Section 1. A bidder identifies a set of auctions on which to bid, with the possibility of that set increasing in size as time passes and new auctions open. She must decide a) whether or not to bid on the auction closing soonest and b) how much, if any, to bid. If the earliest auction does not close for a long time – in this case, three days – there is no benefit to bidding now. Over the next three days, new auctions may arrive, and new bids may be placed on later-closing auctions. Her estimate of the option value of losing the current auction will be updated. Whatever bid she might place initially is unlikely to be the bid she would place in three days, given the additional information she will have at that date. Moreover, since her bid on the auction should not affect the strategies of her competitors for that auction (see Remark 1), she cannot affect the probability of winning the auction with a bid of x by bidding early.

In light of Proposition 1, let us reconsider Assumption 2, which stated that the bidder only considers bidding on a_1 . Suppose the bidder is considering a bid on a later-closing auction a_n . She would do so if it would either improve her expected surplus from eventually bidding on a_1 or if winning a_n with an early bid would yield higher expected surplus than bidding on a_1 . Neither is true. Submitting a bid on

a later-closing auction a_n is not optimal as a strategy to win a_1 , as it increases the minimum bid b_n , which should reduce her competitors' option value of losing a_1 (according to Conjecture 1). One would expect their bids on a_1 mechanically to increase, which would alter the distribution $f_1(\cdot)$ unfavorably and reduce the bidder's expected surplus from bidding on a_1 . Similarly, submitting a bid on a later-closing auction a_n will not yield higher expected surplus than bidding optimally on a_1 . Recall that bidding optimally on a_1 ensures the surplus is at least as large as the expected surplus from bidding optimally on a_2 . Bidding optimally on a_2 should ensure the surplus will be at least as large as the expected surplus from bidding optimally on a_3 , and so forth. The expected surplus of bidding on a_n thus has an upper bound of π_2 , which is equivalently the *lower* bound of bidding optimally on a_1 .

2.1.3 Optimal Bid Strategy

The preceding analysis implies the optimal strategy at t is

1. If $t = a_1 - 1$ and $x_1^* = V - \pi_2 > b_1$, bid x_1^* .
2. Otherwise, wait one period and repeat from Step 1.

Implicit in this strategy is that if a_1 exists and $b_1 < V$, the bidder need only revisit the market at $a_1 - 1$ to potentially bid. Assuming she is notified directly if she has won (i.e., she does not have to visit the market to verify whether she has won), if she does not win she need only visit at $a_2 - 1$ if a_2 existed at $t = a_1 - 1$ or at $a_1 + D - 1$ otherwise.

Three examples are illustrated in Figure 2. In Panel (a), the bidder should revisit the market at $t = 1$ to potentially bid on a_1 . If she does not win, she should revisit the market at $t = a_2 - 1 = 3$. In Panel (b), the bidder should bid on a_1 immediately. If she does not win a_1 , she should revisit the market at $t = a_2 - 1 = 2$. In Panel (c), the bidder should bid on a_1 immediately. If she does not win a_1 , she should revisit the market at $t = 4$, given that the earliest a subsequent auction could close is $a_2 = 5$.

3 A Solution for π_2

In this section I provide a quantitative solution for π_2 under a set of modest assumptions.⁹ Given its computational complexity, the analysis is intended not as

⁹As mentioned previously, the option value π_2 is a subjective calculation by the bidder. To place any explicit structure on it, then, is to offer merely one possibility of how a bidder might




Figure 2: Panels (a), (b), and (c) illustrate optimal bid timing for three different scenarios, all with $D = 4$. The solid arrow in Panel (a) indicates the bidder should optimally revisit the market at that time, while the dashed arrows in all three panels indicate the bidder should revisit the market at the indicated time only if she does not win a_1 .

a realistic solution for bidders to implement but rather as providing mathematical intuition for Conjecture 1 in the preceding section.

π_2 can be solved using dynamic programming. Let Equation (1) be defined as π_1 , the expected surplus of bidding on a_1 . π_1 is a function of π_2 , the expected surplus of bidding on a_2 . Similarly, π_2 is a function of π_3 , the expected surplus of bidding on a_3 . π_3 is a function of π_4 , and so forth. In this manner, one can iterate forward through successive auctions, even those that may not yet exist, calculating the surpluses and option values of each (potential) auction.

Given the uncertainty about whether additional auctions will arrive in the future and what the bids on any current or future auctions will be, the bidder would ideally like to be clairvoyant. She would like to somehow know in advance the sequence of auctions on which she could bid between now and T as well as the bids her competitors will place on each auction. She could pick the auction which will have the lowest maximum bid y amongst all other bidders and, assuming $y < V$, bid some amount $x \in (y, V)$ on that auction. Failing this, she would settle for knowing only the sequence of auctions but not the bids that will be placed on them. For each auction a_n , she could integrate over $f_n(\cdot)$ to determine the option value of losing, which in turn would yield her optimal bid on a_1 . When $t \geq T - D$, she knows this sequence, as no new eligible auctions can arrive before her deadline. When $t < T - D$, however, she does not know this sequence, as it has not been fully realized. New eligible auctions can arrive between $t + 1$ and $T - D$. In this case, the best the bidder can do is to project forward how she would bid on later-closing

calculate it. Different bidders could have different calculations, based in part on their individual endowments (T and V), their beliefs (the distributions $f_n(y_n)$), and the state s .

auctions if certain events were to unfold.

Specifically, in calculating π_2 when $t < T - D$, the bidder can iterate over potential future timelines. This process involves mapping out all possible permutations of auction arrivals and calculating the associated expected surpluses, conditional on $f_n(\cdot)$. What makes $f_n(\cdot)$ potentially thorny to implement is that rational bidders realize the distribution of y_n may be state-dependent. If, for example, there is an unusual glut of auction openings, the distribution of bids on each auction is likely to be pushed downward to account for the decreased competition. Similarly, if the minimum bids on all open auctions are higher than normal, the distribution of bids on each auction is likely to be pushed upward to account for the increased competition. In the interest of expositional clarity, I make the simplifying assumption that $f_n(\cdot)$ is state-independent:

Assumption 3. $f_n(y_n|s) = f(y)$ for all n and s .

Under Assumption 3, the bidder acts as though the state does not affect the distribution of bids on any given auction. A practical interpretation would be that under “average” conditions, however defined, she believes the distribution of maximum bids across auctions for the same good is fixed. When this assumption may fail is when conditions are not average, as may be the case with an unusual influx of open auctions or an abnormal frenzy of bidding on later-closing auctions. In any case, Assumption 3 is not necessary for my solution to π_2 . One could use state-dependent beliefs to calculate π_2 and yield similar results. Without the assumption, however, one must conceptualize what it would mean for $f_n(\cdot|s)$ to potentially vary over every possible state s , a (mental) task which adds complexity but not necessarily depth to the underlying intuition.

As constituted, the game corresponds to one in which a bidder i needs to win an auction by T , auctions arrive exogenously according to some process λ , and bids by her competitors are unaffected by her bidding, with their collective maximum bid on each auction conforming to some distribution $f(\cdot)$. Bidders on an auction are not adjusting their bids as functions of what others have bid on it, nor are sellers arriving to the market as a function of what the current bids are. A narrative for her competitors’ bidding is implied by Result 1, while a narrative for auction arrivals is that sellers provide new listings based on liquidity needs and not as a reaction to the bidding observed in the market.

The process for solving for π_2 will depend on the current time relative to the bidder’s deadline T and the duration of newly opened auctions D . When $t \geq T - D$, the iterative procedure to solve for π_2 is straightforward since the set of auctions on which she can bid is not subject to change. The bidder knows with certainty the

sequence of auctions on which she may potentially bid, and the only uncertainty is what other bidders will bid on those auctions. When $t < T - D$, the procedure becomes considerably more complex since the set of auctions on which she can bid prior to her deadline is random. In this case, she must iterate over all possible permutations of auction openings that could evolve from the current state. I will consider each of these two cases ($t \geq T - D$ and $t < T - D$) separately. In order to distinguish it from a generic π_2 , I will refer to my proposed solution as π_2^* .

3.1 π_2^* when $t \geq T - D$

Let π_n^* represent the expected surplus of bidding on a_n , including the option value of losing a_n and bidding on any subsequent auction:

$$\pi_n^* = \max_{x_n} \int_{b_n}^{x_n} (V - y)f(y)dy + [1 - F(x_n) + F(b_n)] \pi_{n+1}^* \quad (3)$$

Equation (3) is the auction a_n analog of Equation (1) for a_1 . From the bidder's point of view, π_n^* is equivalent to removing auctions a_1 to a_{n-1} from consideration and calculating the expected surplus of bidding on the remaining auctions. The resulting first-order condition yields a familiar optimal bid:

$$x_n^* = V - \pi_{n+1}^*$$

x_n^* represents the hypothetical bid the bidder would place on a_n in the event that she does not win any auction prior. The option value π_{n+1}^* is equivalent to the expected surplus from bidding optimally on the next auction a_{n+1} and all subsequent eligible auctions. When $t + D \geq T$, there is no uncertainty about the number of auctions on which she can bid, so the last auction on which she might potentially bid is a_N . $\pi_{N+1}^* \equiv 0$ since there are neither eligible auctions after a_N nor the possibility of such auctions opening. For $t \geq T - D$, the solution for π_n is given by

- $\pi_{N+1}^* = 0$
- if $V - \pi_{n+1}^* > b_n$

$$\pi_n^* = \int_{b_n}^{x_n^*} (V - y)f(y)dy + [1 - F(x_n^*) + F(b_n)] \pi_{n+1}^*$$

where $x_n^* = V - \pi_{n+1}^*$. Otherwise $\pi_n^* = \pi_{n+1}^*$

The solution yields a preferred bid of V on the last auction a_N , exactly the Vickrey solution for a single auction. The reader may have anticipated this result: bidding




Figure 3: Timeline for Example 1, with $D = 3$ and $T = 3$.

on a_N is conditional on losing all previous auctions and having no subsequent auctions on which to bid. With only one auction to bid on, the game collapses into the canonical game in Vickrey (1961). Consistent with the existing literature on sequential auctions (see, for example, Milgrom and Weber, 2000), the solution implies that her preferred bid is in expectation increasing after each successive auction ($x_1^* \leq x_2^* \leq \dots \leq x_{N^*}^*$).

Example 1. The bidder's deadline (T) is in 3 periods, and newly-listed auctions close (D) in 3 periods. There are currently two eligible auctions, with the first (a_1) closing in one period and the second (a_2) closing in three periods. The timeline is illustrated in Figure 3. To keep the intuition as simple as possible, I set the current time equal to 0, all minimum bids b_n equal to 0, and $f(y)$ equal to 1 for all y .

Given $D = 3$, a_1 and a_2 are the only two auctions on which she will be able to bid before her deadline. Her optimal bid on a_2 would be V , since if she is still in the market at $t > 1$, she will have only that one auction on which to bid. Knowing this, she can calculate the expected surplus of winning a_2 :

$$\pi_2^* = \int_{b_2}^V (V - y)f(y)dy = \frac{V^2}{2}$$

This determines her the optimal bid on a_1 , which she submits immediately:

$$x_1^* = V - \pi_2^* = V - \frac{V^2}{2}$$

3.2 π_2^* when $t < T - D$

When $t < T - D$, the bidder is uncertain about the set of auctions on which she could potentially bid before her deadline T . After each successive unit of time, either an auction arrives or it does not. She can map out the potential permutations of auction arrivals over the interval t to $T - D$ and iteratively calculate the expected surplus from bidding on each auction, conditional on the state that is realized.

The arrival process governed by λ allows me to model the evolution of states as a binomial tree, with each node corresponding to a unique state. Though the set of future states is technically infinite, since for every auction there is a continuum of minimum bids that could exist at a given point in the future, I can abstract from potential changes in minimum bids since their effect on π_2^* will largely be captured by $f(y)$. Changes in minimum bids will affect the expected surplus from winning a given auction, but the expected surplus is calculated using $f(y)$, which already assumes changes in minimum bids will occur.

The tree is not recombining. An auction arrival followed by a non-arrival may generate a different strategy than a non-arrival followed by an arrival. Suppose, for example, that over a 3 period interval, a single auction arrives. If the bidder's deadline is in 7 periods, her optimal strategy may be different if the auction arrived in the first period versus if it arrived in the third period.

Figure 4 illustrates the tree over four periods, with the current time set to 0. "Up" branches correspond to an auction arriving at t and "down" branches correspond to no auction arriving. The initial state is s_0 and contains all the information currently available to the bidder. One period later, the state will be either s_1 if an auction opens or s_2 otherwise. The state s_5 , for example, is the state s_0 updated with an auction opening at $t = 2$ (and closing at $2 + D$). The non-recombining nature of the binomial tree means that future states are path-dependent. A state with successive states is a *parent*, and a state descended from an existing state is a *child*. s_0 is a parent of s_1 and s_2 , while s_8 is a child of s_3 . A state s at some t prior to $T - D$ has two children, denoted as s^+ if an auction opens in the next period and s^- if an auction does not open. From s , the probabilities of reaching s^+ and s^- are λ and $1 - \lambda$, respectively. Each potential state s is associated with a state-time $t(s)$. For example, $t(s_0)$ is 0, while $t(s_1)$ and $t(s_2)$ are 1, and so on.

The tree has $T - D - t$ terminal nodes, each corresponding to a unique state. A *terminal state* \bar{s} has state-time $t(\bar{s}) = T - D$. A terminal state is significant in that, upon being realized, there is no further uncertainty regarding the set of auctions on which to bid. At a terminal state \bar{s} , the solution for the optimal bids and expected surpluses of the remaining auctions follow exactly as in Section 3.1, but with $N^*(\bar{s})$ indexing the last eligible auction.

Given that different sets of auction will exist across different states, I use subscript $n|s$ notation to indicate the n th auction conditional on state s being reached. For example, $a_{3|s_3}$ represents the auction closing 3rd from now if an auction were to open in both of the next two periods, while $\pi_{4|s_3}^*$ is the option value of losing that auction.




Figure 4: Binomial Tree when $t < T - D$. s_0 represents the current state, with the subsequent nodes corresponding to potential future states. Up nodes indicate an auction opens at t , while down nodes indicate an auction did not open at t .

I can now provide a generic solution for π_2^* , having defined a nomenclature for the evolution and timing of states. At the current time t , π_2^* is equivalent to $\pi_{2|s_0}^*$. s_0 is, by assumption, a non-terminal state, so I will initially focus on a generic solution for $\pi_{n|s}^*$ in non-terminal states. In conceptualizing how to calculate $\pi_{n|s}^*$, the bidder must consider whether an n th auction exists at state s and what her bidding strategy on that auction would be if it existed. If the auction does not exist, she can simply restate $\pi_{n|s}^*$ as a probability-weighted average of the expected surplus of bidding on the n th auction next period:

$$\pi_{n|s}^* = \lambda \pi_{n|s^+}^* + (1 - \lambda) \pi_{n|s^-}^* \quad (4)$$

If the auction exists and was to close immediately ($t(s) = a_{n|s} - 1$), her objective function would be

$$\pi_{n|s}^* = \max_{x_{n|s}} \int_{b_{n|s}}^{x_{n|s}} (V - y) f(y) dy + [1 - F(x_{n|s}) + F(b_{n|s})] \pi_{n+1|s}^* \quad (5)$$

which yields the state-dependent equivalent of Equation (2): $x_{n|s}^* = V - \pi_{n+1|s}^*$ if

$V - \pi_{n+1|s}^* > b_{n|s}$ and no bid otherwise.

If the auction exists at s but was not closing immediately, she must decide whether she would bid immediately or to wait until the last second. I will assume the latter: that in calculating expected surplus on future auctions, the bidder anticipates waiting until the last second to bid on any auction.

Assumption 4. *Given a non-terminal state s , an associated state-time $t(s)$, and an auction $a_{n|s}$, the bidder anticipates postponing bidding until the last possible period.*

Mathematically, Assumption 4 implies that for $a_{n|s}$ such that $t(s) < a_{n|s} - 1$,

$$\pi_{n|s}^* = \lambda \pi_{n|s^+}^* + (1 - \lambda) \pi_{n|s^-}^*$$

One may question whether this assumption is justified. Note that the scenario is roughly a state-dependent equivalent to the scenario examined in Section 2.1.2. Proposition 1 demonstrated that last-second bidding was optimal conditional on π_2^* satisfying the conditions in Conjecture 1. In this case, at $t = a_{n|s} - 1$ there will be a single optimal bid $x_{n|s}^*$ conditional on the state that is realized at $a_{n|s} - 1$. If the bidder were to bid some $x_{n|s}$ on $a_{n|s}$ prior to $a_{n|s} - 1$, unless she correctly guessed which state is eventually realized, $x_{n|s}$ would not be optimal and she would want to revise it to $x_{n|s}^*$. By assumption, her bid on $a_{n|s}$ has no effect on the distribution $f(\cdot)$, so bidding early offers no strategic advantage to her. Thus it is logical that in projecting how she would bid on future auctions, she anticipates that she will submit only last-second bids.

Having established how the bidder anticipates bidding in the future, I can now explicitly solve for π_2^* . The binomial tree in Figure 4 may remind the reader of a similar representation for stock prices from the options pricing literature.¹⁰ In pricing options, one may use a binomial tree to map out all possible future price paths and work backward from the terminal prices to solve recursively for the option price at each prior period of time. The same approach applies in the present context. The bidder can iterate forward through future states, calculating each $\pi_{n|s}^*$ as a function of $\pi_{n+1|s'}^*$ where $t(s') > t(s)$, until she reaches the terminal states. For each terminal state \bar{s} , she can solve for the expected surplus of the last auction in that state ($a_{N^*(\bar{s})|\bar{s}}$) and then work backward through earlier states having solved for the various state-dependent option values $\pi_{n+1|s}^*$. The resulting solution is given by

- $\pi_{N^*(\bar{s})+1|\bar{s}}^* = 0$

¹⁰See Baxter and Rennie, 1996, for a particularly lucid explanation.




Figure 5: Timelines for Example 2, with $D = 3$ and $T = 4$

- if $t(s) \in \{a_{n|s} - 1, T - D\}$ and $V - \pi_{n+1|s}^* > b_{n|s}$:

$$\pi_{n|s}^* = \int_{b_{n|s}}^{x_{n|s}^*} (V - y)f(y)dy + [1 - F(x_{n|s}^*) + F(b_{n|s})]\pi_{n+1|s}^*$$

where $x_{n|s}^* = V - \pi_{n+1|s}^*$.

If $t(s) \in \{a_{n|s} - 1, T - D\}$ and $V - \pi_{n+1|s}^* \leq b_{n|s}$:

$$\pi_{n|s}^* = \pi_{n+1|s}^*$$

Otherwise:

$$\pi_{n|s}^* = \lambda\pi_{n|s+}^* + (1 - \lambda)\pi_{n|s-}^*$$

Since all future auctions have an opening bid of 0, $b_{n|s}$ equals 0 for all auctions which do not currently exist and b_n otherwise.

Example 2. $t + D + 1 = T$

Consider the same scenario as in Example 1 but with the bidder's deadline extended one period to $T = 4$. Next period, there are two possibilities: an auction opens or it does not. The former is denoted by s_1 and occurs with probability λ , while the latter is denoted by s_2 and occurs with probability $1 - \lambda$. Any auction which arrives after next period is not relevant since it closes after the bidder's deadline. The timelines are illustrated in Figure 5.

If the bidder loses a_1 , she will have either two auctions on which to bid or only one auction on which to bid. More importantly, if she loses a_1 , she will *know* whether there are two or one auction on which to bid. Consequently, the expected surplus

from bidding on a_2 is conditional on which state is reached, s_1 or s_2 . If an auction arrives at $t = 1$, she will be at s_1 ; otherwise, she will be at s_2 . Her optimal bids and expected surplus are

$$\begin{aligned}
x_{3|s_1}^* &= V \\
\pi_{3|s_1}^* &= \int_0^V (V - y)f(y)dy = \frac{V^2}{2} \\
x_{2|s_1}^* &= V - \pi_{3|s_1}^* = V - \frac{V^2}{2} \\
\pi_{2|s_1}^* &= \int_{b_2}^{V - \pi_{3|s_1}^*} (V - y)f(y)dy + [1 - F(V - \pi_{3|s_1}^*) + F(b_2)]\pi_{3|s_1}^* \\
&= V^2 - \frac{V^3}{2} + \frac{V^4}{8}
\end{aligned}$$

If an auction does not arrive at $t = 1$, her optimal bid and expected surplus are

$$\begin{aligned}
x_{2|s_2}^* &= V \\
\pi_{2|s_2}^* &= \int_{b_{2|s_1}}^V (V - y)f(y)dy = \frac{V^2}{2}
\end{aligned}$$

The option value of losing a_1 is the probability-weighted average of the conditional expected surpluses of bidding on a_2 :

$$\pi_2^* = \pi_{2|s_0}^* = \lambda\pi_{2|s_1}^* + (1 - \lambda)\pi_{2|s_2}^* \quad (6)$$

One can solve for π_2 by substituting into Equation (6). Since $t = a_1 - 1$, she should bid $x_1^* = V - \pi_2^*$ immediately.

Example 3. The previous example is conveniently constructed, as there is no timing strategy involved. At $t = 1$, there are no further auctions to consider, and the timeline is static for the remainder of the bidding horizon. But what if $T > 4$? What if, for example, $T = 5$, as in Figure 6?

In this case, if the bidder were to lose a_1 , she knows she can bid on a_2 . Furthermore, she knows that she could wait to submit a bid on a_2 until $t = 2$. By waiting, she observes whether auctions arrive at $t = 1$ and $t = 2$, and she can condition her eventual bid on that information. One can express her option value of losing a_1 as

$$\begin{aligned}
\pi_2^* &= \lambda\pi_{2|s_1}^* + (1 - \lambda)\pi_{2|s_2}^* \\
&= \lambda^2\pi_{2|s_3}^* + \lambda(1 - \lambda)\pi_{2|s_4}^* + (1 - \lambda)\lambda\pi_{2|s_5}^* + (1 - \lambda)^2\pi_{2|s_6}^*
\end{aligned} \quad (7)$$




Figure 6: Timelines for Example 5, with $D = 3$ and $T = 5$

The state-dependent optimal bids and surpluses are

$$\begin{aligned}
x_{4|s_3}^* &= x_{3|s_4}^* = x_{3|s_5}^* = x_{2|s_6}^* = V \\
\pi_{4|s_3}^* &= \pi_{3|s_4}^* = \pi_{3|s_5}^* = \pi_{2|s_6}^* = \frac{V^2}{2} \\
x_{3|s_3}^* &= x_{2|s_4}^* = x_{2|s_5}^* = V - \frac{V^2}{2} \\
\pi_{3|s_3}^* &= \pi_{2|s_4}^* = \pi_{2|s_5}^* = V^2 - \frac{V^3}{2} + \frac{V^4}{8} \\
x_{2|s_3}^* &= V - V^2 + \frac{V^3}{2} - \frac{V^4}{8} \\
\pi_{2|s_3}^* &= \frac{3}{2}V^2 - \frac{3}{2}V^3 + \frac{9}{8}V^4 - \frac{11}{16}V^5 + \frac{5}{16}V^6 - \frac{3}{32}V^7 + \frac{1}{64}V^8
\end{aligned}$$

As before, one can solve for π_2^* by substituting into Equation (7). Since $t = a_1 - 1$, she should bid $x_1^* = V - \pi_2^*$ immediately.

3.3 Comparative Statics

In this section, I establish a series of propositions which characterize π_2^* .

Proposition 2. *Under Assumptions 3 and 4, π_2^* satisfies the conditions in Conjecture 1.*

Proof. Proofs of this and all subsequent propositions are in Appendix B. \square

Recall that the principal purpose of solving for π_2^* is to corroborate Conjecture 1. Insofar as the bidder calculates π_2 according to the procedure outlined in this section, Proposition 1 implies that last-second bidding is optimal.

Proposition 3. *For an infinitely patient bidder, the option value of losing the first auction approaches the bidder's private valuation: $\pi_2 \rightarrow V$ as $T \rightarrow \infty$.*

For a bidder willing to bid some low amount repeatedly, over a long enough timeline there is a high probability that at some point the bid will be accepted, yielding the bidder high surplus (nearly) equal to her valuation. Proposition 3 thus provides justification for “low ball” bids: low bids which have little probability of being accepted.

Proposition 4. *a) π_2 is increasing with respect to V . b) If $t + D < T$, π_2 is increasing with respect to λ .*

The option value increases with the bidder's private valuation V , since a higher V increases both the probability of winning the current auction (by increasing the optimal bid x_1^*) and the expected surplus conditional on winning. Similarly, higher λ implies more future auctions on which to bid, which in turn raises the option value of losing earlier-closing auctions.

4 Determination of $f_n(\cdot)$

Nearly all the computational results from the model depend on $f_n(\cdot)$, the bidder's belief about the highest bid on a_n amongst her competitors. I have to this point been agnostic on the determination of these beliefs. This is less for convenience than for practicality: decoding $f_n(\cdot)$ analytically is nebulous, at best. I present here an attempt to solve for $f_n(\cdot)$ in order to demonstrate the futility of such an endeavor.

The realization of y_n is dependent on, among other things, the number of competitors J_n who are in the market for auction a_n . Assume for a moment that J_n is known. Amongst these other bidders, bidder i may divide them into two categories: those that are specified similarly to i and those that are not. Insofar as some bidder $j \neq i$ is similar to i , j 's bid on a_n will be a function of her patience T_j and her private valuation V_j .

Start with the extreme case where all other bidders j are specified similarly to i and where every bidder is in the market for only one period. The bidders interested in auction a_n submit their full valuations V_j , and the cumulative distribution of the maximum bid y_n is the cumulative distribution of V_j raised to the J_n th power.

This is simple enough to calculate, assuming one has a reasonable estimate of the distribution of V_j .

Now suppose that bidders who enter the market potentially remain there for more than one period. Then, on a given auction, there may be some bidders submitting V_j (if their deadlines are approaching) and some submitting $x_j < V_j$, where x_j takes into account the option value of losing the auction. These option values are contingent on their respective deadlines T_j and their respective beliefs $f_{n,j}(\cdot)$, both of which are unknown to i . i 's beliefs are thus conditional on j 's beliefs, which by symmetry are conditional on i 's beliefs, which are conditional on j 's beliefs, ad infinitum. This mobius strip of higher order beliefs may not easily be resolved analytically. If one were to then include bidders whose strategies did not conform to those of i , then not only must i factor their bids into $f_n(\cdot)$ but also must she consider how these bids factor into j 's beliefs $f_{n,j}(\cdot)$. Modeling the resultant higher-order beliefs becomes even more complex.

Perhaps most relevantly, in actuality J_n is unknown and cannot be estimated, given the anonymous nature of the bidding process. The same group of bidders will not necessarily bid on each successive auction that i is interested in. New bidders may enter the market, and existing bidders may leave it. Some of her competitors for a_1 may not have identified or even be interested in the auction a_2 . The sum total of all these complexities suggests the near-impossibility of "solving" for $f_n(\cdot)$. As a practical matter, it best left as subjectively determined by each bidder.¹¹

4.1 State-Dependent Beliefs

Even if subjective, $f_n(\cdot)$ may nevertheless be state-dependent (contrary to Assumption 3). Consider two states s' and s'' with the same state time ($t(s') = t(s'')$) such that

- There are more auctions in state s' than in s'' , or
- The minimum bids on all auctions are at least as low in s' as in s''

The bidder may speculate that $F(\cdot|s') > F(\cdot|s'')$, given the decreased competition in state s' .¹²

How would state-dependent beliefs affect the solution π_2^* in the previous section? As previously noted, Assumption 3 is not in fact necessary for the solution to remain intact. If one were to replace each $f(y)$ with $f_n(y_n|s)$ in Section 3 (and if Assumption

¹¹ $f_n(\cdot)$ may be informed, for example, by the history of winning bids

¹²In Example 3, for example, it is reasonable for the bidder to believe that $F_2(\cdot|s_3) > F_2(\cdot|s_6)$ given that at s_3 there are more auctions on which to bid than in s_6 .

1 still holds), the necessary refinement to the solution would be a redefinition of terminal states as those corresponding to a state-time of $T - 1$ instead of $T - D$. With state-dependent beliefs about the distribution of y_n , the expected surplus is evolving with the state, even if the set of auctions on which to bid is fixed. The game tree increases by a factor of 2^{D-1} , but the mechanics of the solution remain intact.

5 Extensions

For pedagogical purposes, I have kept the model as simple as possible. Every auction is for a homogeneous good and is open for bidding for a common duration D , and the bidder is able to monitor the auction market continuously. As these specifications may not necessarily apply to real-life auction markets, I consider the effects of relaxing various assumptions of the model.

5.1 Non-homogeneous Goods

I have modeled a sequence of auctions for a homogeneous good, though in practice this is not what one observes. Internet auctions for even the most often-listed items are almost invariably differentiated in one form or another: by model, condition, or location, for example. To accommodate differences across bidders in their tolerances for differentiated goods, I can introduce different types of a substitutable good, any one of which would satisfy the bidder's needs.

Suppose that the bidder is willing to accept any one of $q \in Q$ types of the good, with q corresponding to a specific type. For each type, she has a private valuation $V_q \in [L_q, H_q]$ and belief $f_q(\cdot)$ with support $[L_q, H_q]$ regarding the maximum bids by all other bidders for auctions of that type. At t , at most a single auction opens. A type q auction opens with probability λ_q , with $\sum_{q \in Q} \lambda_q \leq 1$. The bidder has type-specific utility $U_q(p)$, where p is the dollar surplus from winning the auction. The number of elements in Q indicate the degree of flexibility the bidder has. If the size of Q is small, the bidder is discriminating, accepting only a narrow range of items, while if the size of Q is large, the bidder is flexible, willing to consider an array of different items.

Denoting a_1 's type as q_1 , I can write her objective function as

$$\max_{x_1} \int_{b_1}^{x_1} U_{q_1}(V_{q_1} - y_1) f_{q_1}(y_1) dy_1 + [1 - F_{q_1}(x_1) + F_{q_1}(b_1)] \pi_2^U$$

The optimal x_1^* is given by $x_1^* = V_{q_1} - U_{q_1}^{-1}(\pi_2^U)$ where π_2^U is the option value denom-

inated in units of expected utility, not dollars. Given Proposition 4, I conjecture that the bidder's expected surplus is increasing with respect to the size of Q and, consequently, that her preferred bid on the current auction is decreasing with respect to the size of Q . The more willing the bidder is to accept substitutes, the higher the expected surplus she will extract from her bidding.

Example. Consider the example from Section 1, but now suppose that instead that there are two types of bikes, good quality and bad quality. The bidder has found a single auction: a bad quality bike with a current minimum bid of \$50 closing in three days. Bad quality bikes typically sell for anywhere between \$0 and \$200, while good quality bikes typically sell for \$500 to \$1000. She is willing to pay up \$100 for the former and \$800 for the latter. Her utility function for each type is given by $U_q(p) = p/(H_q - L_q)$.

On a given day, there is equal probability that an auction for a good quality bike arrives, a bad quality bike arrives, and no bike arrives (1/3 each). She could bid $x_1 \leq V_{bad} = 100$ immediately, but alternatively she could wait for three days and revisit the auction site just before the bad quality bike auction closes. If she were to bid immediately, she runs the risk of winning and extracting at most $50/200 = 0.25$ units of utility.

Suppose that the bidder submits a bid today of \$100 and revisits the site in three days, just prior to the auction's closing time. At that time, she finds

1. The auction for the bad quality bike has been bid up to \$90, and
2. Two good-quality bikes have been listed, one closing in two days and the other closing in three days.

She estimates that her expected utility from bidding on the two good-quality bikes is 0.25, implying her preferred bid on the current auction for the bad-quality bike is now $x_1^* = 100 - 0.25(200) = 50$. Her pre-existing bid of \$100, however, locks her into the auction for the bad-quality bike, for which her maximum possible surplus if she wins is now 0.05. She would prefer to have bid less on the current auction, which would have simultaneously 1) lowered the ex ante probability she wins the bad-quality bike, 2) raised the surplus utility conditional on winning the bad-quality bike, and 3) raised the probability she can bid on the good-quality bikes and extract high surplus utility. Her earlier bid has proven suboptimal.

5.2 Monitoring Costs

In the optimal strategy outlined in Section 2, it is implicitly assumed that bidders can constantly monitor the auctions and bid at any time. In fact, there may be

exogenous constraints on their time, limiting their access to the auction market. To accommodate this friction, one could assign probability $\gamma_i(t)$ that the bidder will have access to the market at t . With uncertain monitoring, a value can be attached to the ability to submit a non-binding bid.

Again, consider the example from Section 1, but suppose the bidder does not know today if she will be available to revisit the auction market just prior to the current auction's close to determine if new auctions have arrived. She would be willing to bid some relatively high amount x assuming new auctions don't arrive in the next three days. As she is uncertain whether she will revisit the market, she may contract out the bidding to a sniper service, which will submit a last-second bid on behalf of its client. An early high bid effectively handcuffs the bidder from recalibrating her bidding strategy if new auctions open. If $\gamma_i(t = a_1 - 1)$ is low enough, she can bid x through the sniper service, knowing that the sniping contract is non-binding and can be amended if she is able to revisit the market at some point prior to a_1 and incorporate new information into an updated bid.

5.3 Search Intensity

As originally specified, the arrival rate λ would appear to be a common arrival rate. Some bidders, though, may search with higher intensity than others. This could be modeled by simply recasting the arrival rate as $\lambda(\sigma_i)$, where σ_i is bidder i 's search intensity. One might presume the arrival probability is increasing and concave with an upper limit less than or equal to 1. The option value of losing an auction for a high search intensity bidder would be high relative to that for a low-search intensity bidder, since in expectation the former would have more subsequent auctions on which she could bid.

5.4 Discounting

For some goods, it may be that winning an auction sooner is more valuable than winning one later. In this case, one may want to include a discount factor to capture the time value of money and/or intertemporal preferences. The discount factor would be applied in the calculation of the option value so as to lower the expected surplus of later-closing auctions. For example, given a continuously compounded discount rate r , Equation (5) becomes:

$$\pi_{n|s} = \max_{x_{n|s}} \int_{b_{n|s}}^{x_{n|s}} (V - y) f(y) e^{-r(a_{n|s} - t)} dy + [1 - F(x_{n|s}) + F(b_{n|s})] \pi_{n+1|s}$$

where r represents the discount rate. Assuming $r > 0$, discounting decreases the expected surplus of future auctions, which then increases the optimal bid x_1^* . The inclusion of r implies that the bidder's patience has two dimensions. T represents her willingness to stay in the market, while r represents her preference for earlier-closing auctions.

5.5 Multiple values of D

I have assumed all auctions are open for the same duration D . In practice, it may be the case that sellers can choose from a variety of durations. eBay, for example, allows sellers to choose 1-, 3-, 5-, 7-, and 10-day auctions. The model may easily be amended to allow multiple values of D , each with a different arrival probability λ_d . Assuming that the arrival processes are such that two auctions never close at the exact same time, I conjecture the results would remain qualitatively similar as with a single D . The game tree in Section 3.2 would have more branches at every node, though the iterative procedure to solve for π_2^* would not change.

5.6 Sealed Bid Auctions

The model can easily be amended to consider sealed bid second-price auctions. In particular, if the minimum bids are not posted, \mathbf{b} is simply the zero vector. The calculations of π_2^* and x_1^* remain intact with $b(a, t) = 0$ for all a and t .

5.7 Risk-averse bidders

In maximizing expected surplus, bidders are implicitly risk-neutral. This specification may not be controversial in the context of Internet auctions for small-value goods like bikes or television, but it is less tenable for big-value items like housing. To accommodate such scenarios, one could easily recast bidders as risk averse. Conceptually, this is a simple transformation: conditional on winning a single auction, bidders maximize their expected utility rather than expected surplus. For example, the objective function in Equation (1) becomes

$$\max_{x_1} \int_{b_1}^{x_1} U(V - y_1) f_1(y_1) dy_1 + [1 - F_1(x_1) + F_1(b_1)] \pi_2^U$$

where $U(\cdot)$ is her utility function and π_2^U is the expected utility from bidding on all subsequent auctions. Unfortunately, one may quickly observe the analytic problem that arises, in that integrating over the product of $U(\cdot)$ and $f(\cdot)$ may not have a closed-form solution.

6 Conclusion

Taking both the arrival of new auctions and the bidding of other participants as exogenous, I demonstrate that last-second bidding is the optimal strategy if the bidder is sufficiently patient. Losing an auction has an option value, as a bidder can potentially bid on subsequent auctions. A bidder should postpone submitting a bid until as late as possible, as this will allow her to better estimate the option value of losing. The contracting of sniping services is shown to be optimal if the bidder is uncertain whether or not she could revisit the market just prior to the current auction's close.

A solution for the option value is provided. According to this solution, the bidder's preferred bids on each subsequent auction are increasing, with a maximum of her private valuation on the final auction. The more patient the bidder is, the less she prefers to bid on any auction. When her deadline extends to infinity, her optimal strategy is to consistently submit the lowest possible bid, as this will ensure the maximum surplus if and when she wins.

I stress here that this paper is intended not as *the* explanation for observed last-second bidding in online auctions but as *an* explanation for it. I acknowledge the possibility and indeed likelihood that one or many of the previously theorized explanations are in fact correct, particularly those that suggest some sort of behavioral mechanism in play. One may interpret this paper as providing a rationale for why such strategies may nevertheless be optimal.

Appendix

A Summary of Key Variables

Variable	Description
bidder- and time-invariant	
D	Duration of newly arrived auction is open
λ	Probability a new auction opens in a given period
bidder-invariant, time-specific	
t	Current time
$\mathbf{A}(t)$	set of all auctions open at t
$\mathbf{B}(t)$	set of minimum bids for all auctions in $\mathbf{A}(t)$
$s \equiv s(t)$	state at t
bidder-specific, time-invariant	
T	bidder's deadline (time constraint)
V	bidder's private valuation
bidder- and time-specific	
a_n	Unique closing time of the eligible auction ending n th from the current time
$\mathbf{a} = [a_1 a_2 \dots a_N]$	Vector of eligible auctions (i.e., those closing before T). Elements are ordered chronologically by closing time
$b_n \equiv b(a_n, t)$	The current minimum (second highest) bid on auction a_n
$\mathbf{b} = [b_1 b_2 \dots b_N]$	Vector of minimum bids on eligible auctions
x_1^*	Optimal bid on a_1 , the current auction
y_n	Highest bid on a_n amongst all other bidders
$f_n(y_n)$ and $F_n(y_n)$	bidder's beliefs about the p.d.f. and c.d.f. of y_n , respectively
$\pi_2 \equiv \pi_2(s(t))$	Option value of losing a_1 , conditional on the state. Equivalently, the expected surplus from optimally bidding on a_2 (and all subsequent eligible auctions)
bidder- and state-specific	
$a_{n s}$	The n th auction at state s
$\pi_{n s}$	The expected surplus from bidding optimally on $a_{n s}$
$t(s)$	The state-time associated with s

B Proofs

Proof of Proposition 2. Condition 1 corresponds to the inequalities

$$\pi_{2|s}^* < \pi_{2|s^+}^* \quad \text{and} \quad \pi_{2|s}^* > \pi_{2|s^-}^* \quad (8)$$

evaluated at $s = s_0$. Since $t(s_0) < a_{2|s_0} - 1$, I have

$$\pi_{2|s}^* = \lambda \pi_{2|s^+}^* + (1 - \lambda) \pi_{2|s^-}^*$$

Satisfaction of the inequalities in (8) requires

$$\begin{aligned} \lambda \pi_{2|s^+}^* + (1 - \lambda) \pi_{2|s^-}^* &< \pi_{2|s^+}^* \\ \pi_{2|s^+}^* &> \pi_{2|s^-}^* \end{aligned}$$

$\pi_{2|s^+}^* > \pi_{2|s^-}^*$ is true as s^+ and s^- are identical except s^+ has an additional auction closing at $t(s) + 1 + D$.

Condition 2 corresponds to the inequality

$$d\pi_2^*/db_n \leq 0 \quad \forall n > 1 \quad (9)$$

One can show that $d\pi_{2|s}^*/d\pi_{n|s}^* > 0$ for all $n > 2$, implying that the sign of $d\pi_{n|s}^*/db_n$ determines the sign of (9). Differentiation of Equation (5) yields

$$d\pi_{n|s}^*/db_n = \begin{cases} -(V - b_n)f(b_n) \cdot 1 < 0 & \text{if } x_{n|s}^* > b_n \\ 0 & \text{otherwise} \end{cases}$$

Since $d\pi_{n|s}^*/db_n \leq 0$, it must be the case that $d\pi_2^*/db_n \leq 0$. \square

Proof of Proposition 3. Let $Pr[\epsilon]$ denote the probability of winning a single auction with a fixed bid of ϵ . The condition $\pi_2 \rightarrow V$ is equivalent to $Pr[\epsilon] \rightarrow 1$ for some $\epsilon \geq 0$. Suppose that there are no auctions currently open but that until her deadline the bidder plans to bid ϵ on every auction which does arrive. If she knew there would be N auctions on which she eventually could bid, the probability that she wins one

auction with a fixed bid ϵ is

$$\begin{aligned} Pr[\epsilon] &= F(\epsilon) + [1 - F(\epsilon)](F(\epsilon) + [1 - F(\epsilon)](F(\epsilon) + [1 - F(\epsilon)](\dots \\ &= \sum_{n=1}^N F(\epsilon) [1 - F(\epsilon)]^{n-1} \end{aligned} \quad (10)$$

As $T \rightarrow \infty$, it must be the case that $N \rightarrow \infty$ as well. As $N \rightarrow \infty$, $Pr[\epsilon] \rightarrow F(\epsilon)/F(\epsilon) = 1$, as implied by $F(\epsilon) \in [0, 1]$ and Equation (10). \square

Proof of Proposition 4. Proofs by induction: In order to show $d\pi_2^*/dV > 0$, it is sufficient to demonstrate $d\pi_{n|s}^*/dV \geq 0$ for all $n > 1$ and all s . For any terminal state \bar{s} and associated final auction $N^*(\bar{s})$,

$$\frac{d\pi_{N^*(\bar{s})|\bar{s}}^*}{dV} = \begin{cases} F(V) - F(b_{N^*(\bar{s})|\bar{s}}) > 0 & \text{if } V > b_{N^*(\bar{s})|\bar{s}} \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

Similarly, for any state s such that $t(s) \in \{a_{n|s} - 1, T - D\}$

$$\frac{d\pi_{n|s}^*}{dV} = \begin{cases} \overbrace{F(V - \pi_{n+1|s}^*) - F(b_n)}^{>0} + \underbrace{[1 - F(V - \pi_{n+1|s}^*) + F(b_n)]}_{>0} \frac{d\pi_{n+1|s}^*}{dV} & \text{if } V - \pi_{n+1|s}^* > b_{n|s} \\ \frac{d\pi_{n+1|s}^*}{dV} & \text{otherwise} \end{cases} \quad (12)$$

For all other s ,

$$\frac{d\pi_{n|s}^*}{dV} = \underbrace{\lambda}_{>0} \frac{d\pi_{n|s^+}^*}{dV} + \underbrace{(1 - \lambda)}_{>0} \frac{d\pi_{n|s^-}^*}{dV} \quad (13)$$

From (11), (12) and (16), one can induct that $d\pi_{n|s}^*/dV \geq 0$ for all $n > 1$ and all s .

Similarly, note that the inequality $d\pi_2^*/d\lambda > 0$ is equivalent to $d\pi_{2|s_0}^*/d\lambda > 0$. For terminal states \bar{s} ,

$$\frac{d\pi_{n|\bar{s}}^*}{d\lambda} = 0 \quad (14)$$

For non-terminal states s where $t(s) = a_n - 1$

$$\frac{d\pi_{n|s}^*}{d\lambda} = \underbrace{[1 - F(V - \pi_{n+1|s}^*) + F(b_n)]}_{>0} \frac{d\pi_{n+1|s}^*}{d\lambda} \quad (15)$$

For non-terminal states s where $t(s) < a_n - 1$

$$\frac{d\pi_{n|s}^*}{d\lambda} = \underbrace{(\pi_{n|s^+}^* - \pi_{n|s^-}^*)}_{>0 \text{ (by Prop. 2)}} + \underbrace{\lambda}_{>0} \frac{d\pi_{n|s^+}^*}{d\lambda} + \underbrace{(1-\lambda)}_{>0} \frac{d\pi_{n|s^-}^*}{d\lambda} \quad (16)$$

From (14), (15), and (16), one can induct that $d\pi_{n|s}^*/d\lambda \geq 0$ for all $n > 1$ and all s . Since s_0 is a non-terminal state and $t(s_0) = 0 < a_2 - 1$, (16) implies that $d\pi_{2|s_0}^*/d\lambda > 0$. \square

References

- [1] Arora, Ashish, Hao Xu, Rema Padman and William Vogt, 2004. "Optimal Bidding in Sequential Online Auctions." Carnegie Mellon University, working paper.
- [2] Bajari, Patrick and Ali Hortacısu, 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions" *RAND Journal of Economics* 34(2), 329-355
- [3] _____, 2004. "Economic Insights from Internet Auctions." *Journal of Economic Literature* 42(2), 457-486
- [4] Barbaro, Salvatore and Bernd Bracht, 2006. "Shilling, Squeezing, Sniping: Explaining late bidding in online second-price auctions" University of Mainz, working paper.
- [5] Baxter, Martin and Andrew Rennie, 1996. *Financial Calculus*, Cambridge University Press.
- [6] Ely, Jeffrey C. and Tanjim Hossain, 2009. "Sniping and Squatting in Auction Markets." *American Economic Journal: Microeconomics*, 1:2, 68-94
- [7] Engelbrecht-Wiggans, Richard, 1994. "Sequential auctions of stochastically equivalent objects." *Economics Letters* 44, 87-90.
- [8] Gonzalez, Raul, Kevin Hasker, and Robin C. Sickles, 2009. "An Analysis Of Strategic Behavior In eBay Auctions." *Singapore Economic Review* 54(3), 441-472
- [9] Tanjim Hossain, 2008. "Learning by bidding," *RAND Journal of Economics* Vol. 39, No. 2, 509-529
- [10] Milgrom and Weber, 2000. "A Theory of Auctions and Competitive Bidding, II," in Paul Klemperer, ed., *The Economic Theory of Auctions*, Vol. 2, Edward Elgar Publishing, 179-194.
- [11] Ockenfels, Axel and Alvin E. Roth, 2006. "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction," *Games and Economic Behavior* 55, 297-320.
- [12] Roth, Alvin E. and Axel 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* 92:4, 1093-103.

- [13] Said, Maher, 2008. "Sequential Auctions With Random Arrivals," Yale University, working paper.
- [14] Schindler, Julia , 2003. "Late Bidding on the Internet," University of Vienna, working paper.
- [15] Vickrey, William, 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance* 16(1), 8-37
- [16] Wang, Joseph Tao-yi, 2006. "Is Last Minute Bidding Bad?" Caltech, working paper.
- [17] Wilcox, Ronald T., 2000. "Experts and Amateurs: The Role of Experience in Internet Auctions," *Marketing Letters* 11(4), 363-374