

Bidding Dynamics in Auctions^{*}

Hugo Hopenhayn[†]

UCLA

Maryam Saeedi[‡]

Tepper, CMU

September 28, 2016

Abstract

This paper studies bidding dynamics where values and bidding opportunities follow an unrestricted joint Markov process, independent across agents. Bids cannot be retracted, as is frequently the case in auctions. Our main methodological contribution is that we construct a mapping from this general stochastic process into a distribution of values that is *independent* of the type of auction considered. The equilibria of a static auction with this distribution of values is used to characterize the equilibria of the dynamic auction, making this general class very tractable. As a result of the option of future rebidding, early bids are shaded and under mild conditions increase toward the end of the auction. Our results are consistent with repeated bidding and skewness of the time distribution of winning bids, two puzzling observations in dynamic auctions. As an application, we estimate the model by matching moments from eBay auctions.

Keywords: Dynamic auctions, Bid shading, Decision under uncertainty. **JEL codes:** C73, D44, D81, L81

^{*}We would like to thank John Asker, Matt Backus, Tom Blake, Alessandro Bonati, Yingni Guo, Ken Hendricks, Ali Hortacsu, Moritz Meyer-ter-Vehn, Paul Milgrom, Ariel Pakes, Michael Peters, Robert Porter, Andrea Prat, Debraj Ray, Ali Shourideh, Ennio Stacchetti, Steve Tadelis, Robert Town, Anton Tsoy, as well as seminar and conference participants at Northwestern Economics Department Theory seminar, Einaudi Institute of Economics and Finance, Midwest Economic Theory Meetings, International Industrial Organization Conference, Canadian Economic Theory Conference, The Economics of Information and Communication Technologies Conference, Society of Economic Dynamics, SITE, Dynamic Games, Contracts, and Markets Session, eBay, Econometrics Society World Congress, UCLA mini conference, Workshop in Industrial Organization (Toi 8), Econometrics Society Winter Meetings, Dartmouth Winter IO Conference, Carnegie Mellon University, and Tilburg University. We would also like to thank Xiang Hui for excellent research assistantship. Any remaining errors are ours.

[†]hopen@econ.ucla.edu

[‡]msaeedi@andrews.cmu.edu

1 Introduction

Many auctions take place over a considerable length of time. Such is the case of trading platforms (e.g., eBay) and other settings (e.g., procurement, spectrum). As a result, dynamic considerations can be important for understanding bidding behavior and considering questions of auction design. In particular, during these long auctions, bidders' valuations and strategies are likely to be affected by information that arrives during the auction; yet most of the literature has abstracted from this feature. This paper develops a general theory for bidding dynamics in long auctions that incorporates this feature as one of its main building blocks.

The relevance of modeling dynamic bidding can also be motivated from observed behavior in several dynamic auction settings.¹ For instance, the distribution of bidding times for winning bids is observed to be considerably skewed toward the end of the auction, and bidders often place more than one bid during an auction, with their bids differing considerably. Obviously, this behavior cannot be rationalized through the lens of a static model, where bidders should bid based on their valuation only once and at the time of arrival to the auction. In our theory, these two facts arise naturally as a strategic response to dynamic considerations.

Our model has two building blocks. The first one is information arrival during the auction, which changes the expected final value to bidders. This could be due to the existence of alternatives that change the outside value (modeled here exogenously), preference shocks (e.g., change of plans when buying event tickets), cost or capacity shocks (as is the case in electricity markets, see [Ito and Reguant \[2016\]](#)), information regarding complementary goods (as in the case of spectrum auctions, see [Börgers and Dustmann \[2005\]](#)), and alternative demands for use of resources in the face of capacity constraints (as in procurement auctions, see [Jofre-Bonet and Pesendorfer \[2003\]](#)). The accrual of new information has several effects. First, it gives bidders incentives to delay bidding as the estimate of final expected value becomes more precise. In the absence of any other friction, this would lead to extreme sniping, where all bidders place bids close to the end of the auction. Second, the expected final value has the martingale property and thus the distribution of bidders' value is subject to a mean preserving spread over time. Thus, the first-order statistic of the distribution of values increases as the auction progresses, making it more likely that winning bids occur toward the end of the auction.

The second building block is random arrival of bidding opportunities. This stands for frictions that might impede the precise timing of bids and information acquisition or inattention

¹This has been noted for the case of eBay but is shared by other auction settings (e.g. Amazon, procurement auctions in Russia and Brazil, auctions for CD deposits in Texas Treasury).

on the side of bidders. It is a necessary condition for bidding to take place before the end of the auction and, in particular, to explain multiple bids, two features that are frequently observed in these auctions. This assumption is shared by other researchers, such as Groeger and Miller [2015], Coey et al. [2015], Kamada and Kandori [2011, 2015], Ambrus et al. [2014]. We assume bids cannot be retracted or lowered, which is true in many settings, such as in eBay auctions, spectrum auctions, and reverse auctions.²

We focus on the case of independent private values, where signals and final values are independent across bidders and bidding times are exogenous.³ We model this through an otherwise unrestricted joint Markov process of signals and bidding times that are also sufficient statistics for the final expected value. Our process is very general, allowing, for instance, attentiveness to vary with value, and nests the case where a bidder bids at the end of the auction almost surely. At each time of arrival, bidders choose whether to place a bid or increase a previous one. At the end of the auction, the winner and payments are determined from the bids placed according to the rules of the auction. We characterize the equilibrium of this dynamic auction in the case of sealed bids. The results extend to the case of non-sealed-bid second-price auctions with general information structures. In particular, they apply to auctions with proxy bidding, such as eBay auctions.

Our main theorem shows that we can reduce the problem of finding an equilibrium in the dynamic auction to that of an equivalent static auction. This is done by deriving a distribution of values for the static auction that is *independent* of the auction type and holds for *any* auction with monotone bidding functions, and is only a function of the joint stochastic process of valuations and bidding opportunities. Then we show that the equilibrium bidding strategies in the dynamic auction can be derived from those of the equivalent static auction. We explain this result with the aid of Figure 1. The indifference curve represents all points in the value-time space where the bidder chooses to place the same bid b_0 . We characterize these indifference curves, which are independent of the type of auction, and show that to find the optimal bid at any point (v_0, t_0) .

We turn to the determination of the indifference curves. Consider the bid b_0 placed at point (v_0, t_0) . First note that this bid is relevant only as far as it becomes the bidder’s final bid. This will occur only if the agent chooses not to bid higher at any future bidding opportunity. At an extreme, if this were for certain the last bidding opportunity, the expected value would be the unconditional value v_0 .⁴ However, the possibility of rebidding is a source of

²<http://goo.gl/ebaHpR>, accessed January 2016. Withdrawing from spectrum auctions can cause a high withdraw penalty, see Cramton [2002]. <http://goo.gl/1Lfi2q>, accessed July 2016.

³Results extend to the case of pure common values, as explained in Section 10.

⁴For simplicity of exposition, we assume that process for values is a martingale, thus v_0 is also the expected

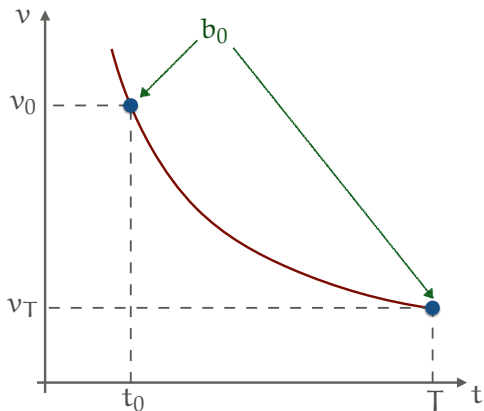


Figure 1: Indifference curves for bidding

Note: All points on this curve result in the same optimal bid value, b_0 .

adverse selection against future self, as the current bid applies when the bidder chooses not to exercise any future options of rebidding, and this is correlated with lower future value.⁵ This explains why the expected value, conditional on this bid being the final one, is lower than v_0 . Since bidders are risk neutral and bids are monotonic, it must be the case that the (conditional) expected final value associated to each point on this indifference curve must also be v_T . This property is independent of the type of auction.

The indifference curve depicted in the Figure 1 has the following self-generating property, which we explain for the point (v_0, t_0) : Starting at this point, consider all paths where the value lies below this indifference curve, at each of the corresponding future bidding opportunities. Those are the paths where the bidder will not increase the bid, so b_0 will be the final bid. The expected final value conditional on this set of paths is precisely v_T , and the same property holds for any point in this indifference curve. This property, which we call self-generated expectation, only depends on the stochastic process for values and bidding opportunities and is thus independent of the auction type.

Moreover, these paths can be used to generate the equivalent static auction. The point (v_0, t_0) is the final bidding point for all paths that, as described, remain below the indifference curve at bidding times. From the point of view of final bidding behavior, these histories can be mapped to a final value v_T . Following the same procedure for all points (v, t) determines a distribution of final values for each bidder and defines the equivalent final auction. Take now

final value at the end of the auction, conditional on the information at time t_0 .

⁵An analogue result is found in [Harris and Holmstrom \[1982\]](#), where initially workers' wages are shaded below marginal products, as the wage is effective in the future only if it is less or equal than the realized marginal product of the worker.

any equilibrium vector of bidding functions $\tilde{B}_i(v_T)$ in the final auction. This can be used to assign strategies $B_i(v, t)$ in the dynamic sealed bid auction by setting $B_i(v, t) = \tilde{B}_i(v_T)$ for the final point v_T in the corresponding indifference curve.

In our setting, bidders have an incentive to shade bids below and beyond what might be the case in a static auction. The indifference curve depicted in Figure 1 can be used again to explain shading. In particular, when placing a bid at time t_0 where the unconditional expected final value is v_0 , the bidder chooses to bid b_0 , which in the static auction corresponds to a lower value, v_T . This reflects the fact that the expected conditional value will be lower than v_0 , unless t_0 was for sure the last bidding opportunity. This source of *value shading* is common to all auctions and translates into *bid shading* according to the specific bidding function.⁶ When indifference curves are downward sloping, as occurs under mild restrictions, the expected conditional value increases toward the end of the auction. This can be seen in Figure 1, by observing that a horizontal line at a given value (e.g., v_0) will cross higher and higher indifference curves. This provides a rationale why bidders might increase their bids over time independently of competitive pressures. The intuition for this result is straightforward and can be traced back to the incentives for shading: As the end of the auction approaches and the option of rebidding becomes less likely, the adverse selection problem mentioned above and the incentives for shading tend to disappear. As the auction progresses and residual uncertainty decreases, bids near the optimal bid for the unshaded expected value. This increase in bids is the third element that contributes to the skewness of the distribution of winning bids.

Further results can be obtained for the case of second-price auctions, where bids equal the corresponding conditional expected value, such as $B(v_0, t_0) = v_T$ in the example depicted in Figure 1. While bids are shaded relative to expected values, the *expected final bid* still equals the expected final value and is thus unbiased. These two results are consistent, since bids, while shaded, are still unbiased in the set where they apply.⁷ As a comparison, an auction that lets bidders retract their bids will have no shading, since a bid only applies when no rebidding takes place or when the final expected value is the same at a new bidding time.⁸ As a result of no retraction, there is a loss of information and efficiency. We prove this by showing that efficient bidding (as in the case of bid retraction) corresponds to a mean preserving spread of values for the equivalent final auction, relative to our baseline case.

Next we demonstrate how the model can be used empirically to estimate the stochastic

⁶For example, in the case of a second-price auction, it translates one-for-one into bid shading.

⁷This observation suggests an avenue to non-parametrically identifying shading as illustrated in Section 6.5.

⁸This corresponds to case of flat indifference curves.

process for values and the final equivalent distribution using eBay data. To do so, we consider the specialized case where values follow a Brownian motion and bidding opportunities a Poisson process. We derive a partial differential equation that is used to solve for the equilibrium bidding function. Estimates of the parameters are obtained using simulated method of moments combined with Laplace-type mcmc estimators. Our estimated model explains more than 50% of the skewness in the timing of winning bids, even though this was not part of the targeted moments. Bid shading accounts for 20% of the skewness generated by the model, while the mean preserving spread effect of new information explains the rest. As a counterfactual, we consider an alternative auction design where bidders are allowed to retract bids at their bidding times, thus removing the incentives to shade. As indicated above, the bids in this scenario are a mean preserving spread of those in the baseline model. This results in a moderate increase in the average final price and total value.

Related Literature Our paper contributes to the literature on auctions and dynamic games in both theoretical and empirical strands. This section provides a brief account of the papers that are closely related to our work.

Following [Ockenfels and Roth \[2006\]](#), a series of recent papers ([Kamada and Kandori \[2011, 2015\]](#), [Ambrus et al. \[2014\]](#)) consider a class of collusive strategies for the case of open dynamic second price auctions that resemble some of the features observed in eBay auctions. As in our setting, bidders have random bidding opportunities. In contrast, values remain constant throughout the auction and in most cases are identical across bidders. Furthermore, the distribution of bidding opportunities is more simplified as it is given by a Poisson process, thus independent of values and time.⁹ [Ambrus et al. \[2014\]](#) prove the existence of a perfect public equilibrium that can involve gradual bidding, periods of inactivity and the late start of bidding in the auction. This form of collusion keeps final price lower and is sustained with reversion to the unique Nash equilibrium of a sealed bid second price auction. The strategic setting in [Ambrus et al. \[2014\]](#) can be inscribed in the class of revision games introduced by [Kamada and Kandori \[2011, 2015\]](#), that consider optimal collusive strategies. These can be characterized by a unique "most cooperative" time path that converges to the unique Nash equilibrium of their stage game. At any random time, players choose the action prescribed by this path, with the threat of Nash reversion. The specific assumptions they make on payoff functions do not strictly apply to the case of second price auctions, but if they could be extended to this setting, would also imply a uniquely defined increasing path of bids.

⁹[Ockenfels and Roth \[2006\]](#) show that in this setting bidding your own valuation is not always a dominant strategy and provide an example where in equilibrium bidders delay all bids until the last moment of the auction.

In contrast to these papers, we focus on a set of simple equilibrium strategies that do not condition bids on other bidders' previous bids. These are the only equilibrium strategies in the case of sealed bid auctions and are also an equilibrium for the case of non-sealed dynamic second price auctions with fairly arbitrary information structure. While missing the strategic elements considered in the above papers, our approach has several advantages. First, as mentioned our restriction is without loss of generality for a general class of sealed bid auctions. Second, with an eye on applications we allow for a very general stochastic structure for values and bidding times. Third, we are able to map our dynamic game into an equivalent static auction, making the derivation of equilibrium strategies very tractable.

Another important building block of our model is bidders' random valuation, that can be interpreted as the result of information arrival throughout the auction. This is a key ingredient to generate bid shading as a result of the future option to raise bids if so desired. [Compte and Jehiel \[2004\]](#) consider the case of an open english button auction, where bidders may receive finer information about their value throughout the auction. The price increases gradually and bidders decide when to drop out. This can result in bidders paying the option by staying in the auction beyond the point at which the price reaches their ex-ante expected value. This contrasts with our result of bid shading. The key difference is that in our setting bidders have always the option to rebid at any of their bidding times independently of the history of their past bids. [Compte and Jehiel \[2007\]](#) extend the analysis to the case of costly information acquisition at a cost.¹⁰

In addition to the above, there are several papers that consider models that are relevant for bidding dynamics. [Bajari and Hortacsu \[2003\]](#) incorporate a model of common-value auction to explain late bidding. [Ockenfels and Roth \[2002\]](#) argue that rational bidders on eBay delay bidding to the end of the auction as an optimal response to other behavioral types ("incremental bidders") that increase their maximum bids as they get outbid.¹¹ As shown in our empirical results, even though incremental bidding occurs in practice, the average observed increase in bids is considerably large. [Daniel and Hirshleifer \[1998\]](#) and [Avery \[1998\]](#) rationalize jump-bidding behavior in an English auction with common values in the presence of bidding costs. [Hörner and Sahuguet \[2007\]](#) consider jump bidding as strategic behavior in a costly signaling game.

To the best of our knowledge, our paper is one of the first to estimate a structural model of dynamic bidding within an auction. A notable exception is the recent paper by [Groeger and Miller \[2015\]](#), who identify a structural model for a first-price auction. The authors consider

¹⁰[Rezende \[2005\]](#) extends the setup of the previous paper to allow for different random costs for acquiring information.

¹¹[Ambrus et al. \[2014\]](#) provide equilibrium foundations for this type of behavior, as explained above.

the case where bidders have random bidding windows, similar to our own framework. As one of our extensions, we propose a method to identify bid shading non-parametrically. There is a large literature strand on non-parametric identification in auctions, following the classic papers by [Guerre et al. \[2000\]](#), [Athey and Haile \[2002\]](#), and [Haile and Tamer \[2003\]](#).

There is also a literature strand on modeling and estimating dynamics across auctions. The classic paper is [Jofre-Bonet and Pesendorfer \[2003\]](#), who estimate dynamic auctions in procurement by controlling for the utilized capacity of participants. More recent papers that consider the option value faced by bidders in sequential auctions include [Zeithammer \[2006\]](#), [Said \[2011\]](#), [Hendricks and Sorensen \[2015\]](#), [Backus and Lewis \[2012\]](#), and [Coey et al. \[2015\]](#). As a result of this option value, changes in the available alternative items can alter the reservation price for bidders over time. This literature strand motivates our reduced-form approach toward the change in valuation to be a result of changes in the outside option.

The paper is organized as follows. [Section 2](#) provides a simple example that conveys the main intuition and results in the paper. [Section 3](#) describes the general model and defines an equilibrium. [Sections 4 and 5](#) prove the existence and uniqueness of equilibrium, characterize bidding behavior, and provide the dynamic programming algorithm. [Section 6](#) gives properties for the case where values are independent of Poisson arrivals for bidding times and solves the two special cases described above. [Section 8](#) shows that the bidding function derived before is still applicable under arbitrary assumptions about the information a bidder observes on the past bids of other bidders. [Section 9](#) applies the model to eBay data and performs a counterfactual. Finally, [Section 10](#) extends the model to that of correlated signals.

2 A Simple Example

We start with a simple example in a second-price sealed-bid auction and derive some key properties of the equilibrium bidding function that extend to the general model. We then show that these results can be extended to other sealed-bid auction formats such as first-price auctions. There are two periods, $t = \{0, 1\}$. Each bidder observes a private signal of his value in the first period, v_0 and can submit a bid for sure at that time. In addition, each bidder can rebid with probability p in the second period after observing his final value v drawn from distribution $F(\cdot|v_0)$. Both signals and final values are drawn independently across bidders. [Figure 2](#) illustrates this example.

A bidding strategy specifies an initial bid $b_0(v_0)$ and a bid $B(v)$ in the second period. The latter bid will only be effective if the bidder has an opportunity to rebid in the second period

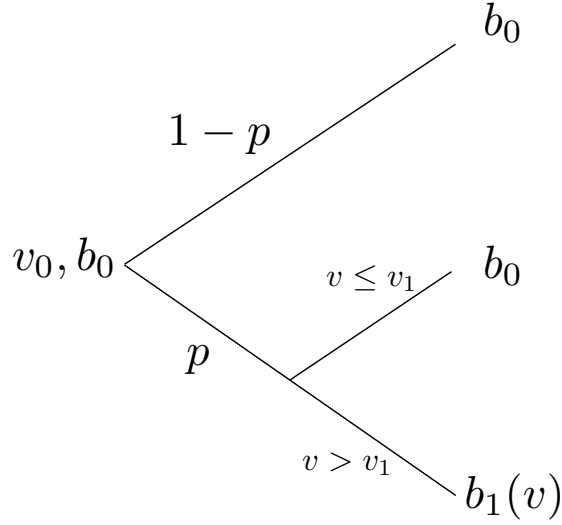


Figure 2: Bidding Tree

and $B(v) \geq b_0$. For any increasing bidding function, this determines a threshold v_1 such that $B(v) > b_0$ whenever $v > v_1$. For example, in the second-price auction where $B(v) = v$, this threshold is $v_1 = b_0$. A final distribution of bids is determined given these strategies and an initial distribution of signals $\Phi(v_0)$. In equilibrium, these bidding strategies must be best responses to the distribution of bids of the remaining players.

2.1 Second-Price Auction

We first analyze the above game in a second-price auction setting. In the second period, the optimal bid will be equal to each bidder's final value. Let b_{-i} denote the maximum final bid among all bidders, excluding i , and let G denote its CDF. In the first period, the expected value for bidder i is

$$\begin{aligned}
 U_i &= (1-p) \int_0^{b_0} (E(v|v_0) - b_{-i}) dG(b_{-i}) \\
 &+ p \int_0^{b_0} \int_0^{b_0} (v - b_{-i}) dG(b_{-i}) dF(v|v_0) + p \int_{b_0}^v \int_0^v (v - b_{-i}) dG(b_{-i}) dF(v|v_0).
 \end{aligned}$$

Differentiating with respect to b_0 and equating to zero results in

$$(1-p) (E(v|v_0) - b_0) dG(b_0) + p \int_0^{b_0} (v - b_0) dG(b_0) dF(v|v_0) = 0. \quad (1)$$

It can be easily verified that, as this equation is strictly decreasing in b_0 , there is a unique solution that satisfies

$$b_0 = \frac{(1-p)E(v|v_0) + p \int_0^{b_0} v dF(v|v_0)}{(1-p) + pF(b_0|v_0)}. \quad (2)$$

This expression has a natural interpretation: The optimal initial bid equals the expected value conditional on no rebidding in the second period. Given a current bid b_0 , the final bid prevailing at the end of the auction is then b_0 with probability $(1-p)$ and $\max(b_0, v)$ with probability p . The initial bid b_0 binds in two cases: 1) there is no opportunity to rebid and 2) there is an opportunity to rebid but $v \leq b_0$; therefore, the bidder decides not to bid; see the first two branches of the tree in Figure 2. The second term in the above equation represents an adverse selection effect. It follows that $b_0 < E(v|v_0)$ as $\int_0^{b_0} v dF(v|v_0) / F(b_0|v_0) < E(v|v_0)$. Indeed, as $p \rightarrow 1$, it is easy to see that b_0 decreases monotonically to zero.

The probability distribution \tilde{F} for the final bids of a player is given by

$$\tilde{F}(b) = \begin{cases} 0 & \text{if } b < b_0 \\ (1-p) + pF(b|v_0) & \text{if } b \geq b_0 \end{cases}, \quad (3)$$

with mean

$$E(b) = b_0 [p + (1-p)F(b_0|v_0)] + p \int_{b_0}^1 v dF(v|v_0).$$

Substituting the first term using (2), it follows that

$$E(b) = (1-p)E(v|v_0) + p \int v dF(v|v_0) = E(v|v_0);$$

therefore, in spite of first-period shading, the expected final bid is unbiased, equaling the expected final value.

2.2 Intuition

We can also interpret the above results in another way. In the first period, bidders with value v_0 bid as if their value was v_1 in the last period. The pair (v_0, v_1) can be interpreted as two values that result in the same optimal bid b_0 (refer to Figure 3). In addition, this shaded value $v_1 < v_0$ is also equal to the expected final value of the bidder conditional on having b_0 prevail at the end of auction. This result can be generalized to other auction mechanisms, as will be shown in the next section.

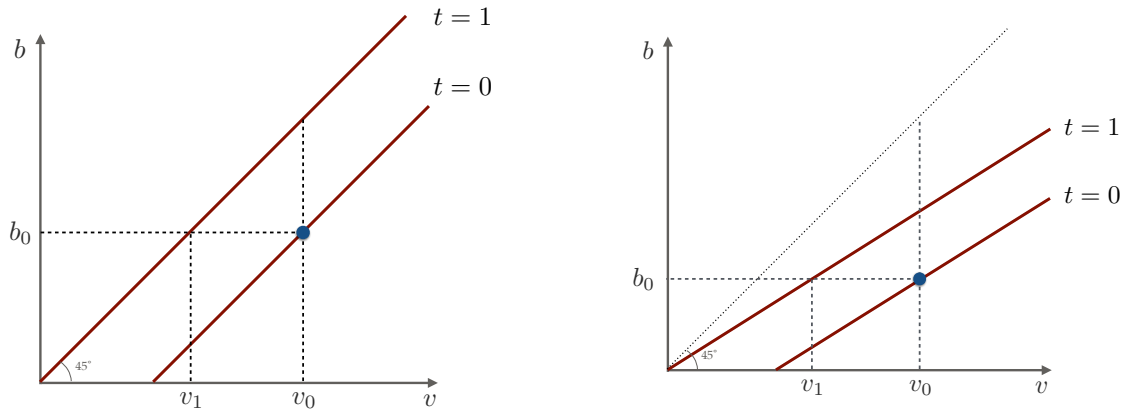


Figure 3: Second-price vs. First-price Auctions

2.3 Extension to General Auctions

Now consider an arbitrary auction. Assume (v_0, b_0) are the value and optimal bid of a bidder in the first period. Given b_0 , the optimal threshold v_1 for increasing this bid in the second period will be such that

$$B(v_1) = b_0, \quad (4)$$

where B is the bidding strategy for this auction in the last period.¹² The threshold also defines a lottery under which the bid b_0 will be the final bid, comprising the events where the agent cannot rebid or has value less than this threshold, v_1 . This lottery has the expected value

$$\frac{(1-p) E(v|v_0) + \frac{F(v_1) \int^{v_1} (v-b) dF(v|v_0)}{F(v_1)}}{1-p + pF(v_1)}.$$

Assuming the agent is risk neutral, this is indeed the value of the lottery to the agent, so it follows that

$$b_0 = B \left(\frac{(1-p) E(v|v_0) + \frac{F(v_1) \int^{v_1} (v-b) dF(v|v_0)}{F(v_1)}}{1-p + pF(v_1)} \right), \quad (5)$$

where B is again the bidding function. Using (4) and (5) and assuming B is strictly increasing, it follows that

$$v_1 = \frac{(1-p) E(v|v_0) + \frac{F(v_1) \int^{v_1} (v-b) dF(v|v_0)}{F(v_1)}}{1-p + pF(v_1)},$$

¹²Note that regardless of the equilibria, this bidding function can be interpreted as the best response in this auction to the distribution of other bidders' final bids. The only property that we use is that it be strictly increasing in value.

so this threshold is defined independently of the auction considered. Once we solve for this value (which is what we do in our example), we can calculate the first period bid by $B(b_0)$. This latter calculation may not be trivial, as we need to derive in equilibrium the corresponding bidding function.

Note that the threshold v_1 depends only on the signal v_0 and the conditional distribution $F(v|v_0)$, so it is independent of the bidding function used (i.e., of the auction format and bidding strategies of other bidders). What is common to all auctions and equilibria is that the bid which will be chosen for v_0 is identical to the one of the corresponding threshold v_1 . Figure 3 illustrates this for the case of a second-price auction and a hypothetical first-price auction. Hence, we can identify pairs (v_0, v_1) such that bids are the same across these pairs but vary across them (for example, increase with v_0 if F is stochastically increasing in v_0). This is analogue to indifference curves, where only the *bidding labeling* would vary across different auctions. What is common for the points (v_0, v_1) is that v_1 is precisely the expected value starting at v_0 and conditional on not having a final value above the threshold v_1 in a rebidding state. Hence v_0 and v_1 are *expected-value equivalent*.

The partition of the space of values into expected-value equivalent pairs (v_0, v_1) can be used to reduce the dynamic auction into an equivalent static one. Starting with an initial distribution $\Phi(v_0)$ and a conditional distribution $F(v|v_0)$, we can construct a new distribution of final values as follows. For any initial v_0 , assign a value $v_1(v_0)$ to the histories where either the bidder does not have a rebidding opportunity in the second period or gets a value $v \leq v_1$. In the complement (histories where the bidder can rebid and $v > v_1$), set the final value v . Assigning the corresponding probabilities for these histories as determined from Φ , F and the rebidding probability p defines a distribution for final values \tilde{F} for each bidder and thus a static auction. Letting \tilde{B} denote the equilibrium bidding function for our bidder in that auction (e.g., $\tilde{B}(v) = v$ in the second-price auction), we can now assign $b_0(v_0) = \tilde{B}(v_1(v_0))$ and $B(v) = \tilde{B}(v)$ as equilibrium strategies in the dynamic auction.

Risk Aversion

Suppose the agent is risk averse over outcomes v . Then, repeating our analysis, it should follow that v_1 is the certainty equivalent of the lottery with threshold v_1 . This implies that the agent is indifferent between this lottery and getting v_1 for certain and would thus bid the same. Risk aversion should thus result in lower thresholds (i.e., more shading).

3 The Model

There are $i=1,\dots,N$ potential bidders in an auction. The auction is sealed-bid and takes place in the time interval $[0, T]$ when bids are submitted. As shown in Section 8, for the case of a second-price auction, the equilibrium bidding functions obtained here are also equilibrium bids for more general information structures and, in particular, for that of eBay auctions.

Each bidder has the option of submitting bids only at random times τ_1, τ_2, \dots . Bids can only be increased at any of these random bidding times and cannot be retracted. Bidders are risk neutral and the valuation of the bidder is modeled as a stochastic process $v_i(t)$ interpreted as the expected value at the end of the auction given the available information. A more formal description of the joint process of signals and values assumed is given in Appendix 11. Since these values are only relevant at bidding times τ_n , we restrict attention to the corresponding values v_n at these dates.¹³ Assume $\{v_{in}, \tau_{in}\}$ follow a joint Markov process with a transition function that is independent across bidders,¹⁴ and inscribed in a common probability space (Ω, σ) with typical element ω . Note that this specification allows for random entry and rebidding, as well as a random number of bidders. It also allows for the degree of inattention to depend on value and accommodates the case where a bidder is sure to come at the end of the auction, the latter by letting $P_i(\tau' \leq T | v, \tau) = 1$ for all v, τ . From now on, for notational convenience, we drop the i index whenever it does not lead to confusion.

Assuming that values and bidding times are Markov, this implies that we can identify information sets with pairs (v, τ) corresponding to the time and value in the last bidding window. A bidding function for bidder i specifies at each possible bidding time τ_n and, given value v_n , a desired bid $B_i(v_n, \tau_n)$. Given that bids can only be increased, $b_i(t) = \max\{B_i(v_m, \tau_m) | \tau_m \leq t\}$ is the bid that prevails at time t and, in particular, $b_i(T)$ is the final bid. Let $b_{-i}(T)$ denote the maximum bid over the remaining bidders at time T and $F_{-i}(b)$ its distribution. Denote the expected utility of bidder i for final value $v(T)$ and bid $b(T)$ by

$$U_i(v(T), b(T), F_{-i}).$$

Canonical examples are the second-price auction, where

$$U_i(v(T), b(T), F_{-i}) = \int^{b(T)} (v(T) - u) dF_{-i}(u);$$

¹³The process for the value $v_i(t)$ can be considered a continuous Markov process sampled at random stopping times τ_n that are Markov with respect to the last stopping time and valuation at that time.

¹⁴As shown in Section 10, results extend to the case of common values.

the first-price auction, where

$$U_i(v(T), b(T), F_{-i}) = F_{-i}(b(T))(v(T) - b(T));$$

and the all-pay auction, where

$$U_i(v(T), b(T), F_{-i}) = F_{-i}(b(T))v(T) - b(T).$$

As seen below, our results generalize to any auction format where bidders' best responses are given by increasing bidding functions.

An optimal bidding function B_i for bidder i is the one that maximizes expected utility at the time the bidder enters the auction, that is,

$$\max_{B(\cdot)} E(U(v(T), b(T), F_{-i}) | \tau_1, v_1), \tag{6}$$

where τ_1 is the time at which the bidder enters the auction, and v_1 is the initial value.

Definition. An *equilibrium* for the auction is a vector of bidding functions B_i and final distributions F_{-i} for each player i such that for every bidder i , the bidding function B_i is the best response to F_{-i} and bidding functions are consistent with the final distribution of bids.

4 Equilibrium Bidding

In this section, we characterize the equilibrium bidding functions and provide a dynamic programming problem that can be used to solve them. The main idea follows the example in section 2. We first provide an intuitive representation of the problem that is common to all auctions and then proceed to the formal analysis. The key insight is that we identify a distribution of valuations for each player that is independent of the auction format and opponents' valuations and strategies. Then we show that equilibria of the dynamic auction correspond one-to-one to equilibria of the static auction with respect to this distribution of values.

4.1 Intuitive Analysis

The examples in section 2 identified initial values v_0 with a threshold v_1 with the property that for *any* auction, the initial bid at v_0 equals the final bid at this threshold. Moreover, this

threshold, v_1 , was defined by the property $E(v|\text{no bidding opportunity with } v > v_1) = v_1$. In the general model where $t \in [0, T]$, a similar representation can be obtained. As illustrated by the example, bidders have an incentive to shade their bids prior to the end of the auction. More generally, we can partition the set of value and time pairs (v, t) into indifference classes such that bids are identical for all pairs in an indifference class, as depicted in Figure 1 in the introduction.¹⁵

As in the example given in Section 2, the pairs (v_0, t_0) along any indifference curve have a common feature. Starting at a point (v_0, t_0) , consider all paths with the property that at no future bidding time t_n the value lies above this indifference curve. Then $e(v_0, t_0)$ is the expected final value v_T conditional on this set of paths and is the same for all points in the indifference curve. The first four panels in Figure 4 illustrate paths that stay in this set, while the latter two panels paths that exit the set. In the first two cases, the bidder has no further chance of bidding. In the next two examples, the bidder has one more opportunity but at the time of arrival his valuation lies under this indifference curve and therefore does not change the outstanding bid. The last two panels of Figure 4 represent cases where the bidder gets a chance of bidding at a time period in which his valuation is above the indifference curve. In these two cases, the bidder will increase the outstanding bid, so these cases are not included in the calculation of conditional expected value.

As in the example in Section 2, we can also use the indifference curves together with the stochastic process for values and bidding times to derive a distribution of values for an equivalent static auction. In the example, this distribution has a point mass probability $(1 - p) + pF(\varepsilon)$ at the threshold ε and cumulative probability $(1 - p) + pF(v)$ for v above this point. For the general case, consider the following procedure for generating values. For any path $\omega = (v_1, t_1), (v_2, t_2), \dots, (v_n, t_n)$ define $v(\omega) = \max\{e(v_1, t_1), e(v_2, t_2), \dots, e(v_n, t_n)\}$, i.e., the value associated to the highest indifference curve reached during the bidding times t_1, \dots, t_n . Taken over all paths, this defines a distribution of values, as in the simple example. Take any auction format and this distribution of values, and assume there is an equilibrium with bidding functions $\tilde{B}_i(v)$ for each bidder. We can assign bids for all points in the dynamic auction by setting $B_i(v, t) = \tilde{B}_i(e(v, t))$.

Note that the indifference curves depicted in Figure 4 are downward sloping. This requires the expected values to be increasing in v and t . While this is not needed for the existence of equilibria, sufficient conditions for this to happen are provided below. Also note that for the case where the auction termination T follows a stationary random process (e.g., Poisson),

¹⁵Under some regularity conditions discussed below, this indifference curve will be downward sloping, as represented.

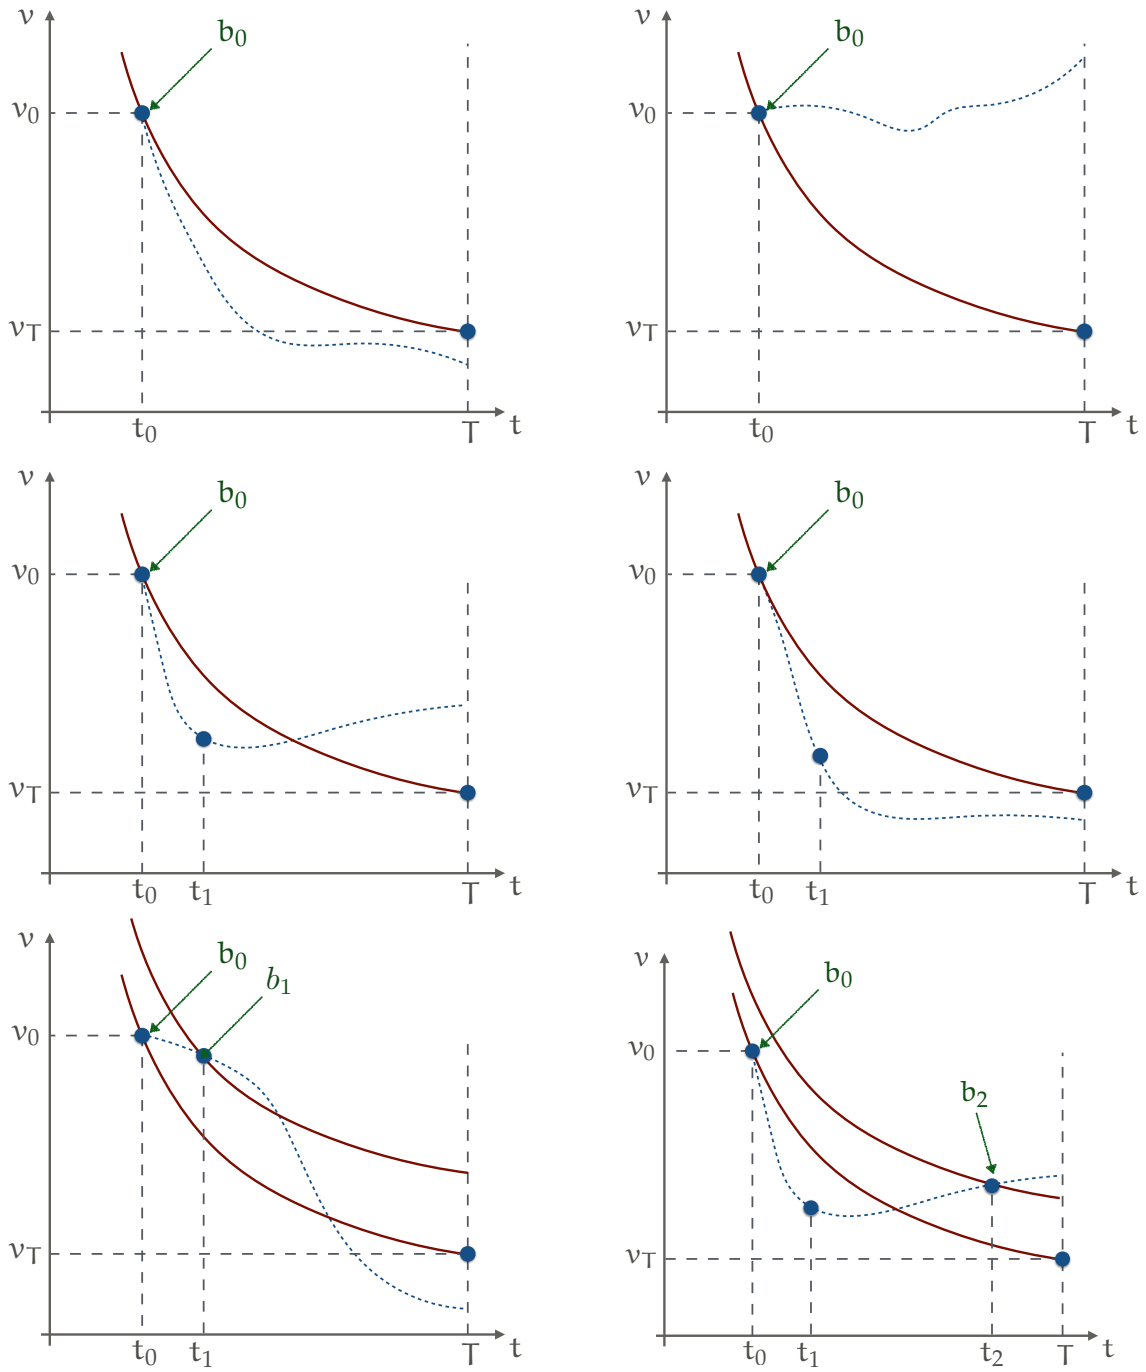


Figure 4: Indifference curves and increasing outstanding bids

Note: The dotted line represents changes in valuation of the bidder. In the top two graphs, the bidder does not get a chance of rebidding. In the middle two graphs, the bidder gets a second chance to update his bid but does not increase the bid. In the bottom two figures, the bidder increases his bid.

calendar time does not matter per se, so the indifference curves are flat.

5 Proof of the Main Results

Our analysis in the previous section suggests a general approach to proving existence and characterizing the equilibrium. The steps to this proof are as follows:

1. Define indifference curves that have the property that expected value conditional on not crossing above is constant. These we call self-generated expectations. We show they are uniquely defined and can be obtained via solving a dynamic programming problem.
2. Assign the corresponding bid to each point in that curve. This is the candidate bidding function.
3. Prove this is an optimal bidding function.

5.1 Self-Generated Expectations

Consider a function $e(v, t)$. Let $\bar{e}(t, \omega) = \max \{e(v_n(\omega), \tau_n(\omega)) \mid \tau_n \geq t\}$ and set this equal to zero if the set is empty. For every state (v, t) such that $0 \leq t \leq T$, we can define $H(v, t)$ as the set of histories ω following (v, t) such that $e(v, t) \geq \bar{e}(t, \omega)$.

Definition. The function $e(v, t)$ is a *self-generated expectation (SGE)* if it satisfies the following property for all (v, t) :

$$e(v, t) = E_{H(v,t)}v(T). \tag{7}$$

Given a self-generated expectation function $e(v, t)$, we can define *iso-expectation* level curves $I(u) = \{(v, t) \mid e(v, t) = u\}$. Intuitively, the level u indicates the conditional expectation of the final value of all paths, starting from a given point in $I(u)$, that never cross above this iso-expectation curve. These are the indifference curves described in the previous section.

5.2 Existence and Uniqueness of a Self-Generated Expectation

The derivation of a self-generated expectation follows a recursive structure. First note that $e(v, T) = v$, since this is a terminal node. Intuitively, working backwards from that point

using (7) should give a unique self-generated expectation $e(v, t)$. Consider the following dynamic programming problem:

$$W(\varepsilon, v, \tau) = \int_{\tau}^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau'|v, \tau) + P(\tau' > T|v, \tau) (E[v_T|v, \tau, \tau' > T] - \varepsilon), \quad (8)$$

where τ' denotes the following bidding time. We show below that $W(e(v, \tau), v, \tau) = 0$ if and only if $e(v, \tau)$ is a self-generated expectation.

Assumption 1. *Assume $P(\tau' > T|v, \tau) > \delta > 0$ for all (v, τ) . In addition, assume all conditional probabilities and expectations are continuous in the state.*

By Assumption 1 and using standard dynamic programming arguments, it follows that there is a unique function satisfying functional equation (8) and that it is strictly decreasing in ε and continuous. Moreover, it is greater or equal to zero when $\varepsilon = 0$ and negative for large ε . It follows by the intermediate value theorem that there is a unique value $e(v, \tau)$ such that $W(e(v, \tau), v, \tau) = 0$.

Theorem 1. *Given Assumption 1, the function $e(v, \tau)$ defined implicitly by $W(e(v, \tau), v, \tau) = 0$ is the unique self-generated expectation.*

Proof. See Appendix. □

5.3 Reduction to a Static Auction

For every history ω , we can identify a unique value corresponding to the highest iso-expectation level approached, $v(\omega) = \max\{e(v_n(\omega), \tau_n(\omega))\}$ for the corresponding path. This procedure determines uniquely a distribution of (final) values for each bidder that only depends on the corresponding Markov process P_i for that bidder. These distributions of values together with the auction format (e.g., first-price, second-price) define an equivalent final static auction. Letting $\{\tilde{B}_i(v) | i \in N\}$ denote equilibrium bidding functions in this final auction, we can define candidate bids in each state (v, t) for bidder i as $B_i(v, t) = \tilde{B}_i(e(v, t))$. The next step is to prove that the bidding functions B_i so defined are equilibrium functions in the dynamic auction.

5.4 Existence Proof

We make the following assumption.

Assumption 2. $U(v, b)$ is linear in v for fixed b and supermodular in (v, b) .

Supermodularity is a standard condition in mechanism design and holds for standard auctions. The linearity assumption follows for all standard auctions from risk neutrality.¹⁶

Proposition 1. *The optimal bid $B_i(v, t) = \tilde{B}_i(e(v, t))$.*

We provide an intuitive argument here, while the formal proof is in the Appendix. Take the given candidate optimal bidding function B_i for bidder i in all future instances and consider bid b in state (v, t) . Let $H(b)$ denote all paths $\omega = \{\tau_m, v_m\}_{\tau_m > t}$ starting from the current history, where $B(\tau_m, v_m) \leq b$ including $\tau_1 > T$ (i.e., no rebidding opportunity). Using the candidate bidding function for future states, this corresponds to states where $\tilde{B}(e(\tau_m, v_m)) \leq b$. Let $b(T, \omega)$ denote the final bid when applying this strategy in the future, so $b(T, \omega) = b$ for $\omega \in H(b)$ and is strictly greater than b in the complement. We take $U(v, b)$ to denote the expected utility of a bidder with value v and bid b in the final auction.¹⁷ The expected value of bidding b in the dynamic auction in state (v, t) equals

$$V(v, t, b) = \int_{H(b)} U(v(T), b) dP(\omega) + \int_{H(b)^c} U(v(T), b(T, \omega)) dP(\omega).$$

First, note that the boundary of the set $H(b)$ (and $H(b)^c$) consists of all those paths starting from (v_n, τ_n) for which the final bid $b(T)$ is equal to b . So, given the envelope condition, when considering the derivative of the above, we can ignore the effect of the change in the supports of the two integrals. The first-order condition is then $\partial V / \partial b = \frac{\partial}{\partial b} \int_{H(b)} U(v(T), b) dP(\omega) = 0$. By Assumption 2 (linearity), this is equivalent to the condition

$$\frac{\partial}{\partial b} U(E_{H(b)}v(T), b) P(H(b)) = 0, \tag{9}$$

so the optimal bid satisfies $b = \tilde{B}(E_{H(b)}v(T))$. It only remains to show that $E_{H(b)}v(T) = e(v, t)$. When $b = \tilde{B}(e(v, t))$, this follows easily; the set $H(b)$, is exactly the set such that for all future bidding states $e(v_m, t_m) \leq e(v, t)$. On the other hand, suppose some other b satisfies equation (9). Then for all future values where $B(v', t') = b$, it must be the case that $e(v', t') = E_{H(b)}v(T)$. It then follows from the definition of a self-generating expectation that $e(v, t) = E_{H(b)}v(T)$. Since by Theorem 1, there is a unique self-generating expectation, the optimal bid is uniquely determined by the condition given in Proposition 1.

¹⁶An interesting question is if our results extend to the case of risk aversion by defining the above indifference curves in terms of certainty equivalent of final values, as suggested in the analysis of our simple example.

¹⁷We suppress the distribution of other players' bids as an argument to simplify notation.

Corollary 1. *For the case of a second-price auction, the optimal bid $B(v, t) = e(v, t)$.*

This last result is intuitive, as in a second-price auction, a bidder submits a bid that equals the expected final value conditional on that bid being the final one, i.e., $b = E(v(T) | v, b(T) = b)$. Notice that this is exactly the condition that determined the initial bid b_0 for the example in Section 2.

6 Properties

In this section, we derive the general properties on equilibrium bidding behavior and consider some special cases.

6.1 Bid Shading

Our leading example suggests that bidders will shade bids as a consequence of adverse selection given by the option of future rebidding. In this section, we prove this is true for an arbitrary process.¹⁸ We first establish the following result:

Lemma 1. *$W(\varepsilon, v, \tau) \leq E[v_T | v, \tau] - \varepsilon$ with strict inequality if and only if rebidding occurs with positive probability.*

Proof. See Appendix. □

Since $W(\varepsilon, v, \tau)$ is decreasing in ε , the next proposition follows.

Proposition 2. *$e(v, \tau) \leq E[v_T | v, \tau]$ with strict inequality if and only if rebidding occurs with positive probability.*

As a consequence of this Proposition, for any auction where bids are an increasing function of value, $B(v, \tau) \leq B(E[v_T | v, \tau])$ with strict inequality if and only if rebidding occurs with positive probability.

¹⁸Bid shading can occur for strategic reasons; for example, in [Hortaçsu et al. \[2015\]](#), it is a result of dealer market power in uniform-price Treasury bill auctions. Our source of shading is distinct and fundamentally non-strategic.

6.2 Unbiased Expected Final Bids for Second-Price Auctions

Bids are shaded as a result of conditioning on the adverse selected set of no-rebidding. However, in the complement of this set, the current bid is replaced by a higher one. What is the overall effect on the final expected bid? The answer is unambiguous for the case of second-price auctions.

Proposition 3. *Consider a bidding time τ with value v and bid $b = e(v, \tau)$. Then the expected final bid equals v , that is, $E(B(T) | v, \tau, b = e(v, \tau)) = v$.*

Proof. Let H_0 consist of all histories ω starting from (v, τ, b) , where $B(T) = b$. By Proposition 1, $E(v(T) | H_0) = e(v, \tau) = b$. Now consider all histories ω on the complement H_0^c where rebidding occurs with probability one. Let τ' be the stopping time corresponding to the first time rebidding occurs. Assume, by way of induction, that $E[B(T) | v(\tau'), \tau', b = e(v(\tau'), \tau')] = v(\tau')$. By the optional stopping theorem, it follows that $E(B(T) | H_0^c) = E(v(\tau') | H_0^c) = E(v(T) | H_0^c)$. \square

In particular, this Proposition implies that when a bidder places the initial bid, the expected final bid is unbiased and equals the expected final value. This result is consistent with bid shading, since bids, while shaded, are still unbiased in the set where they *apply*. Section 6.5 derives a method for non-parametric identification of shading based on this result.

6.3 Bid Shading Over Time

We now apply our model to the case where rebidding times are independent of values. Assume arrival times are given by a distribution $F(\tau' | \tau)$. The following Proposition gives conditions such that the bidding function $B(v, \tau)$ is increasing in t .

Proposition 4. *Assume that values and time arrivals are independent and that $F(\tau' | \tau)$ is stochastically increasing in τ . Then $W(\varepsilon, v, \tau)$ and $e(v, \tau)$ are increasing in τ .*

Proof. See Appendix. \square

It follows immediately that in the case of a second-price auction, under the conditions of the Proposition, bid shading decreases with time. This implies that bids increase over time even in the absence of competitive pressure. The condition for the proposition follows immediately when the distribution of arrivals is independent of the initial time τ , so $F(\tau' | \tau) = G(\tau' - \tau)$.

While the assumptions of this proposition cover many relevant cases, some other useful ones are excluded. In particular, this proposition assumes that the new arrival τ' and the new value v' are independent. This will not be the case when information is accumulated over time prior to the next bidding opportunity, as in the case considered in Section 7 and in our application. The following proposition provides an alternative set of sufficient conditions that apply to those cases.

Proposition 5. *Suppose that*

1. $P(v', \tau' | v, \tau) = P_v(v' | v, \tau' - \tau) P_\tau(\tau' - \tau)$ and
2. $E(v' | v, \tau' - \tau)$ is nondecreasing in $(\tau' - \tau)$.

Then, $W(\varepsilon, v, \tau)$ and $e(v, \tau)$ are increasing in τ .

Proof. See Appendix. □

The assumptions in this proposition require that the time to the next bidding opportunity be independent of current calendar time τ , and that the next value v' only depend on v and the time elapsed until this next arrival.

6.4 Random Termination

We have assumed that the auction lasts for a fixed time $[0, T]$. Our formulation allows for random termination without modification. Consider equation (8), which is the key equation used to find the self-generated expectation. We can consider T in this equation as a random termination without any changes, with a slightly different interpretation: The term $P(v', \tau' | v, \tau)$ can be interpreted as the probability of the event that the next bidding time is t' and that $t' < T$ (i.e., the auction has not ended by then). Similarly, the second term, $P(\tau' > T | v, \tau)$, can be interpreted as the probability that the random termination occurs before the next arrival. With this change of interpretation, the same equation applies and so are all the results that follow.

It is useful to examine the conditions of Proposition 4 in light of this reinterpretation. Rewriting the assumption as $F(\tau' - T, \tau)$ stochastically *increasing* in τ (which, in the case of deterministic T , is an identical condition), the result follows. This is now an assumption regarding the difference between the two random variables, τ' and T . The following Corollary gives conditions that generalize Proposition 4 for the case of random T .

Corollary 2. Let $H(y|\tau) = P(T - \tau \leq y|\tau)$ denote the CDF for the remaining time of the auction conditional on $T \geq \tau$. Let $G(x|\tau) = P(\tau' - \tau \leq x|\tau)$ denote the CDF of the time to the next arrival. Assume

1. $H(y, \tau)$ is weakly increasing in τ and
2. $G(x|\tau)$ is weakly decreasing.

Then, $F(\tau' - T|\tau)$ is (weakly) decreasing in τ and $e(v, \tau)$ is increasing in τ .

Proof. Note that $\tau' - T = \tau' - \tau - (T - \tau)$. Let $(T - \tau) = x$ so $P(T - \tau' \leq z|\tau, x) = P(\tau' - \tau \leq z + x) = G(z + x|\tau)$. Integrating over x , it follows that $F(z|\tau) = \int G(z + x|\tau) dH(x|\tau)$. By the second assumption, the integrand is pointwise decreasing in τ . By the first assumption, the distribution H is stochastically decreasing in τ and, since G is an increasing function, it also implies that the integral is decreasing in τ . This proves that $F(z|\tau)$ is decreasing in τ . The second conclusion follows directly from Proposition 4. \square

The conditions of the Proposition have an intuitive interpretation. The second one is the analogue of the assumption made in Proposition 4. The first assumption simply states that the hazard rate for termination of the auction increases with duration, which seems a natural assumption in the case of random termination. These assumptions imply that the the level curves for self-generated expectations are decreasing, as depicted in Figure 1. In the special case where both conditional distributions G and H are independent of τ , $W(\varepsilon, v, \tau)$ and $\varepsilon(v, \tau)$ will also be independent of τ , so the level curves will be flat. Shading will still occur, but will not change over time.

6.5 Identifying Shading

This section shows that shading can be identified non-parametrically for the case of a second-price dynamic auction such as eBay. As usual, we assume the data consists of a large set of symmetric auctions drawn from the same underlying stochastic process P . The key idea for identification follows from Proposition 3 stating that while the bid $b = e(v, t)$ at time t is shaded, the expected final bid for this bidder equals v . If all bids were recorded, then taking all auctions where a bidder places a bid b at time t , the average final bid for those bidders should equal the unshaded value v .

In case the auction is not a sealed bid, as is the case for eBay auctions, this procedure is complicated by the censoring of bids, leading to a downward bias in the estimate of v .¹⁹ We now show a method that can be used to control for censoring that gives unbiased estimates. The information required is a record of the highest and second highest final bids as well as the time and value of all bids that when placed were highest.

Let $f(x|b, t)$ denote the probability density for the final bid of a bidder (for example, bidder one) when placing a winning bid b at time t , and $F(x|b, t)$ the CDF. Let $G(x|b, t)$ denote the CDF for the corresponding second highest final bid among the remaining bidders. To simplify notation, we suppress the conditioning on (b, t) that should be understood when constructing the moments.

Consider the following observations:

1. Frequency that bidder one's final bid is *observed* to be equal to x , either as first or second highest bid. The analogue according to the distributions is $f(x)G(x)$.
2. Frequency that bidder one and the second highest bid are below x . The analogue according to the distributions is $F(x)G(x)$.

Dividing these two moments, we eliminate $G(x)$ and get the inverse hazard rate $h(x) = f(x)/F(x)$. Integrating out this hazard rate gives

$$\ln F(x) = \int_0^x h(z) dz + \text{constant}.$$

The intuition here is precisely that when the final bid of bidder one is censored, we still observe the second highest bid, so we know that this censored bid is below the second highest bid, so we can still identify the CDF for this censored bid. Note that there is a point mass $f(b) = F(b) > 0$ at b due to non-retraction, so

$$F(x) = F(b) + \exp\left(\int_b^x h(z) dz\right).$$

$F(b)$ can be identified using

$$\lim_{x \rightarrow \infty} F(x) = F(b) + \exp\left(\int_b^{\infty} h(z) dz\right) = 1.$$

¹⁹The problem arises because in most cases bidders can only submit their bid if it is higher than the current second highest bid.

The expected final bid for bidder one is

$$\begin{aligned}
& \int_b x dF(x) + bF(b) \\
= & b + \int_b (x - b) dF(x) \\
= & b + \int_b (1 - F(x)) dx \\
= & b + \int_b (1 - F(b) - \exp(\int_b^x h(z) dz)) dx \\
= & b + \int_b (\exp(\int_b h(z) dz) - \exp(\int_b^x h(z) dz)) dx
\end{aligned}$$

So, bid shading equals

$$\int_b \left(\exp\left(\int_b h(z) dz\right) - \exp\left(\int_b^x h(z) dz\right) \right) dx.$$

In order to properly estimate shading from data, it would be necessary to account for heterogeneity across auctions as done in [Krasnokutskaya \[2011\]](#). This is an extension and an application worth pursuing, but beyond the scope of this paper.

6.6 Independent Increments

In this section, we consider two cases where bidding behavior is simplified: 1) increments independent of the current value v and 2) increments proportional to v . In particular, these apply to the cases where v follows an arithmetic and geometric Brownian motion, respectively. In both cases, we assume that bidding arrival times are given by a homogeneous Poisson process. In addition, we provide a result connecting shading to the variance of innovations.

Proposition 6. *Assume $P(v' = v + \delta | v, \tau)$ is independent of v for all δ and all τ , and arrival rates for bidding are independent of v . Then $W(\varepsilon + \delta, v + \delta, \tau) = W(\varepsilon, v, \tau)$.*

Proof. Follows by standard induction argument on the Bellman equation. Assume the function W has this property. Then evaluate

$$\begin{aligned}
W(\varepsilon + \delta, v + \delta, \tau) &= \int_t^T \min(W(\varepsilon + \delta, v' + \delta, \tau'), 0) dP(v' + \delta, \tau' | v + \delta, \tau) \\
&\quad + P(\tau' > T | v, \tau) (E[v_T | v + \delta, \tau, \tau' > T] - (\varepsilon + \delta)) \\
&= \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, \tau) \\
&\quad + P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] + \delta - (\varepsilon + \delta)) \\
&= W(\varepsilon, v, \tau).
\end{aligned}$$

□

The following Corollary simplifies shading to a simple one-dimensional shading function $s(t) = v - e(v, \tau)$ that is independent of v . Since this difference also corresponds to shading of bids $v - B(v, \tau)$ in a second-price auction, it implies that such shading is independent of v .

Corollary 3. $e(v + \delta, \tau) = e(v, \tau) + \delta$.

Consider now the case where $P(\gamma v', \tau' | \gamma v, \tau) = P(v', \tau' | v, \tau)$. As the next Proposition shows, this implies that $W(\gamma \varepsilon, \gamma v, \tau) = \gamma W(\varepsilon, v, \tau)$.

Proposition 7. *Assume $P(\gamma v', \tau' | \gamma v, \tau) = P(v', \tau' | v, \tau)$. Then $W(\gamma \varepsilon, \gamma v, \tau) = \gamma W(\varepsilon, v, \tau)$ and consequently $e(\gamma v, \tau) = \gamma e(v, \tau)$.*

Proof. Follows by standard induction argument on the Bellman equation. Assume the function W has this property. Then evaluate

$$\begin{aligned} W(\gamma \varepsilon, \gamma v, \tau) &= \int_t^T \min(W(\gamma \varepsilon, \gamma v', \tau'), 0) dP(\gamma v', \tau' | \gamma v, \tau) \\ &\quad + P(\tau' > T | v, \tau) (E[v_T | \gamma v, \tau, \tau' > T] - \gamma \varepsilon) \\ &= \int_t^T \min(\gamma W(\varepsilon, v', \tau'), 0) dP(v, \tau' | v, \tau) \\ &\quad + P(\tau' > T | v, \tau) \gamma (E[v_T | v, \tau, \tau' > T] - \varepsilon) \\ &= \gamma W(\varepsilon, v, \tau). \end{aligned}$$

The second property follows immediately from the definition of a self-generated expectation. □

Shading over Time with Independent Increments

A natural question is how the variance of new values affects shading, as it affects the option value of future rebidding. In the extreme, if variance were zero so $v(T) = v$ with probability one, there is no shading. We prove a monotonicity result for the case of independent increments considered in Proposition 6

Proposition 8. *Under the assumptions of Proposition 6, W is concave in v . A mean preserving increase in spread of the distribution of increments decreases $e(v, t)$.*

Proof. See Appendix. □

Note that while the assumptions require that the next arrival τ' is independent of current value v , they do not require that next value v' be independent from either τ or τ' .

In a Bayesian learning environment, the weight of new information decreases over time and so does the variance of the change in the posterior, which gives another reason for decreasing level of shading over time.

7 Brownian Motion and Poisson Arrivals

In this section, we introduce a special case, which we later estimate using eBay data and run counterfactual studies. Suppose that $v(t)$ follows a Brownian motion with drift μ and volatility σ , and the process for rebidding is Poisson with arrival ρ , as in the previous case. The functional equation (8) can be rewritten as

$$W(\varepsilon, v, \tau) = \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W\left(\varepsilon, v + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) d\Phi(z) d\tau' \quad (10) \\ + \exp(-\rho(T - t))(v + \mu(T - t) - \varepsilon).$$

Note that for this case, $E[v(T) | v, \tau] = v + \mu(T - t)$, so if the drift $\mu \leq 0$, then $E[v(T) | v, \tau]$ is increasing with t ; so, by Proposition 4, $e(v, \tau)$ is increasing in t .

Our specification satisfies the condition in 6; therefore, the proposition applies. Specifically, $W(\varepsilon, v, \tau) = W(0, v - \varepsilon, t)$. In consequence, we can write the value function $W_\mu(x, t)$ where $x = v - \varepsilon$,

$$W_\mu(x, t) = \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W_\mu\left(x + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) d\Phi(z) d\tau' \quad (11) \\ + \exp(-\rho(T - t))(x + \mu(T - t)).$$

It is worth mentioning scaling as another property that is used in Section 9.1 as well. If we scale σ and μ of the Brownian motion by a factor λ , the corresponding W function also scales proportionally and so are the self-generated expectation (and bids, in case of a second-price auction).

Lemma 2. $W(\lambda x, t; \lambda \mu, \lambda \sigma) = \lambda W(x, t, \mu, \sigma)$ and further $e(\lambda v, \tau; \lambda \mu, \lambda \sigma) = \lambda e(v, \tau; \mu, \sigma)$.

Proof. The proof is by induction. Assume the right hand side of (11) satisfies this property. It verifies immediately that $W(\lambda x, t; \lambda \mu, \lambda \sigma) = \lambda W(x, t, \mu, \sigma)$. This also implies that the shading function $x(t; \lambda \mu, \lambda \sigma) = \lambda x(t; \mu, \sigma)$ or, equivalently, $e(\lambda v, \tau; \lambda \mu, \lambda \sigma) = \lambda e(v, \tau; \mu, \sigma)$. \square

The effect of drift on the value function can also be easily determined, as shown by the following Lemma.

Lemma 3. *Let W_μ denote the value function, with drift μ and W the value function with drift 0. Then $W_\mu(x, t) = W(x + \mu(T - t), t)$.*

Proof. This is proved inductively using

$$\begin{aligned}
W(x + \mu(T - t), t) &= \\
&\rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W\left(x + \mu(T - \tau') + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) d\Phi(z) d\tau' \\
&\quad + \exp(-\rho(T - t))(x + \mu(T - t)) \\
&= \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W_\mu\left(x + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) d\Phi(z) d\tau' \\
&\quad + \exp(-\rho(T - t))(x + \mu(T - t)) = W_\mu(x, t).
\end{aligned}$$

It also follows that the bid shading function is of the form $x(t) + \mu(T - t)$, where $x(t)$ is the shading function when $\mu = 0$. In what follows, we assume $\mu = 0$. \square

7.1 Partial Differential Equation

In Section 9.1, we need to solve the functional equation 12 to find the bidding function. We do that by numerically solving its corresponding PDE. To find the PDE for the Bellman equation (the Hamilton-Jacobi equation), first subtract $W(x, t)$:

$$\begin{aligned}
0 &= \rho \int_t^{t+\Delta} \exp(-\rho(\tau' - t)) \int \left[\min\left(0, W\left(x + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) - W(x, t) \right] d\Phi(z) d\tau' \quad (12) \\
&\quad + \exp(-\rho\Delta) \left[\int W\left(x + \sqrt{\Delta}\sigma z, t + \Delta\right) - W(x, t) \right] d\Phi(z).
\end{aligned}$$

Taking derivative with respect to Δ and evaluating at $\Delta = 0$ result in

$$\begin{aligned}
0 &= \rho [\min(0, -W(x, t))] + \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial x^2} W(x, t) + \frac{\partial}{\partial t} W(x, t) \\
\rho \max(W(x, t), 0) &= \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial x^2} W(x, t) + \frac{\partial}{\partial t} W(x, t).
\end{aligned}$$

This PDE belongs to a standard class and can be easily solved with numerical methods. Once this function is derived, optimal shading can be obtained by finding the root $W(s, t) = 0$.

8 Bidding with Partially Observed Competing Bids

The analysis in this section applies to second-price auctions and in particular to the application considered in Section 9. Since for this case $B(v, t) = e(v, t)$, we will refer to bidding functions instead of self-generated expectations.

In our previous analysis, we have assumed that the auction is sealed bid, so there is no information on competing bids. However, if bidders can observe other bidders' previous bids, they may use that information strategically. In this section, we show that the previous derived equilibrium will remain valid given any information structure in case of second-price auctions. However, we cannot rule out that the existence of other equilibria. Throughout this section, we maintain our assumption of independent values across bidders. We first provide details of the information structure and define an equilibrium. We next show that the equilibrium derived before remains an equilibrium of the extended game.

At any point in time, each agent's information set includes a common public history $h(t)$ in addition to the private history $h_i(t)$ of all past realizations $\{\tau_{in}, v_{in}\}$ and bids for $\tau_{in} \leq t$. A strategy s_i for player i specifies at every bidding time τ_{in} and corresponding information set $(h(\tau_{in}), h_i(\tau_{in}))$ a bid b_{in} with the restriction that $b_{in} \geq b_{i,n-1}$ for all $n \geq 2$. Let $v_i(\omega, T)$ denote bidder i 's value for the object at the end of the auction, $b_i(T, s, \omega)$ denote the final bid of player i , and $b_{-i}(T, s, \omega)$ the highest bid among the remaining bidders, where s is the vector of strategies $s = (s_1, \dots, s_N)$. For simplicity of notation, we drop ω from the bidding function. Finally, let $u(v_i, b_i, b_{-i}) = \chi_{\{b_i \geq b_{-i}\}}(v_i - b_{-i})$ denote the final payoff for player i . We can define the final expected payoff for player i under strategy s , after history $h(\tau_{in}), h_i(\tau_{in})$, by

$$E_{\tau_{in}}(u(v_i(T), b_i(T, s), b_{-i}(T, s)) | h(\tau_{in}), h_i(\tau_{in})), \quad (13)$$

where beliefs are derived from information and other players' strategies. An equilibrium is a vector of strategies s such that for all players and at all information sets, strategy s_i maximizes (13) given s_{-i} .

We now show that the equilibrium derived previously remains an equilibrium in this game by showing that if all bidders use strategies that are functions of their private history only, then it is a best response to do so.

Proposition 9. *Let $B_i(v_n, \tau_n)$ denote the equilibrium bidding functions derived in Section 4. Then the strategies $s_i(v_n, \tau_n; h_{in}, h_n) = \max_{m \leq n} B_i(v_m, \tau_m)$ are an equilibrium of the game.*

Proof. Suppose for all remaining players s_{-i} is defined as above. We now show the following strategy is a best response for player i . It is sufficient to check that the bidder will not deviate from the bidding strategy at any arbitrary information set, $(v_n, \tau_n; h)$. Let player i follow this prescribed strategy in all future information sets. Without loss of generality, assume the optimal bid chosen at the information set considered is greater than the outstanding bid of the player. Let $F_{-i}(u|h)$ denote the distribution of the highest bid of b_{-i} (h_{-i} is enough). The optimal bid at (v_n, τ_n) solves

$$V(v_n, \tau_n, h) = \max_b \int_{H(b)} \int^b (v(T) - u) dF_{-i}(u|h) dP(\omega) + \int_{H(b)^c} \int^{b(\omega, T)} (v(T) - u) dF_{-i}(u|h) dP(\omega),$$

where, as before, $H(b)$ denotes the sets of histories where, given future bidding behavior, bid b is the final bid. So, the public history h is relevant only when considering its impact on the expected distribution for the highest bid of the remaining players in that set. Given the optimality of future bidding and the definition of $H(b)$, the effect on the integration set can be ignored, so it immediately follows that the optimal bid is independent of the distribution F_{-i} and consequently independent of h . \square

8.1 Censoring and Observed Behavior

Let $B_i(v, \tau)$ denote the maximum bid that bidder i is willing to place at time t . When taking the model to the data, one must take into account that when a bidder observes an outstanding bid greater than $B(v, \tau)$, the bidder will choose not to bid. Hence, the process for observed bids is censored, and this censoring depends on the available information. In eBay auctions, bidders can see the outstanding second highest value, thus defining the threshold below which bids are censored. The model has implications for the observed timing of bids. If at time t a bidder has a high probability of returning to the auction right before its end, that bidder's desired bid $B(v, \tau)$ will be very low and thus it is likely to be censored by an existing higher second bid. As a consequence, it is unlikely to see any bids from this bidder until the end of the auction, consistent with the observed sniping behavior.

9 Data and Estimation

In the following section, we identify the model using the eBay auction mechanism as a motivating example. The auction mechanism on eBay is similar to a second-price or Vickery auction. A seller sets the starting bid of an auction, and bidders can bid repeatedly for the item until the end of the auction. Each bidder observes all previous bids, except for the current highest bid. A bidder should bid an amount higher than the current second highest bid, plus some minimum increment.²⁰ If this value is higher than the current highest bid, the bidder becomes the new highest bidder. Otherwise, the bidder becomes the second highest bidder. The winner has to pay the second highest bid, plus the increment or his own bid, whichever is smaller. Auctions last for one to ten days, and they have a predetermined and fixed ending time that cannot be changed once the auction is active. As a general rule, bidders cannot retract or cancel a bid.²¹

By using eBay data on successful auction sales during the first week of June 2014, we observe the following bidding behavior that suggests the importance of dynamics for these auctions. First, a considerable number of bidders submit multiple bids. For example, over 30% of bidders place more than one bid. Second, the increase in the bid amount placed by the same bidder is considerable. As an example, in the case of tablets, the median increase in bids for bidders that place a bid in the first 50% of the auction time and the last 10% of time is 33%. Third, as noted in the literature, the timing of winning bids is skewed toward the end of the auction. More than 40% of all bids, and more than 60% of winning bids, are submitted in the last 10% of the duration of an auction.

We consider three subcategories of the eBay marketplace: Tablets and iPads, Tickets, and Artworks.²² We chose these categories because they are more consistent with the private value assumption in our setup.

9.1 Estimation

In this section, we use the simulated method of moments to estimate the Brownian motion setting described in section 7 using eBay data. In our setting, bidders are symmetric and have independent draws for their arrival times, values, and bidding times. Arrival times

²⁰The increment is a function of the second highest bid, 1–5% of the second highest bid, fixed for all auctions, and is set by eBay.

²¹<http://goo.gl/ebaHpR>, accessed July 2016.

²²We consider only artworks with a final price lower than \$500 to eliminate incentives from resale and therefore common value.

are given by a Poisson process with arrival rate η . This implies that both the total number of bidders and their arrival times are random. We normalize the total time of the auction $T = 1$ so that both the expected number of bidders participating in an auction and the variance are equal to η . It is interesting to note that conditional on the number of bidders participating in an auction, arrival times are uniform. Succeeding bidding opportunities are given by a Poisson process with arrival rate ρ .

The process for bidders' values is modeled as follows. For a bidder arriving at time $t \in [0, 1]$, the initial value is drawn from a Normal distribution with mean μ and standard deviation $\sigma_0 + \sigma t$. After this arrival, it follows a Brownian motion with zero drift and volatility σ . This process implies that the unconditional distribution of values for a bidder bidding at time τ is also Normal with mean zero and standard deviation $\sigma_0 + \sigma\tau$ *independently of that bidder's initial arrival time to the auction*. Our assumption captures the idea that the amount of information a bidder has in an auction only depends on calendar time and not on the time of entry. In particular, this would be true if all relevant information to determine valuation were publicly available to all bidders.

In this simple estimation method, we identify four parameters, $(\eta, \sigma_0, \sigma, \rho)$. To identify these parameters, we use the simulation method of moments. We first use grid search on the parameter space to find an initial guess for the parameters. Next, we run the Laplace-type estimator introduced in [Chernozhukov and Hong \[2003\]](#) to generate standard errors. We choose moments that we believe are informative of the above parameters:

1. Average number of bidders, defined as those that placed either the highest or second highest bid at some point during the auction. We refer to these as *recorded* bids.
2. Average number of bids per bidder.
3. Variance of bids recorded within the first 50% of the auction time normalized by final price.
4. For bidders who bid both in the first 50% of the auction time and in the last 10% of it, the increment between the two bids normalized by the final price.

While we have a total of five parameters, $(\eta, \mu, \sigma_0, \sigma, \rho)$, three of them (μ, σ_0, σ) can be identified with our moments only up to a scalar. (There is no problem identifying the other two.) This is a consequence of [Lemma 2](#) and the normalization used to calculate our two last moments. Intuitively, scaling μ, σ_0 , and σ by a factor λ is equivalent to an equal change in scale of bidders' values. This in turn implies that for any realization of values and bidding

Table 1: Moments

	#of Bidders	# Bids per Bidder	Var. of Early Bids	Average Increment
Tablets	6.28	1.38	0.053	36.7%
Tickets	2.97	1.53	0.063	29.7%
Artworks	2.03	1.27	0.077	44.1%

times in an auction, all bids will be scaled by λ and so will be the final price. Given that in computing the last two moments', bids are normalized by this final price, the corresponding ratios are unaffected by the scale change. Hence our moments allow us to identify only the ratios σ_0/μ and σ/μ . This is not a problem, as these ratios are also sufficient for our counterfactual exercises.

In each step we should estimate a loss function defined as the sum of squared of percentage deviations from the above moments. To find the shading function, we need to first numerically solve the PDE described in Section 7 for a set of parameters (σ, ρ) . Next, we can simulate the model and calculate the optimal bidding function of bidders at any arrival time, taking into account censoring.

9.2 Results

The estimated parameters are the ones that minimize the loss function. Table 1 gives the moments for the three categories. Given that gradual bidding cannot be generated by our model, we remove all such bids by considering only the last bid a bidder placed within two minutes of the previous bid. This reduces the average number of bids per bidder. Table 2 shows the estimated parameters, $(\eta, \rho, \frac{\sigma_0}{\mu}, \frac{\sigma}{\mu})$, and their standard errors.

It is instructive to consider the order of magnitude of some of these parameters and how they relate to our moments. As discussed in Section 9.1, the expected number of bidders in an auction is equal to the Poisson arrival rate, η . However, some of these bidders would never place a bid, as their intended bid at any arrival instance is lower than the second highest bid (refer to Section 8.1). The expected number of observed bidders as a function of total bidders arrived can be estimated analytically, assuming that bidders' intended bid is independent of the time of arrival.²³ A bidder who arrives after k bidders already have arrived will bid if his bid is higher than the second highest bid at the moment. If $k = 0, 1$, this probability is 1; otherwise, it will be $2/(k + 1)$. This is because his bid should be either the highest bid or

²³This is just an approximation given that shading increases a bidder's bid toward the end of the auction and also given that one should also consider the distribution of number of bidders rather than just the average.

the second highest bid. Adding these together, the expected number of observed bidders if n bidders show up is $b(n) = 2 + \sum_{m=3}^n \frac{2}{m}$. Calculating this using the closest integer to the estimated η for the three groups results in the expected number of observed bids equal to 5.5, 3.2, and 2 in Tablets, Tickets, and Artworks, respectively.

Bidders entering the auction can place a bid immediately. Future arrival times follow a Poisson distribution with parameter ρ . We can approximate the average number of bids per bidder using the following argument. On average, bidders arrive at midpoint. The number of arrivals starting at t has a Poisson distribution with parameter $\rho(T - t)$, which is also the mean number of arrivals. Taking the average entry time for a bidder $t = 1/2$, this implies a mean number of arrivals (in addition to the first one) equal to $\rho/2$. Since for each arrival there is a probability $1/2$ that the value increases, ignoring shading, the expected number of bids is $1 + \rho/4$. This is the case without censoring, as bids have to be recorded to be in our observations. An estimated $\rho = 1.03$ (Artworks) gives $1 + 1.03/4 = 1.257$ and an estimated $\rho = 1.72$ gives $1 + 1.72/4 = 1.43$. These are in the order of magnitude of the moments above.

The two normalized standard deviations, $\left(\frac{\sigma_0}{\mu}, \frac{\sigma}{\mu}\right)$, are jointly identified by using the variance of normalized early bids and average increments of the bids between initial 50 percentile and last 10 percentile for the same bidder. Both of these normalizations are with respect to the final price of the auction. Variance of early bids is affected by both parameters, as we assumed that if bidders arrive at t , their valuation follows a Normal distribution with standard deviation $\sigma_0 + t\sigma$. Given the uniform distribution for the arrival of bidders, the average of t is $1/4$. So, we can say $(\sigma_0/\mu + \sigma/4\mu) \sim \frac{p}{\mu} \sqrt{\text{var of early bids}}$, where p is the final price. The average increment on the bids is identified by σ . The Artworks category, on average, has the highest standard errors and the Tablet category the least.

9.3 Information and the Timing of Bids

In our setup, conditional on the number of bidders, arrival times are uniform. It follows that in the absence of information arrival (i.e., $\sigma = 0$) where bidders' policy is to bid their valuation upon first arrival, the time distribution of winning bids is also uniform. Information arrival ($\sigma > 0$) skews the distribution of winning bids toward the end of the auction for two reasons. First, given the change of information, bidders shade their bids, and this shading decreases monotonically to zero as bidders get closer to the end of the auction, thus giving an upward trend to bids. Second, the arrival of new information increases the variance of bidders' valuations as time progresses. This leads to an increase in the first-order statistic of the distribution of values, making it more likely that the winning bid be placed later. These

Table 2: Estimated Parameters

Variable	Mean	Median	Std. Dev.	95% Confidence Interval	
Tablets					
η	13.87	13.85	0.49	12.91	14.94
ρ	1.72	1.72	0.10	1.54	1.91
$\frac{\sigma_0}{\mu}$	0.15	0.15	0.01	0.13	0.17
$\frac{\sigma}{\mu}$	0.43	0.43	0.01	0.40	0.45
Tickets					
η	3.78	3.78	0.14	3.48	4.04
ρ	2.12	2.13	0.16	1.77	2.44
$\frac{\sigma_0}{\mu}$	0.30	0.30	0.02	0.25	0.34
$\frac{\sigma}{\mu}$	0.46	0.46	0.03	0.40	0.51
Artworks					
η	2.08	2.07	0.13	1.86	2.35
ρ	1.03	1.02	0.16	0.72	1.33
$\frac{\sigma_0}{\mu}$	0.11	0.09	0.08	0.01	0.32
$\frac{\sigma}{\mu}$	1.11	1.12	0.07	0.97	1.23

two effects can be separated by considering a counterfactual where buyers do not shade their bids but still have the effect of the increase in information as the auction proceeds.

We define early winning bids as those that are placed in the first 50% of the auction time and late bids as those placed in the last 10% of the auction time. As noted, in the standard environment (considered in the literature) where there is no information arrival ($\sigma = 0$) and bidders' arrival to the auction is stationary, the timing of winning bids is uniform, implying a 50% share of early winning bids and 10% of late ones. In contrast, as seen in Table 3, our baseline estimation for Tablets gives shares of early winning bids close to 9% and late bids close to 30%. A standard measure of inequality is the Gini coefficient, which would be zero with uniform distribution of winning time, and is approximately 0.55 for our benchmark. In the absence of bid shading, the asymmetry is lower but still considerable: 15.2% for early bids, 26.5% for late bids, and a Gini coefficient estimate of 0.47. While the actual distribution of bids is significantly more skewed (5% early winning bids, 85% late winning bids, and Gini coefficient of 0.87), our benchmark model goes a long way toward closing the gap (63% of the data, as measured by the Gini coefficient). This is the case in spite of the simplicity of our structure and the lack of heterogeneity across bidders. In particular, we assume ρ is the same across bidders and time, while in reality, it is a likely source of heterogeneity.²⁴

²⁴As an example, suppose there are n bidders. One of them can bid with probability one at the end of the auction, while the other $n - 1$ bidders only bid at the beginning of the auction. To make things extreme, suppose the final value is uniform between $[0, 1]$ and all the $n - 1$ initial bidders have no information, and

Table 3: Timing of Winning Bids

	$\sigma = 0$	Baseline	No Shading	Data
Tablets				
Early Bids	50%	9.4%	15.2%	5%
Late Bids	10%	30.0%	26.5%	85%
Gini Coefficient	0	0.55	0.47	0.87
Tickets				
Early Bids	50%	19.9%	26.4%	9%
Late Bids	10%	24.1%	20.8%	77%
Gini Coefficient	0	0.41	0.32	0.77
Artworks				
Early Bids	50%	28.6%	33.0%	16%
Late Bids	10%	19.0%	16.9%	63%
Gini Coefficient	0	0.29	0.22	0.62

Note: Early bids are the winning bids placed in the first 50% of the auction time, and late bids are the winning bids placed in the last 10% of the auction time.

We have also assumed that the arrival time of bidders to the auction is uniform, while it is more likely that bidders pay more attention to auctions ending soon. In addition, the use of sniping programs that submit bids toward the end of the auction creates a wedge between the time a bid is placed, or revised, and the time it is submitted to eBay. In the other two categories, our model explains approximately half of skewness observed in data based on Gini coefficient.

9.4 Bid Retraction

In this section, we consider bid retraction. This exercise has a series of motivations. First, in the absence of restrictions on bid retraction, bidders have no incentive to shade bids, and this allows us to quantify the effect of shading. Second, the question might be of interest as a matter of auction design. Third, in practice, some bidders are able to work around the restrictions on bid retraction by using a sniping program, which submits bids automatically in the last few seconds of an auction. Bidders using these programs can change their bids at any point in time and thus can retract their intended bids. We compare the results with both our baseline and an alternative benchmark given by an ex post efficient allocation, as would be given by a second-price auction that takes place at the end of the auction after bidders know their final valuation.

thus approximately the same expectation, equal to $1/2$; so, the initial winning bid is also around $1/2$. Thus, the late bidder wins the auction $1/2$ of the time, even though representing $1/n$ of the number of bidders.

For the bid retraction case, we assume that bidders can choose to retract or lower their bids during any one of their bidding times but not outside them. This is a natural restriction, because if bidders were allowed to retract at the end, bidding would become meaningless. The possibility of retraction eliminates the source of adverse selection mentioned earlier and thus the incentives for bid shading. Given that values follow a martingale, at every bidding time τ , a bidder will choose a bid equal to the current value v_τ .

The effect of bid retraction on the distribution of final bids for a given bidder can be further characterized.

Proposition 10. *The final bids under bid retraction are a mean preserving spread of those without retraction.*

Proof. Consider an information set (v_n, τ_n) where bid b_n is placed. From Proposition 3, it follows that $E b(T|v_n, \tau_n) = v_n$, and the same holds true under bid retraction. Let $H_0(b)$ denote the set of histories starting at this node where bid b is the final bid. On $H_0(b)^c$, the expected final bid must be the same with and without retraction. This is because the first time there is rebidding in that set, final bids are unbiased and equal to the corresponding value, for both cases. It must then follow that the expected final bid in $H_0(b)$ under bid retraction is the same as that under no retraction, which by definition is b , implying that bid retraction gives a mean preserving spread of final bids in this set. Applying this argument forward on $H_0(b)^c$ proves that final bids under bid retraction are a mean preserving spread of those without retraction. \square

The mean preserving spread under bid retraction has two implications. First, it represents the fact that bids are more closely connected to values. Second, it is likely to lead to an increase in the second-order statistic and therefore price, particularly in the presence of many bidders.

Table 4 gives the average performance of this bidding scenario compared to the baseline case. The expected value in the bid retraction final allocation is 0.76% higher than that in the baseline in Tablet category. In turn, average prices are 1.27% higher under bid retraction, a consequence of the mean preserving spread discussed above. The increase in prices exceeds the increase in value, so bid retraction decreases the expected payoffs for bidders by approximately 1.67% while increasing the sellers' revenue by 1.27%. However, in case of Tickets, both bidders and sellers are better off as the effect of increase in efficiency overcomes the downward impact of higher prices for bidders. The impact for Artworks follows the trends in Tablets.

Table 4: Value and Revenues Compared to Baseline, Tablets

	Bid Retraction Tablets	Bid Retraction Tickets	Bid Retraction Artworks
Expected Value	0.76%	1.10%	1.40%
Expected Price	1.27%	1.26%	3.17%
Expected Payoff	-1.67%	0.68%	-1.00%

10 Correlated Information

In many cases, it is likely that information or signals observed are correlated across bidders. For example, the arrival of a competing auction with a similar product is an event that creates an opportunity cost and is likely to affect in a correlated way the value of all bidders keeping track of that information. Our results extend easily to the case of pure common values, where all agents have the same values but different bidding windows.

10.1 Pure Common Values

We restrict our analysis again to the case of dynamic second-price auctions. The setting is identical to that described in our general model, but where all bidders' values v_i are identical, bidders observe the same signals but have independent bidding windows. As before, it follows that the optimal bid at state (v, τ) satisfies $b_i = E(v(T) | v, \tau, b(T) = b_i)$, where now $b(T) = \max_{ij} b_i(v_j, t_j)$ is the maximum over all bids, including those of other bidders. The recursive representation given before still holds,

$$W(\varepsilon, v, \tau) = \int_{\tau}^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, \tau) + P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] - \varepsilon),$$

where, now, the interpretation of arrivals is for the arrival of any bidder. As an example, if opportunities for bidding are Poisson with arrival ρ , then the total arrival rate that would be used in this dynamic programming equation is $N\rho$, where N is the number of potential bidders. The effect on bidding behavior is similar to an increase in ρ that intuitively should result in greater shading of bids. This follows naturally the interpretation of increased adverse selection, as now the current bid will win not only when that bidder decides not to bid any higher but also when any other bidder chooses not to do so.

11 Final Remarks

This paper considered a general theory for bidding dynamics in long auctions. The theory relies on a simple yet rich structure with potentially broad applicability. The problem of finding equilibria in the dynamic auction is reduced to solving for the corresponding equilibria in an equivalent static auction with respect to a distribution of values that is independent of the auction format, making this general class very tractable.

Values and bids in this static auction are shaded because of the option of rebidding and, under mild conditions, the incentives for shading decrease as the auction progresses. This contributes to the concentration of winning bids toward the end of the auction. In addition, information arrival throughout the auction results in increased dispersion of values over time, which also contributes to account for late winning bids. Our application to eBay shows that in a very simple version of the theory, these two factors are capable of generating considerable skewness in the distribution of bids, as observed in the data.

While bids are shaded relative to expected values, the *expected final bid* still equals the expected final value and is thus unbiased. These two results are consistent, since bids, while shaded, are still unbiased in the set where they *apply*. This property is exploited in the paper to derive a promising method for non-parametrically estimating shading.

The theory could be extended in several directions. It seems that risk aversion can be introduced relatively easily by defining the self-generated expectations in terms of certainty equivalent values. Our methods might also extend to the case of correlated values, with the obvious complication that self-generated expectations would require to be solved jointly across all bidders. In particular, the methods extend to the extreme case of perfect correlation, as shown in the paper. Our methods should also generalize to more general stochastic processes for signals and bidding opportunities. Finally, our techniques could also be applied to other dynamic games, such as patent races, given their similarity to all pay auctions.

The analysis of dynamic games has proven to be a difficult problem. There are obvious tradeoffs in research and corners to cut. Our paper is no exception and we have our share of strong assumptions. In particular, we have chosen to represent the impact of information on values, alternatives and opportunities, and the existence of bidding frictions in a reduced form, given by the stochastic process for values and bidding opportunities. There are obvious shortcomings, but the payoff is that of a parsimonious representation of equilibria and a very tractable general structure that could be easily used in further applications.

References

- Attila Ambrus, Yuhta Ishii, and James Burns. Gradual bidding in ebay-like auctions. *Economic Research Initiatives at Duke*, July 2014. 1, 1, 11
- Susan Athey and Philip A Haile. Identification of standard auction models. *Econometrica*, 70(6):2107–2140, 2002. 1
- Christopher Avery. Strategic jump bidding in english auctions. *The Review of Economic Studies*, 65(2):185–210, 1998. 1
- Matthew Backus and Gregory Lewis. A demand system for a dynamic auction market with directed search. *Harvard University*, October, 2012. URL <http://goo.gl/zdxan3>. 1
- P. Bajari and A. Hortacsu. The winner’s curse, reserve prices, and endogenous entry: empirical insights from eBay auctions. *RAND Journal of Economics*, pages 329–355, 2003. 1
- Tilman Börgers and Christian Dustmann. Strange bids: bidding behaviour in the united kingdom’s third generation spectrum auction. *The Economic Journal*, 115(505):551–578, 2005. 1
- Victor Chernozhukov and Han Hong. An mcmc approach to classical estimation. *Journal of Econometrics*, 115(2):293–346, 2003. 9.1
- Dominic Coey, Bradley Larsen, and Brennan C Platt. A theory of bidding dynamics and deadlines in online retail. 2015. 1, 1
- Olivier Compte and Philippe Jehiel. The wait-and-see option in ascending price auctions. *Journal of the European Economic Association*, 2(2-3):494–503, 2004. 1
- Olivier Compte and Philippe Jehiel. Auctions and information acquisition: sealed bid or dynamic formats? *The Rand Journal of Economics*, 38(2):355–372, 2007. 1
- Peter Cramton. Spectrum auctions. *Handbook of Telecommunications Economics*, 2002. 2
- Kent D Daniel and David A Hirshleifer. A theory of costly sequential bidding. *University of Michigan Business School Working Paper*, (98028), 1998. 1
- Joachim R Groeger and Robert A Miller. Bidding frictions in ascending auctions. 2015. URL <https://goo.gl/3VKNGY>. 1, 1

- Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000. 1
- Philip A Haile and Elie Tamer. Inference with an incomplete model of english auctions. *Journal of Political Economy*, 111(1):1–51, 2003. 1
- Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *The Review of Economic Studies*, 49(3):315–333, 1982. 5
- Kenneth Hendricks and Alan Sorensen. The role of intermediaries in dynamic auction markets. Technical report, Working Paper, 2015. 1
- Johannes Hörner and Nicolas Sahuguet. Costly signaling in auctions. *The Review of Economic Studies*, 74(1):173–206, 2007. 1
- Ali Hortaçsu, Jakub Kastl, and Allen Zhang. Bid shading and bidder surplus in the u.s. treasury auctionsystem. 2015. URL <https://goo.gl/PC9crD>. 18
- Koichiro Ito and Mar Reguant. Sequential markets, market power and arbitrage. *American Economic Review*, 106(7):1921–1957, 2016. 1
- Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489, 2003. 1, 1
- Yuichiro Kamada and Michihiro Kandori. Asynchronous revision games. Technical report, mimeo, 2011. 1, 1
- Yuichiro Kamada and Michihiro Kandori. Revision games part i: Theory. *Working paper*, December 2015. 1, 1
- Elena Krasnokutskaya. Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies*, 78(1):293–327, 2011. doi: 10.1093/restud/rdq004. URL <http://restud.oxfordjournals.org/content/78/1/293.abstract>. 6.5
- Axel Ockenfels and Alvin E Roth. 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), 2002. 1
- Axel Ockenfels and Alvin E Roth. Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic behavior*, 55(2):297–320, 2006. 1, 9

Leonardo Rezende. Mid-auction information acquisition. *University of Illinois mimeo. www.econ.puc-rio.br/lrezende/noise.pdf*, 2005. 10

Maher Said. Sequential auctions with randomly arriving buyers. *Games and Economic Behavior*, 73(1):236–243, 2011. 1

Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476, 2006. 1

Appendix 1. Signals vs. Values

We modeled our process for $v(t)$ as values. An alternative is to consider signals, $s(t)$. This appendix discusses the implicit assumptions used on the process of signals and in particular the joint distribution of final value $v(T)$ and these signals.

Values

We modeled $v(t)$ as a stochastic process that is observed at random times $0 \leq \tau_n \leq T$. Further we assume that $(v(\tau_n), \tau_n)$ is a Markov process. Assuming that at all time τ , $E(v(T)|v(\tau), \tau)$ is strictly increasing in $v(\tau)$, we can substitute this value by the corresponding conditional expectation. Under this assumption, there is no loss of generality of assuming that the value $v(\tau)$ is the conditional expectation of $v(T)$ at time τ .

Note. This does not imply that the process $v(\tau_n)$ is a martingale, for our model allows for the possibility that values are correlated with arrivals. Thus while the unconditional expectation $E(v'|v, t) = v$ (integrating over all future arrival times), the conditional expectation for a specific arrival time $E(v'|v, t, t')$ might be different. For example, late arrivals could be correlated with lower values.

Signals

Suppose there is a process for signals s_n , each observed at bidding time τ_n . (We could in principle have a more general structure where the arrival of signals and bidding times are not the same.) Hence at the n^{th} arrival time, the vector of signals is $s = (s_1 \dots s_n)$ and arrival times $\tau = (\tau_1 \dots \tau_n)$. There is a joint probability function $P(s, \tau, v_T)$, where v_T is interpreted as the final value, so

$$E(v_T|s, \tau) = \frac{\int v_T P(s, \tau, dv_T)}{\int P(s, \tau, dv_T)}.$$

The question might be phrased as follows: Under what conditions is the process (v_n, τ_n) given by $v_n = E(v_T|s_1 \dots s_n, \tau_1 \dots \tau_n)$ a Markov process? So, given (v_n, τ_n) , the distribution of (v_{n+1}, τ_{n+1}) should be independent of the history of signals and times. Consider

$$\begin{aligned} v_{n+1} &= E(v_T|s_1 \dots s_{n+1}, \tau_1 \dots \tau_{n+1}) \\ &= E(v_T|v_n, s_1 \dots s_{n+1}, \tau_1 \dots \tau_{n+1}) \\ &= E(v_T|v_n, \tau_n, s_{n+1}, \tau_{n+1}). \end{aligned}$$

We want to be able to write the last equality so that v_n, τ_n, s_{n+1} and τ_{n+1} are sufficient for the expectation of v_T , and then also want s_{n+1} and τ_{n+1} , given v_n and τ_n , to be conditionally independent of $(s_1 \dots s_n, \tau_1, \dots, \tau_n)$. In other words, (v_n, τ_n) have to be sufficient statistics for s_{n+1}, τ_{n+1} and in turn $(v_n, \tau_n, s_{n+1}, \tau_{n+1})$ sufficient statistics for v_{n+1} .

A Bayesian learning environment can be easily accommodated in this setup. For example, take the case of a Kalman filter with exogenous Poisson arrivals for information and bidding opportunities.

Appendix 2. Proofs

Proof of Theorem 1

Let $e(v, \tau)$ be a threshold function with the property that

$$E[v(T) | v, \tau \text{ and } \bar{e}(T) = e(v, \tau)] = e(v, \tau). \quad (14)$$

We will prove that there is a unique self-generating function satisfying property (14) and that it solves $W(e(v, \tau), v, \tau) = 0$, where $W(e, v, \tau)$ is the unique solution to the Bellman equation:

$$\begin{aligned} W(e, v, \tau) &= \int_{\tau}^T \min(W(e, v', \tau'), 0) dP(v', \tau' | v, \tau) \\ &+ P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] - e). \end{aligned} \quad (15)$$

The first step is to establish *necessity*, that is, any bidding function $e(v, \tau)$ with property (14) corresponds to a function $W(e, v, \tau)$ satisfying this functional equation and $W(e(v, \tau), v, \tau) = 0$. The second step is to show that the Bellman equation is a contraction mapping and that the unique solution $W(e, v, \tau)$ is strictly decreasing in e . Consequently, there is a unique function $e(v, \tau)$ for which $W(e(v, \tau), v, \tau) = 0$.

Step 1. Necessity

Take a self-generating function $e(v, \tau)$, i.e., that satisfies property (14). For any state (v, τ) and starting threshold ε , consider the random variable defined by $e(T) = \max\{e(t_n, v_n), \varepsilon\}$

for all future bidding time $t_n > \tau$. Let $Q(\varepsilon, v, \tau)$ denote the resulting probability that $e(T) = \varepsilon$. Now define

$$W(\varepsilon, v, \tau) = \{E[v(T) | v, \tau, e(T) = \varepsilon] - \varepsilon\} Q(\varepsilon, v, \tau). \quad (16)$$

We now show that $W(\varepsilon, v, \tau)$ is a solution to functional equation (15). Substituting in the right hand side of (15) results in

$$\begin{aligned} W(\varepsilon, v, \tau) &= \int \min(\{E[v(T) | v', \tau', e(T) = \varepsilon] - \varepsilon\} Q(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, \tau) \\ &\quad + P(\tau' > T | v, \tau) (E[v(T) | v, \tau, \tau' > T] - \varepsilon). \end{aligned}$$

As indicated, ε will only remain valid in the next arrival state v', τ' when $e(v', \tau') \leq \varepsilon$ and, by Lemma 4, this will occur if and only if $E[v(T) | v', \tau', e(T) = \varepsilon] \leq \varepsilon$. It follows that

$$Q(\varepsilon, v, \tau) = \int_{\tau}^T \mathbb{I}_{E[v(T) | v', \tau', e(T) = \varepsilon] - \varepsilon \leq 0} Q(\varepsilon, v', \tau') dP(v', \tau' | v, \tau) + P(\tau' > T | v, \tau)$$

and, using the law of iterated expectations, it follows that

$$W(\varepsilon, v, \tau) = \{E[v(T) | v, \tau, e(T) = \varepsilon] - \varepsilon\} Q(\varepsilon, v, \tau).$$

To complete the first step, we need to show that $W(e(v, \tau), v, \tau) = 0$; this follows immediately from the definition of W and of a self-generated expectation $e(v, \tau)$ (14).

Step 2. Sufficiency and Uniqueness We first show that there is a unique function $W(\varepsilon, v, \tau)$ satisfying (15) by establishing it is a contraction mapping. Check Blackwell sufficient conditions: Monotonicity is trivially satisfied. To check discounting, consider the function $W(\varepsilon, v, \tau) + a$ for $a \geq 0$ on the right hand side of the Bellman equation (15).

$$\begin{aligned} &\int_{\tau}^T \min(W(\varepsilon, v', \tau') + a, 0) dP(v', \tau' | v, \tau) \\ &+ P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] - \varepsilon) \\ &\leq \int_{\tau}^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, \tau) + a(1 - P(\tau' > T | v, \tau)) \\ &+ P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] - \varepsilon) \\ &= W(\varepsilon, v, \tau) + a(1 - P(\tau' > T | v, \tau)). \end{aligned} \quad (17)$$

By Assumption 1 $P(\tau' > T | v, \tau) > \delta$ for some $0 < \delta < 1$, proving the second Blackwell sufficient condition.

To complete the proof, we need to show that there is a unique bidding function $e(v, \tau)$ satisfying $W(e(v, \tau), v, \tau) = 0$ for all (v, τ) . We show this recursively by establishing that W is strictly decreasing in the first argument, for any (v, τ) . The proof is then by induction. So, assume that the W function in the right hand side of the Bellman equation (15) is weakly decreasing. Letting $\varepsilon' > \varepsilon$,

$$\begin{aligned} TW(\varepsilon', v, \tau) &= \int_{\tau}^T \min(W(\varepsilon', v', \tau'), 0) dP(v', \tau' | v, \tau) \\ &\quad + P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] - \varepsilon') \\ &\leq \int_{\tau}^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, \tau) \\ &\quad + P(\tau' > T | v, \tau) (E[v_T | v, \tau, \tau' > T] - \varepsilon'), \end{aligned}$$

which is strictly less than $TW(b, v, \tau)$ since by Assumption 1 $P(\tau' > T | v, \tau) > 0$, thus proving sufficiency and uniqueness.

Supporting Lemmas

Lemma 4. *Consider a self-generated expectation $e(v, \tau)$. Starting with an arbitrary threshold ε at state (τ_0, v_0) , define $\varepsilon_{\tau'} = \max(\varepsilon, e(\tau', v(\tau')))$ for all $\tau' > \tau_0$ and define $e(T)$ as the final value. Then, $E[v(T) | \tau_0, v_0 \text{ and } e(T) = \varepsilon] - \varepsilon$ is decreasing in ε .*

Proof. Take the initial threshold ε at some point (τ_0, v_0) and $\delta > 0$. For any path ω following (τ_0, v_0) , that is, $\omega \in \Omega(\tau_0, v_0)$ and $\tau > \tau_0$, define $e(T, \omega, \varepsilon) = \max\{e(t_n(\omega), v_n(\omega)), \varepsilon\}$. Let $H(\varepsilon)$ denote the set of paths where $e(T, \omega, \varepsilon) = \varepsilon$ and $H(\varepsilon + \delta)$ the set where $e(T, \omega, \varepsilon + \delta) = \varepsilon + \delta$ when applying the same rule to the initial value bid $\varepsilon + \delta$. It follows immediately that $e(T, \omega, \varepsilon + \delta) = \max\{e(T, \omega, \varepsilon), \varepsilon + \delta\}$. Then

$$\begin{aligned} \int_{H(\varepsilon + \delta)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega) &= \int_{H(\varepsilon)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega) \\ &\quad + \int_{\{\varepsilon < e(T, \omega, \varepsilon) \leq \varepsilon + \delta\}} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega) \\ &\leq \int_{H(\varepsilon)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega) \\ &\quad + \int_{\varepsilon < e(T, \omega, \varepsilon) \leq \varepsilon + \delta} (v(\omega, T) - (e(T, \omega, \varepsilon))) dP(\omega) \\ &= \int_{H(\varepsilon)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega), \end{aligned}$$

where the last equality follows from Lemma 5. It follows that

$$\begin{aligned} & \int_{H(\varepsilon+\delta)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega) - \int_{H(\varepsilon)} (v(\omega, T) - \varepsilon) dP(\omega) \\ & \leq \int_{H(\varepsilon)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega) - \int_{H(\varepsilon)} (v(\omega, T) - \varepsilon) dP(\omega) < 0 \quad . \end{aligned}$$

Now

$$\begin{aligned} E[v(T) | \tau_0, v_0 \text{ and } e(T) = \varepsilon + \delta] - (\varepsilon + \delta) &= \frac{\int_{H(\varepsilon+\delta)} (v(\omega, T) - (\varepsilon + \delta)) dP(\omega)}{P(H(\varepsilon + \delta))} \\ &\leq \frac{\int_{H(\varepsilon)} (v(\omega, T) - \varepsilon) dP(\omega)}{P(H(\varepsilon + \delta))} \\ &\leq \frac{\int_{H(\varepsilon)} (v(\omega, T) - \varepsilon) dP(\omega)}{P(H(\varepsilon))} \\ &= E_{H(\varepsilon)} v(\omega, T) - \varepsilon. \end{aligned}$$

□

Lemma 5. $\int_{e(T, \omega) \in B} (v(\omega, T) - (e(T, \omega))) dP(\omega) = 0$ for any (Borel set) B

The proof of this Lemma uses the following property of conditional expectation in Ash (Theorem 5.3.3. pg. 210):

Theorem. Let Y be an extended random variable on (Ω, \mathcal{F}, P) , and $X : (\Omega, \mathcal{F}) \rightarrow (\Omega', \mathcal{F}')$ a random object. If $E(Y)$ exists, there is a function $g : (\Omega', \mathcal{F}') \rightarrow (\bar{\mathbb{R}}, \mathcal{B})$ such that for each $A \in \mathcal{F}'$,

$$\int_{\{X \in A\}} Y dP = \int_A g(x) dP_x(x), \quad (18)$$

where $P_x(A) = P(\omega | X(\omega) \in A)$. The function $g(x)$ is interpreted as $E(Y | X = x)$.

Proof of Lemma Define a series of random variables that are a function of ω :

$e(\omega)$: final expected value indifference curve reached in that path

$\tau'(\omega)$: time at which it was reached

$v(\omega)$: final value for that path ω

Let $X(\omega) = (\tau(\omega), e(\omega))$. The random variable that we will consider is $Y(\omega) = v(\omega) - e(\omega)$. Let $A = \{\tau_0 < \tau \leq T\} \times B$, where B is a Borel set (of expected values) in \mathbb{R} . We want to show that

$$\int_{X \in A} Y(\omega) dP(\omega) = 0.$$

We know that for all $\tau > \tau_0$ and ε , $E(Y|\tau, \varepsilon) = 0$, where the pair (τ, ε) are interpreted as the final expected value and time at which it was reached. Substituting $x = (\tau, \varepsilon)$ and using (18),

$$\int_{\{\tau(\omega) > \tau_0, \varepsilon(\omega) \in B\}} (v(\omega) - \varepsilon(\omega)) dP(\omega) = \int_{\{\tau > \tau_0, \varepsilon \in B\}} E(Y|\tau, \varepsilon) dP_x(\tau, \varepsilon) = 0.$$

Proof of Proposition 1

We need to show that the bidding function defined in Proposition 1 is a best response for any bidder (say bidder one) given that the other bidders follow the corresponding bidding function. First note that from the perspective of bidder one, the distribution of the final bids of the other bidders is exactly the same that would be obtained in the final auction. So all we need to show is that if \tilde{B}_1 is a best response to this final distribution of bids, so is B_1 in the dynamic auction. In what follows, we suppress the index of bidder one to simplify notation.

Consider a node (v, t) and some alternative bid $b_2 \neq b_1 \equiv B(v, t)$. We will show that bid b_2 cannot be an improvement over b_1 . Let $V(v, t, b)$ denote the expected utility bidding b at this state and following the (candidate) optimal bidding policy for the future. A convenient expression can be obtained as follows:

$$\begin{aligned} V(v, \tau, b) &= \int_{H(b)} U(v(T), b) dP(\omega) \\ &+ \int_{H(b)^c} U(v(T), b(\omega, T)) dP(\omega), \end{aligned}$$

where $b(\omega, T)$ is the final bid for path ω when following the optimal bidding function, and $H(b)$ denotes the set of paths starting from this node such that b is the final bid, that is, $H(b) = \{\omega \in \Omega(v, t) : b(\omega, T) = b\}$. By abuse of notation, we will let $H(b)^c = \{\omega \in \Omega(v, t) : b(\omega, T) > b\}$. Given the definition of b_1 and the linearity assumption, it follows that

$$U(E_{H(b_1)} v(T), b_1) \geq U(E_{H(b_1)} v(T), b_2).$$

Suppose first that $b_2 > b_1$. Then

$$\begin{aligned}
V(v, t, b_2) &= \int_{H(b_1)} U(v(T), b_2) dP(\omega) \\
&+ \int_{b(\omega, T) > b_2} U(v(T), b(\omega, T)) dP(\omega) \\
&+ \int_{b_1 < b(\omega, T) < b_2} U(v(T), b_2) dP(\omega). \tag{19}
\end{aligned}$$

Consider a relaxed problem where this bidder is allowed to follow the original bidding function in any future node. This is a relaxed constraint since it requires that $b(\omega, T) \geq b_1$ that is lower than b_2 . As a result, the continuation value in the set $H(b_1)^c$ is the same as under the original plan and thus

$$\begin{aligned}
V(v, \tau, b_2) - V(v, \tau, b_1) &\leq \int_{H(b_1)} [U(v(T), b_2) - U(v(T), b_1)] dP(\omega) \tag{20} \\
&= U(E_{H(b_1)}v(T), b_2) - U(E_{H(b_1)}v(T), b_1) \leq 0.
\end{aligned}$$

Now suppose instead that $b_2 < b_1$.

Consider the decision node (v, t) . Let $v_1 = e(v_1, t) = E_{H(b_1)}v(T)$ denote its associated conditional expected value, so $b_1 = \tilde{B}(v_1)$. For the alternative bid $b_2 < b_1$, it easily follows that

$$\begin{aligned}
V(v, t, b_2) &= \int_{H(b_2)} U(v(T), b_2) dP(\omega) + \int_{H(b_2)^c} U(v(T), b(\omega, T)) dP(\omega) \tag{21} \\
&= \int_{H(b_1)} U(v(T), b_2) dP(\omega) + \int_{H(b_1)^c} U(v(T), b(\omega, T)) dP(\omega) \\
&\quad + \int_{b_2 < b(\omega, T) \leq b_1} (U(v(T), b(\omega, T)) - U(v(T), b_2)) dP(\omega).
\end{aligned}$$

By Assumption 2 (linearity)

$$\begin{aligned}
\int_{H(b_1)} U(v(T), b_2) dP(\omega) &= E_{H(b_1)}U(v(T), b_2) P(H(b_1)) \tag{22} \\
&= U(E_{H(b_1)}v(T), b_2) P(H(b_1)) \\
&= U(v_1, b_2) P(H(b_1)).
\end{aligned}$$

By the definition of $V(v, t, b)$ and using (21) and (22), we can write

$$\begin{aligned}
V(v, t, b_1) - V(v, t, b_2) &= (U(v_1, b_1) - U(v_1, b_2)) P(H(b_1)) \\
&- \int_{b_2 < b(\omega, T) \leq b_1} (U(v(T), b(\omega, T)) - U(v(T), b_2)) dP(\omega).
\end{aligned} \tag{23}$$

Applying Lemma 6,

$$\begin{aligned}
V(v, t, b_1) - V(t, v, b_2) &= (U(v_1, b_1) - U(v_1, b_2)) P(H(b_1)) \\
&- \int_{v_2 < v(\omega) \leq v_1} U(v(\omega), B(v(\omega))) - U(v(\omega), b_2) dP(\omega).
\end{aligned} \tag{24}$$

Consider the last integral. Supermodularity of the U function implies that B must be an increasing function, so $b_1 \geq B(v(\omega)) \geq B(v_2) \geq b_2$. It also follows from supermodularity and $v_1 \geq v(\omega)$ that

$$U(v(\omega), B(v(\omega))) - U(v(\omega), b_2) \leq U(v_1, B(v(\omega))) - U(v_1, b_2). \tag{25}$$

Finally, since $b_1 = B(v_1)$ it follows that

$$U(v_1, B(v(\omega))) - U(v_1, b_2) \leq U(v_1, b_1) - U(v_1, b_2). \tag{26}$$

Combining (24), (25) and (26), it follows that

$$\begin{aligned}
V(v, t, b_1) - V(t, v, b_2) &\geq (U(v_1, b_1) - U(v_1, b_2)) P(H(b_1)) \\
&- \int_{v_2 < v(\omega) \leq v_1} U(v_1, b_1) - U(v_1, b_2) dP(\omega),
\end{aligned} \tag{27}$$

and since the set of paths $\{v_2 \leq v(\omega) \leq v_1\}$ is a subset of $H(b_1)$, then $V(v, t, b_1) - V(t, v, b_2) \geq 0$, so the proof is complete.

Definitions and Lemma 6

For every path ω and its associated decision nodes, we say that (t_n, v_n) is the *final decision node* if and only if $e(v_n, t_n)$ is maximal on these decision nodes and it is the earliest occurrence of the maxima. Intuitively, this will correspond to the point where the final bid $b(\omega, T)$ for that path is placed. It easily follows that each point (v, t) defines a partition of the set of paths Ω with the property that the conditional expected final value in that partition

is $e(v, t)$. Here we are interested in path ω that follows a particular state (\hat{v}, \hat{t}) , so the definition can be made relative to the corresponding sub σ -algebra and for $t_n > \hat{t}$. This takes into account only bidding times that occur after this decision node.

Consider the bidding function defined by $B(v, t) = \tilde{B}(e(v, t))$. Take a node (v, t) and its corresponding sub σ -algebra. For each path $\omega \in \Omega(v, t)$, we can define

$$b(\omega, T) = B(e(v_n, t_n)), \text{ where } (v_n, t_n) \text{ is the final decision node for } \omega.$$

We will focus here on all decision nodes that follow (v, t) and construct the partition of $\Omega(v, t)$ as described above. This defines for each ω in a partition its corresponding final decision node (v_n, t_n) and the conditional expectation $v(\omega) = e(v_n, t_n)$.

Lemma 6. *Consider any bids $b_1 > b_2$. Let v_1 and v_2 be the associated values such that $b_1 = B(v_1)$ and $b_2 = B(v_2)$. (If these don't exist, then take $v_2 = \inf \{v | B(v) \geq b_2\}$ and $v_1 = \sup \{v | B(v) \leq b_1\}$). Then the following holds:*

$$\begin{aligned} & \int_{b_2 < b(\omega, T) \leq b_1} U(v(T), b(\omega, T)) dP(\omega) \\ &= \int_{v_2 < v(\omega) \leq v_1} U(v(\omega), B(v(\omega))) dP(\omega). \end{aligned}$$

Proof. Follows from the above definition and application of iterated expectation (omitted). □

Proof of Lemma 1

First note that by the law of iterated expectation, $E(E([v_T | v', \tau'] - \varepsilon | v', \tau') | v, \tau) = E[v_T | v, \tau] - \varepsilon$ and that

$$\begin{aligned} E(E([v_T | v', \tau'] - \varepsilon | v', \tau') | v, \tau) &= \int_{\tau} E([v_T | v', \tau'] - \varepsilon | v', \tau') dP(v', \tau' | v, \tau) \\ &= \int_{\tau}^T E([v_T | v', \tau'] - \varepsilon | v', \tau') dP(v', \tau' | v, \tau) \\ &\quad + \int_T E([v_T | v', \tau'] - \varepsilon | v', \tau') dP(v', \tau' | v, \tau) \\ &= \int_{\tau}^T E([v_T | v', \tau'] - \varepsilon | v', \tau') dP(v', \tau' | v, \tau) \\ &\quad + P(\tau' > T | v, \tau) E([v_T | v', \tau'] - \varepsilon | v, \tau, \tau' > T). \end{aligned}$$

We now prove the Lemma by induction. So suppose that $W(\varepsilon, v', \tau') \leq E(v(T) | v', \tau') - \varepsilon$. Then

$$\begin{aligned}
W(\varepsilon, v, \tau) &= \int_{\tau}^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau'|v, \tau) + P(\tau' > T|v, \tau) (E[v_T|v, \tau, \tau' > T] - \varepsilon) \\
&\leq \int_{\tau}^T W(\varepsilon, v', \tau') dP(v', \tau'|v, \tau) + P(\tau' > T|v, \tau) (E[v_T|v, \tau, \tau' > T] - \varepsilon) \\
&\leq \int_{\tau}^T E([v_T|v', \tau'] - \varepsilon|v', \tau') dP(v', \tau'|v, \tau) + P(\tau' > T|v, \tau) (E[v_T|v, \tau, \tau' > T] - \varepsilon). \\
&= E[v_T|v, \tau] - \varepsilon,
\end{aligned}$$

where the first inequality is strict iff rebidding occurs with positive probability. .

Proof of Proposition 4

Consider the functional equation:

$$\begin{aligned}
W(\varepsilon, v, \tau) &= \int_{\tau}^T \left[\int \min(W(\varepsilon, v', \tau'), 0) dP(v'|v) \right] dF(\tau'|\tau) \\
&\quad + (1 - F(T|\tau)) [E(v_T|v, \tau) - \varepsilon].
\end{aligned}$$

By induction hypothesis, assume that the right hand side W is increasing in τ' . Given the assumption of stochastic dominance, $F(\tau'|\tau)$ is stochastically increasing in τ and so is the integral. The other two effects of increasing τ involve reducing the area of integration on the first term, a positive effect, and increasing the probability of the last term. Since by the martingale property the second term is $(1 - F(T|\tau))(v - \varepsilon)$ and from Lemma 1 and the martingale property, $\int \min(W(\varepsilon, v', \tau'), 0) dP(v'|v) \leq \int v' dP(v'|v) \leq v$ for all τ' , it follows that this shift of mass will increase the expected value. Finally, since W is decreasing in the first argument, $e(v, \tau)$ must increase with τ .

Proof of Proposition 5

Assume by way of induction that $W(\varepsilon, v', \tau')$ is increasing in its last argument. Using the first assumption in the proposition, equation (8) can be rewritten as

$$\begin{aligned}
W(\varepsilon, v, \tau) &= \int_0^{T-\tau} \int \min(W(\varepsilon, v', \tau + x), 0) dP_v(v'|v, x) dP_{\tau}(x) \\
&\quad + \int_{x>T-\tau} \int (v' - \varepsilon) dP_v(v'|v, x) dP_{\tau}(x).
\end{aligned}$$

Take $\tau_2 > \tau_1$.

$$\begin{aligned}
W(\varepsilon, v, \tau_2) &= \int_0^{T-\tau_2} \int \min(W(\varepsilon, v', \tau_2 + x), 0) dP_v(v'|v, x) dP_\tau(x) \\
&\quad + \int_{x>T-\tau_2} \int (v' - \varepsilon) dP_v(v'|v, x) dP_\tau(x) \\
&\geq \int_0^{T-\tau_2} \int \min(W(\varepsilon, v', \tau_1 + x), 0) dP_v(v'|v, x) dP_\tau(x) \\
&\quad + \int_{x>T-\tau_2} \int (v' - \varepsilon) dP_v(v'|v, x) dP_\tau(x) \\
&= \int_0^{T-\tau_1} \int \min(W(\varepsilon, v', \tau_1 + x), 0) dP_v(v'|v, x) dP_\tau(x) \\
&\quad + \int_{x>T-\tau_1} \int (v' - \varepsilon) dP_v(v'|v, x) dP_\tau(x) \\
&\quad + \int_{T-\tau_2}^{T-\tau_1} \int \min(W(\varepsilon, v', \tau_1 + x), 0) dP_v(v'|v, x) dP_\tau(x) \\
&\quad - \int_{T-\tau_2}^{T-\tau_1} \int (v' - \varepsilon) dP_v(v'|v, x) dP_\tau(x) \\
&\geq \int_0^{T-\tau_1} \int \min(W(\varepsilon, v', \tau_1 + x), 0) dP_v(v'|v, x) dP_\tau(x) \\
&\quad + \int_{x>T-\tau_1} \int (v' - \varepsilon) dP_v(v'|v, x) dP_\tau(x) \\
&\quad W(\varepsilon, v, \tau_1).
\end{aligned}$$

The first inequality follows from the induction hypothesis. The second inequality follows from $W(\varepsilon, v', \tau_1 + x) \leq v' - \varepsilon$, which is proved in Lemma 1.

Proof of Proposition 8

Consider the dynamic programming equation for W ,

$$\begin{aligned}
W(\varepsilon, v, \tau) &= \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau'|v, \tau) \\
&\quad + P(\tau' > T|\tau) (E[v_T|v, \tau, \tau' > T] - \varepsilon).
\end{aligned}$$

Given the assumption of scale invariance and that τ' is independent of v , it follows that $E[v_T|v, \tau, \tau' > T] = v + E[v_T|0, \tau, \tau' > T]$, which is linear in v . Hence we just need to prove that the first term is concave in v . Assume by way of induction that the right hand side function W is concave in v' . Then

$$\begin{aligned}
& \int_t^T \min (W (\varepsilon, v', \tau'), 0) dP (v', \tau' | (\alpha v_2 + (1 - \alpha) v_1), \tau) \\
= & \int_t^T \min (W (\varepsilon, \alpha (v' + v_2) + (1 - \alpha) (v' + v_1), \tau'), 0) dP (v', \tau' | 0, \tau) \\
\geq & \int_t^T (\alpha \min \{W (\varepsilon, (v' + v_2), \tau'), 0\} + (1 - \alpha) W (\varepsilon, (v' + v_1), \tau')) dP (v', \tau' | 0, \tau) \\
= & \alpha \int_t^T \min \{W (\varepsilon, v', \tau'), 0\} dP (v', \tau' | v_2, \tau) \\
& + (1 - \alpha) \int_t^T \min \{W (\varepsilon, v', \tau'), 0\} dP (v', \tau' | v_1, \tau).
\end{aligned}$$

The second result follows immediately from concavity and the definition of a mean preserving increase in spread.