

Robust Bidding and Revenue in Descending Price Auctions*

Sarah Auster[‡] Christian Kellner[§]

This Version: March 2019

Abstract

We study the properties of Dutch auctions in an independent private value setting, where bidders face uncertainty over the type distribution of their opponents and evaluate their payoffs by the worst-case from a set of probabilistic scenarios. In contrast to static auction formats, participants in the Dutch auction gradually learn about the valuations of other bidders. We show that the transmitted information can lead to changes in the worst-case distribution and thereby shift a bidder's payoff maximizing exit price over time. We characterise the equilibrium bidding function in this environment and show that the arriving information leads bidders to exit earlier at higher prices. As a result, the Dutch auction systematically generates more revenue than the first-price auction.

Keywords: Auctions, Ambiguity, Consistent Planning

1 Introduction

The descending price (or Dutch) auction and the first-price sealed-bid auction are known to be strategically equivalent in the canonical model, regardless of whether bidders are risk-neutral or risk-averse, whether values are affiliated or independent, etc. An important

*We would like to thank Piotr Dworczak, Peter Klibanoff, Sujoy Mukerji, Frank Riedel, Andrzej Skrzypacz and Marciano Siniscalchi for very helpful comments.

[‡]Department of Decision Sciences and IGIER, Bocconi University, email: sarah.auster@unibocconi.it

[§]Department of Economics, University of Southampton, email: c.kellner@soton.ac.uk

assumption underlying the canonical model is that bidders are subjective expected utility maximizers and, as such, quantify the uncertainty over other bidders' private information with a single prior belief (or behave as if they do). Yet, in practice, bidders often do not know with precision the distribution of their opponent's private valuation, especially when transactions are infrequent and past data points are scarce. This paper revisits the properties of the descending price auction and shows that the consideration of uncertainty about the distribution of bidders' valuations has important implications for equilibrium bidding and revenue properties.

We study a setting with two bidders, whose private valuations are drawn independently from a continuum of values, according to some unknown distribution. From the viewpoint of a bidder, the descending price auction takes the form of an optimal stopping game in continuous time. At each point in time, the bidder decides whether to terminate the auction and pay the currently displayed price or whether to wait for a lower price. We depart from the benchmark model by assuming that bidders demand robustness with regard to the uncertainty they perceive. Their preferences are represented by the maxmin expected utility model (Gilboa and Schmeidler, 1989). Accordingly, bidders evaluate each bid by the worst case from a set of distributions that they view as possible.

An important feature of the Dutch format is that it gradually transmits information to the participants: given a bidding strategy, players learn about their opponent's valuation from the observation that the auction has not yet been terminated. In particular, as time progresses, bidders revise their beliefs towards lower and lower types. We assume that bidders update each of the distributions and consider the worst case over the set of updated distributions. Updating is thus prior by prior, also referred to as full Bayesian updating. Apart from a smoothness condition, we do not impose any restrictions on the set of distributions over which bidders minimise and, therefore, do not rule out that the worst-case scenario for a bidder, identified by one of the distributions in the set, changes over time. When such changes occur, bidders face a time-inconsistency problem: their current optimal exit point may no longer be optimal at a later point in time.¹ We assume that bidders are sophisticated in that they correctly anticipate how their future selves will act and optimise accordingly.

We start the analysis with an important preliminary result: provided the other bidder follows a strictly monotone strategy, a bidder locally evaluates her expected payoff from

¹Also Riedel (2009) studies optimal stopping problems with multiple-prior preferences but, in contrast to us, restricts attention to sets of priors under which dynamic consistency is ensured.

continuing the auction at a given point in time with the distribution function that has the *highest reversed hazard rate*. The reversed hazard rate for a probability distribution with a cumulative distributions function F and density function f is described by the ratio f/F (see Figure 1). It determines the rate by which the probability of winning the auction falls at a given point in time conditional on reaching that time.

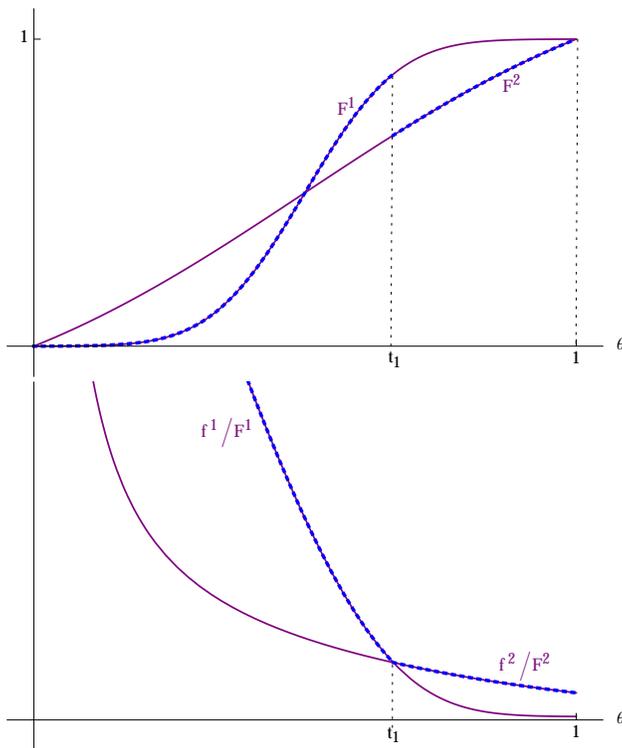


Figure 1: Two cumulative distributions functions and their reversed hazard rates. Both distributions are truncated normals with mean 0.5 and standard deviation 0.135 (F^1) and 0.5 (F^2).

When the maximal reversed hazard rate belongs to different distributions at different points of the support, bidders' preferences over exit prices change over time. The fact that time is continuous makes such preference reversals challenging. To address this issue, a new approach is followed: we construct an auxiliary distribution function with the property that its reversed hazard rate agrees with the maximal reversed hazard rate in the set of distributions over which bidders minimise, illustrated in Figure 2. Using the auxiliary probability distribution, we follow a standard first-order condition approach to derive the candidate equilibrium bidding function; as if both bidders had a single prior belief corresponding to the auxiliary distribution. Finally, to verify the equilibrium, we leverage the fact that bidders foresee the behavior of their future selves and use a backward induction argument.

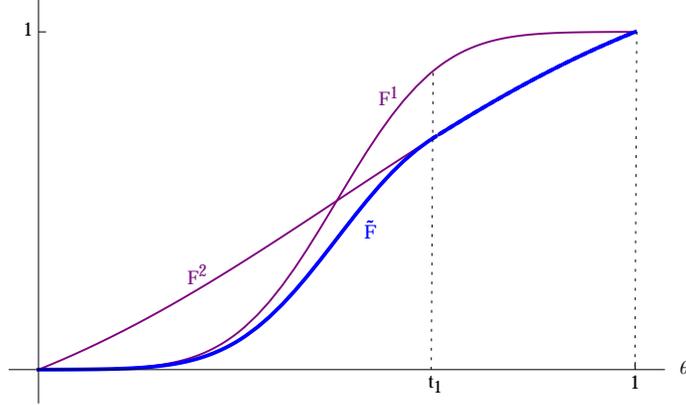


Figure 2: Auxiliary distribution function \tilde{F} for $\{F^1, F^2\}$.

The auxiliary distribution function method not only provides a surprisingly simple characterization of the equilibrium bidding function but also allows for a clean comparison of the properties of the equilibrium outcome of the Dutch auction with that of other auction formats. We show that the auxiliary distribution first-order stochastically dominates the lower envelope of the set of distributions over which bidders minimize; strictly so whenever there is no single distribution in the set that maximizes the reversed hazard rate at each point of the support. Using this feature, we can show that the Dutch auction generates more revenue than the first-price auction, where the ranking is strict if and only if the first-order stochastic dominance is strict. Combining our result with those of the earlier literature (e.g. Lo, 1998) allows us to conclude that the Dutch auction dominates the other standard auction formats in terms of revenue; strictly so for a large class of sets of probability distributions.

Why does strategic equivalence between the Dutch auction and the first-price auction fail? At each point in time, a bidder evaluates her utility of ending the auction at a later time with the conditional distribution under which winning the auction is least likely. Hence, she uses the lower envelope of the set of updated distribution functions to evaluate her continuation payoff. In the first-price auction, bidders receive no information, which means that they use the lower envelope of the set of unconditional distributions. The same is true at the start of the Dutch auction. However, once the price starts descending, bidders discard the possibility of the opponent's valuation being very high and, hence, no longer care about the upper tail of the distribution. Instead they are concerned about the probability of losing the auction if they wait a little longer and evaluate their continuation payoffs with the distribution that maximizes this probability. Since this distribution is in general not the same as the one used in the beginning of the auction, exiting immediately becomes more attractive than it would

have been from an ex-ante point of view. The gradual disclosure of information in the Dutch auction thus leads bidders to end the auction earlier at higher prices, which then benefits the auctioneer. Finally, we extend our results to the case where the auctioneer sets a reserve price and show that the optimal reserve price in the Dutch auction is smaller than in the first-price auction.

The contribution of this paper is twofold. First, we provide a new rationale for the use of descending prices. Besides actual Dutch auctions (e.g., for flowers, treasury bonds, etc.), descending prices appear informally in many situations, such as in the housing market, where asking prices of listed homes decline until a buyer is found. Lucking-Reiley (1999) provides an interesting empirical study on the revenues generated by the different auction formats, using field experiments that auction off collectible trading cards over the Internet. In each experiment, two identical cards are sold: one using a dynamic format (Dutch or English auction), the other using the strategically equivalent static format (first- or second-price auction). While the English auction and the second-price auction produce roughly the same revenue, the Dutch auction generates 30 percent higher revenues than the first-price auction, in line with our theory.² Given that our revenue result holds regardless of the set of probability distributions representing bidders' beliefs—or the seller's knowledge thereof, choosing a descending price auction over another standard format can be viewed as a robust way of increasing revenue. Secondly, we make a methodological contribution by demonstrating a new way of deriving optimal strategies for games in continuous time with preference reversals related to worst-case robustness: after identifying the characteristics of the distribution functions that matter for local optimality conditions, an auxiliary distribution function that shares these characteristics with the worst case is found, mapping the problem into a simpler one with a single distribution. This method is not restricted to dynamic auctions but could, for instance, be applied to a general optimal stopping problem.

The paper proceeds as follows. In Section 2, we describe the environment, in particular the auction game and the solution concept. Section 3 derives the bidding equilibrium, providing the main constructive steps, while Section 4 presents the revenue result and explains the role of reserve prices. Section 5 discusses the updating rule and explains in detail the connection of our work with the existing literature. Section 6 concludes.

²Although some of the revenue differences between Dutch and first-price auction in the experiments can be attributed to different participation rates, Lucking-Reiley (1999) shows that individual bidders violate strategic equivalence in a way that favors the Dutch auction in terms of revenue.

2 Environment

A good is sold using a descending price auction. There are two identical bidders, who perceive the distribution of their valuation to be subjectively unknown. They have maxmin preferences, identified by a set of distribution functions \mathcal{F} . We assume that \mathcal{F} is the convex hull of a set of differentiable distribution functions $\{F^n(\theta)\}_{n \in \mathcal{N}}$, indexed by a compact set \mathcal{N} . We further assume that $F^n(\theta)$ is continuous in n , $f^n(\theta) \equiv F^{n'}(\theta)$ and $f^{n'}(\theta)$ are jointly continuous in (θ, n) and $f^n(\theta) > 0, \forall \theta \in [0, 1]$.³

The auction starts at a price weakly greater than one; the price then descends continuously until either of the two bidders stops the auction. The bidder who terminates the auction receives the object at the current price. We focus on symmetric equilibria in pure strategies, which can be described by a mapping b from a bidder's type $\theta_i \in [0, 1]$ to a price at which the bidder ends the auction, provided the other bidder has not already stopped the auction. As we discuss at the end of this section, a strategy is a *reduced form* of a complete contingent plan of what to do at every information set (a combination of type and time) in the game.

We assume that bidders use prior-by-prior (or full) Bayesian updating in response to new information.⁴ Under this rule, bidders update each distribution in \mathcal{F} according to Bayes' rule and consider the worst case from the set of updated distributions. Prior-by-prior updating does not satisfy dynamic consistency, which means that the sequentially optimal plan in a decision tree need not coincide with the optimal plan based on prior preferences. Indeed, when retaining consequentialism—the property whereby only the outcomes that are still possible matter for updated preferences—dynamic consistency must be dropped if ambiguity attitudes are not to be restricted (Ghirardato, 2002).⁵ We follow the ‘consistent planning’ approach in the spirit of Strotz (1955-56), and assume that agents are sophisticated in that they correctly anticipate their future preferences. This means that bidders will discard all current plans which their future selves are not willing to carry out. An axiomatic foundation of such sophisticated dynamic choice for maxmin expected utility and prior-by-prior Bayesian updating is provided by Siniscalchi (2011). The assumption that a bidder's optimal plan may change over time plays an important role in the analysis that follows.

Next, we formulate the agents' conditional expected payoffs, given an increasing strategy

³The case of a finite index set \mathcal{N} is thus included as a special case.

⁴See Pires (2002) for an axiomatization of this updating rule.

⁵Another well-known updating rule for multiple priors is maximum likelihood updating (Gilboa and Schmeidler, 1993). As the prior-by-prior updating rule, it violates dynamic consistency. For a further discussion see Section 5.1.

$b(\theta_i)$,⁶ representing the exit price of the other bidder. To keep the notation simple, we will not measure time in terms of the price displayed on the clock but instead in terms of the type who would exit according to the strategy of the other bidder. That is, when we say a decision maker exits at time θ , we mean that she exits when the clock displays price $b(\theta)$.

Consider a bidder of type θ_i and suppose the current time is θ . The bidder's expected payoff from exiting at time $\hat{\theta} \leq \theta$ is given by the difference between her valuation θ_i and the price $b(\hat{\theta})$, weighted by the worst-case conditional probability that the other bidder does not end the auction before $\hat{\theta}$. Provided $\theta_i \geq b(\hat{\theta})$, the latter probability is equal to the minimal conditional probability that the other bidder's value is smaller than $\hat{\theta}$, given by $\min_{F \in \mathcal{F}} F(\hat{\theta})/F(\theta)$.⁷ A bidder's conditional expected payoff can thus be written as:

$$V(\theta_i, \hat{\theta}, \theta) \equiv (\theta_i - b(\hat{\theta})) \min_{F \in \mathcal{F}} F(\hat{\theta})/F(\theta). \quad (1)$$

The fact that the set of distributions over the type of the other player is not conditional on a player's own valuation reflects independence of the valuations of the two players.⁸

An important feature of our environment is that bidders face constraints in their exiting choice imposed by the preferences of their future selves. In particular, viewed from time θ , there may be exit times $\hat{\theta} \leq \theta$ that are not feasible because they are not compatible with the bidder's preferences at later points in time. This can occur either because upon reaching time $\hat{\theta}$, the bidder's future self strictly prefers to continue the auction or because there is some future self that strictly prefers to end the auction before $\hat{\theta}$. In equilibrium, each bidder responds optimally not only to the other bidder's equilibrium strategy but also to the equilibrium strategy of her own future selves.

To formalise this notion of equilibrium for our environment, we fix an increasing bidding function $b(\cdot)$ and define for each type θ_i an exit strategy $x(\cdot, \theta_i) : [0, 1] \rightarrow \{0, 1\}$, indicating whether at time θ type θ_i exits ($x(\theta, \theta_i) = 1$) or continues ($x(\theta, \theta_i) = 0$). Hence, x is a contingent plan of what to do at every information set of the game. Given a plan $x(\theta, \theta_i)$,

⁶Standard arguments can be used to show that the equilibrium bidding function for each player must be strictly increasing.

⁷For ease of exposition, we will use the minimal conditional probability for all exit prices, even though for $b(\hat{\theta}) > \theta_i$ the worst-case scenario is described by the maximal conditional probability. This will not affect any of the results, as no bidder would end the auction at a price strictly above her valuation.

⁸Our assumption is satisfied if before learning their own valuations, both players consider a set of joint distributions which treats players symmetrically and has a product structure. That is, players preferences are described by a set of joint distributions, $\mathcal{F}_{1,2}$, such that for each $F \in \mathcal{F}_{1,2}$, $F(\theta_1, \theta_2) = F_1(\theta_1)F_2(\theta_2)$ and if $F_1(\theta_1)F_2(\theta_2) \in \mathcal{F}_{1,2}$, then $F_2(\theta_1)F_1(\theta_2) \in \mathcal{F}_{1,2}$ (see Gilboa and Schmeidler (1989) and Bade (2008) for a discussion of independence under maxmin preferences).

we can define a number $T(\theta, \theta_i)$, indicating how long a bidder of type θ_i would remain in the auction, provided that she continues at time θ . Setting $T(\theta, \theta_i) \equiv \sup\{\hat{\theta} < \theta : x(\hat{\theta}, \theta_i) = 1\}$, we are ready to define an equilibrium for our game.

Definition 1. *Fix an increasing and symmetric bidding strategy $b(\theta_i)$ which satisfies $b(0) = 0$. The strategy constitutes an equilibrium if for every $\theta_i \in [0, 1]$, there exists an upper-semi continuous function $x(\cdot, \theta_i)$ with $x(0, \theta_i), x(\theta_i, \theta_i) = 1$ and $\forall \theta > \theta_i, x(\theta, \theta_i) = 0$, which satisfies the following conditions.*

1) For all $\theta \in (0, 1]$, $x_i(\theta, \theta_i) = 1$ only if either

(i) $T(\theta, \theta_i) < \theta$ and $V(\theta_i, \theta, \theta) \geq V(\theta_i, T(\theta, \theta_i), \theta)$, or

(ii) $T(\theta, \theta_i) = \theta$ and $V(\theta_i, \hat{\theta}, \theta) \geq V(\theta_i, \theta', \theta)$ for all θ' in a left-neighbourhood of θ .

2) For all $\theta \in (0, 1]$, $x(\theta, \theta_i) = 0$ only if $V(\theta_i, \theta, \theta) \leq V(\theta_i, T(\theta, \theta_i), \theta)$.

Requiring $x(\theta_i, \theta_i) = 1$ and $\forall \theta > \theta_i, x(\theta, \theta_i) = 0$ ensures that type θ_i exits the auction at time θ_i . The two additional conditions require that a bidder has no profitable one-stage deviations at any information set, given the plan of future selves. To check such deviations, we compare the payoff at θ with the payoff at $T(\theta, \theta_i)$, the time type θ_i would exit, provided that she remains at time θ .⁹

While this definition considers only one-stage deviations, one could strengthen it to requiring that, given any information set, a bidder cannot gain by deviating to an alternative continuation plan that she is willing to follow through (i.e. that satisfies the no-one-stage-deviation conditions). We call equilibria that satisfy this additional requirement ‘‘consistent-planning’’ equilibria. The requirement is satisfied if, keeping fixed the equilibrium strategy of the other player, there is a unique plan that satisfies conditions 1) and 2). We will show that the unique equilibrium in differentiable strategies we characterise in Proposition 4 is also a consistent-planning equilibrium.

3 Equilibrium

We look for a symmetric equilibrium, where the strategies of both players are described by a differentiable bidding function $b(\theta)$. Towards this goal, a closer inspection of the bidders’

⁹If there is a left neighbourhood of θ on which the plan specifies ‘exit’, then $T(\theta, \theta_i) = \theta$. Condition (ii) of Definition 1 guaranties that a bidder cannot gain by waiting an instance longer.

conditional payoff $V(\theta_i, \hat{\theta}, \theta)$ is needed. We first consider the question of when a bidder can improve her utility locally by remaining in the auction a little longer. We show that at each point in time θ , a bidder evaluates her utility of continuing the auction for a marginal amount of time with the distribution function $F \in \mathcal{F}$ that maximizes the reversed hazard rate $f(\theta)/F(\theta)$.

Lemma 1. *For all $\theta \in (0, 1]$,*

$$\frac{d^- \min_{F \in \mathcal{F}} (F(\hat{\theta})/F(\theta))}{d\hat{\theta}} \Big|_{\hat{\theta}=\theta} = \max_{n \in \mathcal{N}} \frac{f^n(\theta)}{F^n(\theta)}.$$

The proof of this and all other results can be found in the Appendix. The lemma is essentially an application of an envelope theorem in Milgrom and Segal (2002).

For future reference, let $m(\theta)$ denote the function that assigns to any θ the index of the distribution which maximises the reversed hazard rate (making an arbitrary selection if there is more than one distribution). Lemma 1 establishes that at time θ , the decision maker evaluates her utility of continuing the auction—at least locally—according to the distribution $F^{m(\theta)}$. To see why a bidder uses the distribution with the maximal reversed hazard rate, notice that the marginal cost of waiting for a lower price is the marginal decrease in the conditional probability of winning the auction, $\min_{F \in \mathcal{F}} F(\hat{\theta})/F(\theta)$. Since at $\hat{\theta} = \theta$ we have $F(\hat{\theta})/F(\theta) = 1$ for all $F \in \mathcal{F}$, the distribution function that locally minimizes the bidder's utility from waiting is the one with the largest derivative of $F(\hat{\theta})/F(\theta)$ at $\hat{\theta} = \theta$. This derivative is given by the reversed hazard rate $f(\theta)/F(\theta)$.

3.1 Candidate equilibrium bidding function

With this observation, we can construct a candidate equilibrium bidding function. We look for a function $b(\theta_i)$ that is strictly increasing and differentiable in the bidder's type. By Lemma 1, $\min_{F \in \mathcal{F}} F(\hat{\theta})/F(\theta)$ is differentiable in $\hat{\theta}$ on a left neighbourhood of θ , so we consider the following first-order condition:

$$\frac{\partial V^-(\theta_i, \hat{\theta}, \theta_i)}{\partial \hat{\theta}} \Big|_{\hat{\theta}=\theta_i} = 0. \tag{2}$$

Using Lemma 1, the first-order condition can be written as:

$$b'(\theta_i) = (\theta_i - b(\theta_i)) f^{m(\theta_i)}(\theta_i) / F^{m(\theta_i)}(\theta_i). \tag{3}$$

The fact that $m(\theta_i)$ depends on θ_i implies that the condition $f^{m(\theta_i)}(\theta_i) = dF^{m(\theta_i)}(\theta_i)/d\theta_i$ is not in general satisfied. Due to the failure of this property, the standard approach of solving for the corresponding bidding function cannot be directly applied. The complication arises because the distribution that locally minimizes a bidder's continuation payoff changes throughout the auction. To address the issue of shifting worst-case scenarios in continuous time, we follow a new approach: we describe a solution to the first-order condition in terms of an as-if distribution, denoted by \tilde{F} , which has a reversed hazard rate equal to $\max_{n \in \mathcal{N}} f^n(\theta)/F^n(\theta)$ for every $\theta \in [0, 1]$. The following lemma establishes that such distribution function exists.

Lemma 2. *Let $\tilde{F}(0) \equiv 0$ and for $\theta > 0$,*

$$\tilde{F}(\theta) \equiv \exp\left(-\int_{\theta}^1 \max_{n \in \mathcal{N}} \frac{f^n(t)}{F^n(t)} dt\right).$$

Then $\tilde{F}(\theta)$ is a cumulative distribution function with a well-defined density function $\tilde{f}(\theta) \equiv \tilde{F}'(\theta)$. For every $\theta \in (0, 1]$, the reversed hazard rate of \tilde{F} satisfies:

$$\frac{\tilde{f}(\theta)}{\tilde{F}(\theta)} = \max_{n \in \mathcal{N}} \frac{f^n(\theta)}{F^n(\theta)}.$$

As an immediate corollary of the previous lemma, we can rewrite the first-order condition in terms of the auxiliary distribution function \tilde{F} :

$$b'(\theta_i) = (\theta_i - b(\theta_i))\tilde{f}(\theta_i)/\tilde{F}(\theta_i). \quad (4)$$

Hence, the first-order condition (3) is the same as the one of an expected-utility decision maker who believes \tilde{F} is the distribution over types. Using standard techniques, we can then characterise the solution to the first-order condition in terms of \tilde{F} .

Lemma 3. *Suppose an equilibrium bidding strategy $b^*(\theta_i)$ satisfies the first-order condition (3). Then*

$$b^*(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \tilde{F}(\theta) d\theta}{\tilde{F}(\theta_i)}. \quad (5)$$

Figure 3 shows $F^{m(\theta)}(\theta)$ and $\tilde{F}(\theta)$ for a set of truncated normals with support $[0, 1]$, fixed mean, and variances σ belonging to a closed interval $[\underline{\sigma}, \bar{\sigma}]$. The distribution function \tilde{F} is obtained by multiplying the function $F^{m(\theta)}(\theta)$ with an increasing function $\alpha : [0, 1] \rightarrow [0, 1]$.

The transformation of $F^{m(\theta)}(\theta)$ to $\tilde{F}(\theta)$ is particularly easy to understand for the case where the index $m(\theta)$ changes discretely in θ , as we show in the following example.

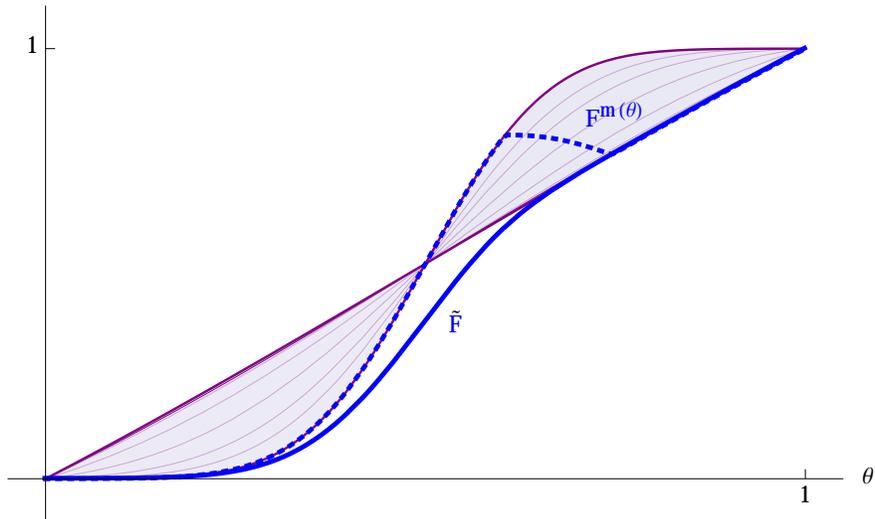


Figure 3: The functions $F^{m(\theta)}(\theta)$ and $\tilde{F}(\theta)$

Example 1. Assume that the type space can be partitioned into $M \in \mathbb{N}$ intervals such that on each interval the reversed hazard rate is maximized by a single distribution belonging to \mathcal{F} . We consider the coarsest partition and label the distribution that maximises the reversed hazard rate on the m 'th interval by F^m , $m = 1, 2, \dots, M$. Hence, $m(\theta)$ is the index of the interval to which θ belongs. We define \bar{t}_m as the least upper bound of the interval indexed by m . The auxiliary distribution function is then described as follows:

$$\tilde{F}(\theta) = \alpha_{m(\theta)} F^{m(\theta)}(\theta),$$

where $\alpha_M = 1$ and for $m < M$,

$$\alpha_m = \prod_{i=m}^{M-1} \frac{F^{i+1}(\bar{t}_i)}{F^i(\bar{t}_i)}.$$

When there is a finite number of distributions in \mathcal{F} maximizing the reversed hazard rate, the auxiliary distribution $\tilde{F}(\theta)$ is thus obtained by multiplying $F^{m(\theta)}(\theta)$ with an increasing step function. It is easy to verify that \tilde{F} is a CDF of a probability distribution with the same reversed hazard rate as $F^{m(\theta)}(\theta)$.¹⁰ Figures 1 and 2 in the Introduction provide a graphical illustration for the case $M = 2$.

¹⁰See Appendix A.8 for details.

3.2 Verifying equilibrium

The following proposition establishes that there is a symmetric equilibrium where the strategies of the two bidders are described by our candidate bidding function.

Proposition 4. *The bidding function $b^*(\theta_i)$, as described in (5), constitutes the unique symmetric equilibrium in differentiable bidding strategies.*

As argued above, the candidate equilibrium bidding function b^* satisfies the local optimality condition for each type of bidder. It thus remains for us to verify that it is also globally optimal for type θ_i to terminate the auction at time θ_i . Existing results for the model without ambiguity show that this is true if bidders evaluate their utility with \tilde{F} at each point in time. In this case, a bidder's expected payoff as a function of the exit time is single peaked in equilibrium, meaning that if the bidder's type is θ_i , her payoff increases from waiting at any point in time before θ_i and decreases at any point in time after θ_i . Bidders in our framework, however, do not evaluate their utility with \tilde{F} but with the distribution that minimizes $F(\hat{\theta})/F(\theta)$. Locally, \tilde{F} is simply a scaled version of that distribution, but globally this need not be the case. As a result, there can be situations where the function V is no longer single peaked.

In order to verify the equilibrium in our environment, we use the feature whereby bidders anticipate the optimal behavior of their future selves. This allows us to follow a backward induction argument. We know that since bidders locally evaluate their utility as if they would be using \tilde{F} , a bidder's utility must be locally decreasing from continuing the auction once the time reaches her type. Now suppose the bidder's type is θ_i and there is a point in time $\theta < \theta_i$, which, viewed from time θ_i , would promise a strictly higher expected payoff than exiting at time θ_i . Then at time $\theta + \varepsilon$, with ε sufficiently small, the bidder will strictly prefer not to wait until time θ and terminates the auction immediately. Anticipating this, the same bidder at a slightly earlier time will find it optimal to exit immediately, and so on, until we reach time θ_i . The backwards induction procedure is particularly simple in the case of a finite set of distributions in \mathcal{F} maximizing the reverse hazard rate.

Example 1. *(continued) As we show in Appendix A.8, conditional on reaching an interval indexed by $m = 1, 2, \dots, M$, a bidder evaluates exit times belonging to m with the distribution function that maximizes the reversed hazard rate on m , F^m . In other words, for all $\hat{\theta} \in$*

$[\bar{t}_{m(\theta)-1}, \theta]$, the following property holds:

$$\min_{F \in \mathcal{F}} \frac{F(\hat{\theta})}{F(\theta)} = \frac{F^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\theta)} = \frac{\tilde{F}(\hat{\theta})}{\tilde{F}(\theta)}.$$

The latter equality follows from the fact that in a given interval, the distribution function \tilde{F} is just a scaled version of $F^{m(\theta)}$. Hence, within each interval, a bidder acts like an expected utility maximizer whose subjective prior is \tilde{F} . Assuming the other player bids according to b^* , this implies that the expected payoff function is single-peaked within the interval. Consider now a bidder of type θ_i at time $\theta < \theta_i$. If θ belongs to the first interval, the bidder optimally ends the auction immediately, as her payoff is strictly decreasing on the interval below θ_i . Anticipating this, the same argument can be applied to the second interval, and so on, until we reach time θ_i .

Finally, we can show that type θ_i has no incentives to terminate the auction before time θ_i . To this end, note that at time $\theta > \theta_i$, the payoff obtained when exiting immediately is $\theta_i - b^*(\theta)$, whereas the anticipated payoff from waiting for time θ_i is $(\theta_i - b^*(\theta_i)) \min_{F \in \mathcal{F}} F(\theta_i)/F(\theta)$. As discussed earlier, if the bidders were to evaluate their expected payoffs with \tilde{F} , waiting for θ_i would be optimal. It is then useful to recall that reversed hazard rate dominance implies conditional stochastic dominance (see, for instance, Shaked and Shanthikumar, 1994). This means that the conditional probability $\tilde{F}(\theta_i)/\tilde{F}(\theta)$ is smaller than $F(\theta_i)/F(\theta)$ for all $F \in \mathcal{F}$. The expected payoff associated with exiting at time θ_i , evaluated at some $F \in \mathcal{F}$, is thus weakly greater than that evaluated at \tilde{F} . This implies that no matter which distribution function type θ_i uses at time θ to assess the option of waiting, the associated payoff is strictly greater than the payoff from ending the auction immediately. Remaining in the auction until the time is θ_i is therefore optimal and consistent with the behaviour of later selves.

Remark: The proof of Proposition 4 shows that in equilibrium, at any point in time, the unique contingent plan without profitable one-shot deviations is to remain before the suggested exit time and to exit from then on. This means that at no point in time, there is an alternative consistent plan that would improve a bidders' contingent payoff. Hence, the bidding strategy b^* not only satisfies conditions 1) and 2) of Definition 1 but also the stronger notion of consistent planning. To prove this, we show that for any type θ_i and time $\theta < \theta_i$, there is an interval $(\underline{\theta}, \bar{\theta})$ with $\underline{\theta} < \theta < \bar{\theta}$ such that within the interval a bidder anticipating

to exit no later than time $\underline{\theta}$ strictly prefers to exit immediately. Thus, despite the continuous time setup, we can use backward induction over a set of intervals, similar to the case of finitely many distributions.¹¹

Though bidders have no desire to end the auction before the clock indicates their type, they may wish to continue longer than their future selves permit. This feature arises whenever the worst-case distribution changes over time. At the start of the game, the worst-case scenario associated to ending the auction at time θ_i is described by the cumulative distribution function with the smallest value at θ_i . At time θ_i , on the other hand, the worst-case scenario is the probability distribution that has the highest reversed hazard rate at θ_i . When these two distributions do not coincide, the ex-ante optimal exit time differs from the contingent optimal exit time, as illustrated in Figure 4. In this case, a bidder would benefit from a commitment device that precludes her from ending the auction too early. However, in the absence of such a device, the bidder will act optimally with regard to the set of updated distribution functions, which is to exit when the clock indicates her type. This discrepancy plays an important role for our results on the revenue properties of the Dutch auction, presented in the following section.

4 Revenue Comparison

In the first-price auction, bidder's receive no information about the other bidder's type. Their worst case scenario is thus described by the lower envelope of unconditional distribution functions in \mathcal{F} , defined as $F^{min}(\theta) \equiv \min_{n \in \mathcal{N}} F^n(\theta), \theta \in [0, 1]$. Indeed, Lo (1998) shows that the equilibrium bidding function for the first-price sealed bid auction is:

$$b^{FPA}(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F^{min}(\theta) d\theta}{F^{min}(\theta_i)}.$$

Bidders thus behave as if they share a common single prior belief F^{min} . Lo (1998) further shows that the first-price auction generates more revenue than the second-price auction. In the second-price auction, bidding one's valuation remains a dominant strategy, even if bidders face ambiguity about their opponent's valuation. Likewise, in an English auction

¹¹The continuous time setup presents a challenge as that, at every given point in time, there is no "next" (respectively, previous) instance of time. We overcome this challenge by showing that, even without knowing what will be done in the immediate future, there will be a uniquely optimal action at any point in time.

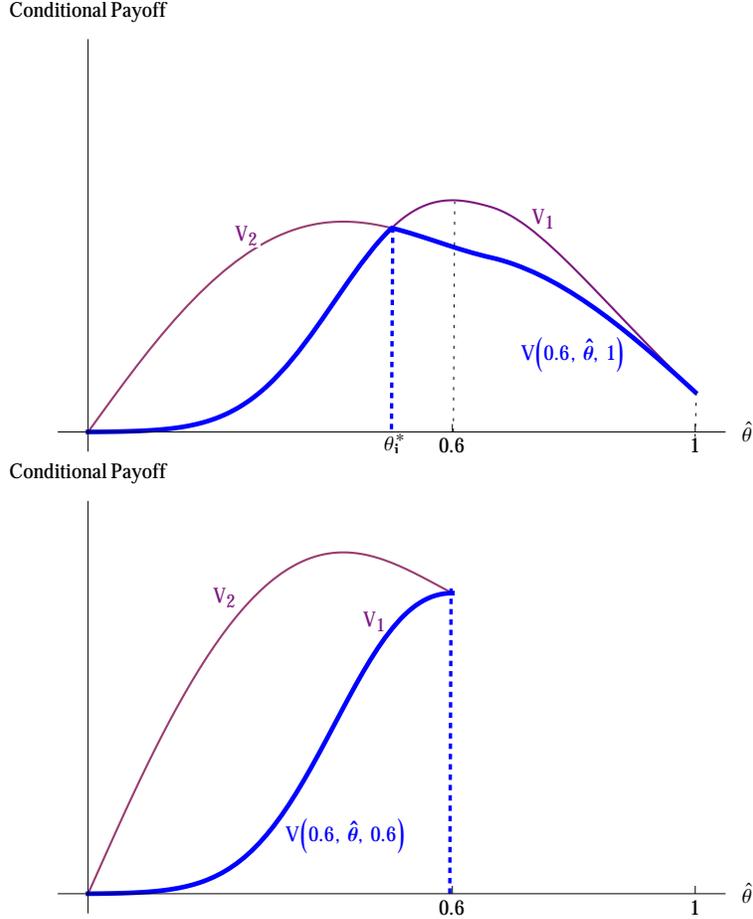


Figure 4: Max min expected utility at two different clock times and $\mathcal{N} = \{1, 2\}$. V_1 and V_2 are the conditional payoffs, evaluated respectively at F^1 and F^2 , for type $\theta_i = 0.6$. The upper figure shows the bidder's maxmin payoff (in blue) at time $\theta = 1$, which is maximized at θ_i^* to the left of 0.6. The lower figure shows the maxmin payoff at time 0.6, where the bidder optimally exits.

with an ascending price, exiting the auction when the price reaches one's valuation is the dominant strategy, no matter how beliefs are updated.¹² English and second-price auctions therefore continue to generate the same revenue. We now show that the Dutch auction dominates all of the other standard auction formats in terms of revenue, in particular, the first-price auction.

Proposition 5. *For all $\theta_i \in [0, 1]$,*

$$b^*(\theta_i) \geq b^{FPA}(\theta_i).$$

¹²See also Karni and Safra (1989).

The inequality holds as equality for all $\theta_i \in [0, 1]$ if and only if there exists a distribution in \mathcal{F} that stochastically dominates all other distributions in the set according to the reversed hazard rate order.

By construction, the distribution function \tilde{F} dominates the distribution function F^{min} in terms of the reversed hazard rate. As discussed earlier, stochastic dominance according to the reversed hazard rate implies conditional stochastic dominance. Hence, we have:

$$\frac{\tilde{F}(\theta)}{\tilde{F}(\theta')} \leq \frac{F^{min}(\theta)}{F^{min}(\theta')} \quad \text{for all } \theta, \theta' \text{ with } \theta \leq \theta'.$$

Bidders participating in the Dutch auction thus behave as if they are faced with a probability distribution which, conditional on having reached a given point in time, first-order stochastically dominates F^{min} . The distribution functions \tilde{F} and F^{min} are the same if there is a distribution in \mathcal{F} that maximizes the reversed hazard rate for all values of θ . When this is not the case, the worst-case scenario changes over time and bidders in the Dutch auction exit sooner than would be optimal from an ex-ante point of view. Intuitively, as time progresses, the bidder learns that the other bidder's type is below the currently displayed time and evaluates her payoff with the distribution that—given this information—makes it most likely that the opponent ends the auction if she waits a little longer. This leads bidders to drop out earlier at higher prices compared to the first-price auction, where no information is conveyed.

The question is in which situation is there a single distribution that maximizes the reversed hazard rate over all distributions in the set. Eeckhoudt and Gollier (1995) show that the monotone likelihood ratio (MLR) order implies the reversed hazard rate order (but is not implied by it). The latter order will often be satisfied if the distributions in \mathcal{F} only vary in their mean. For instance, Milgrom (1981) states that the families of normal, exponential and Poisson distributions with different means—all else being equal—satisfy the MLR order, and hence the reversed hazard rate order. However, the case where agents only perceive ambiguity about the first moment of the distribution is rather special. When there is ambiguity about higher moments, MPR will typically fail. Figure 3 illustrates this for the case of truncated normal distributions with changing variances.¹³

¹³The case where agents know the mean but perceive ambiguity about higher moments of the distribution has been the subject of analysis in several recent papers, for example, Wolitzky (2016) and Carrasco et al. (2018).

4.1 Reserve Prices

We know that an auctioneer can increase her expected revenue by setting a reserve price r below which the good will not be sold. In our environment, the construction of the candidate equilibrium bidding function can be extended to the case where the Dutch auction ends at time $r \in [0, 1]$. A slight modification of the proof of Proposition 5 verifies that this function indeed constitutes an equilibrium.¹⁴

Corollary 6. *Suppose the Dutch auction ends at price r . The equilibrium bidding function is described by:*

$$b_r^*(\theta_i) = \theta_i - \int_r^{\theta_i} \frac{\tilde{F}(\theta)}{\tilde{F}(\theta_i)} d\theta.$$

With regard to the revenue comparison between the Dutch auction and the first-price auction, we can also adapt the proof we used in the model without reserve prices, and show that for any given reserve price r , the Dutch auction generates weakly more revenue than the first-price auction. The revenue ranking is strict if the function \tilde{F} strictly first-order stochastically dominates F^{min} on the restricted domain $[r, 1]$. This happens if there is no distribution in the set \mathcal{F} that maximizes the reserve hazard rate on the interval $[r, 1]$. Thus, if the Dutch auction generates strictly higher revenues than the first-price auction without reserve prices, it also yields strictly higher revenues with a reserve price, as long as the reserve price is not too high.

The fact that the Dutch auction generates higher revenues for any given reserve price implies that it must also lead to higher revenues if we set the reserve price optimally. An interesting question is how the optimal reserve prices in the two auction formats relate. Thus far, we have not had to make any assumptions about the beliefs or preferences of the seller. Since the equilibrium bid in the Dutch auction exceeds the equilibrium bid in the first-price auction for each type, it is irrelevant how the seller aggregates her utility over the possible types. For the optimal reserve price, however, the seller's preferences and beliefs matter.

For simplicity, we consider the case where the seller is risk- and ambiguity-neutral. Her subjective belief about a bidder's type is described by the cumulative distribution function F^S , assumed to be differentiable on its domain $[0, 1]$ with density function f^S . For instance, the seller might consider an average over all distributions in \mathcal{F} (according to some unspecified

¹⁴The modification of the equilibrium definition to accommodate reserve prices is straightforward. Given a reserve price r , types below r will never participate, making the effective type space become $[r, 1]$.

second-order belief).¹⁵ The probability that the higher of the two valuations is at most θ is given by $(F^S(\theta_i))^2$. Thus, the seller's expected revenue Π can be written as:

$$\Pi(r) = \int_r^1 \left(\theta_i - \frac{\int_r^{\theta_i} \tilde{F}(\theta) d\theta}{\tilde{F}(\theta_i)} \right) d(F^S(\theta_i))^2,$$

where $d(F^S(\theta_i))^2 = 2F^S(\theta_i)dF^S(\theta_i)$. The following proposition characterizes the reserve price that maximizes this revenue function and relates the optimal reserve prices in the two auction formats to each other.

Proposition 7. *The optimal reserve price in the Dutch auction, denoted by r^* , is characterized by the condition:*

$$\frac{F^S(r^*)}{\tilde{F}(r^*)} r^* = \frac{1}{f^S(r^*)} \int_{r^*}^1 \frac{F^S(\theta_i)}{\tilde{F}(\theta_i)} dF^S(\theta_i). \quad (6)$$

The reserve price r^ is weakly smaller than the optimal reserve price in the first-price auction, strictly so whenever $\tilde{F}(r^*) < F^{min}(r^*)$.*

The optimal reserve price r^* solves the auctioneer's tradeoff between increasing the probability of selling the item and increasing the price in the event that he does. More specifically, a marginal increase in r raises the equilibrium bids of all types greater than r , but also expands the set of types who prefer not to participate in the auction. The condition characterizing the price that solves this tradeoff, (6), depends on the distribution functions F^S and \tilde{F} . When these two functions are the same, equation (6) simplifies to the familiar condition:

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

The optimal reserve price in the first-price auction is characterized by the same condition as the one in the Dutch auction, with the function F^{min} taking the place of \tilde{F} . Proposition 7 states that in the first-price auction the seller optimally sets a reserve price higher than the optimal reserve price in the Dutch auction. This result comes from the property that bidders in the Dutch auction terminate the auction at higher prices compared to the first-price auction. As a consequence, the cost of excluding certain types is greater, shifting the

¹⁵It is straightforward to extend the characterization to the case where the seller has maxmin preferences. If she minimizes over the same set of distributions as the bidders, she evaluates her payoff according to $F^{max}(\theta) = \max_{n \in \mathcal{N}} F^n(\theta)$, the distribution that is first-order stochastically dominated by all other distributions in the set.

auctioneer’s tradeoff in favor of lower reserve prices.

5 Discussion

5.1 Prior-by-prior updating

The non-equivalence between the Dutch auction and the first-price auction in our setting is closely linked to the dynamic inconsistency of the prior-by-prior updating rule. There are alternative updating rules of ambiguous beliefs that restore dynamic consistency, at the expense of other desirable properties.¹⁶ For instance, Klibanoff and Hanany (2007) propose a dynamically consistent updating rule for the maxmin model, where some distributions from the initial set of distributions over which the decision maker minimizes are discarded. For the Dutch auction this means that as the price descends, bidders have to dispose of those distributions under which exiting before the ex-ante optimal time is preferable. It should be noted, however, that eliminating a subset of distributions in \mathcal{F} cannot solely be defended on the grounds of received information. Indeed, bidders with different valuations discard different sets of probability distributions after the same amount of time has passed (and hence after receiving exactly the same information). For an illustration, consider Example 1 with two subintervals ($M = 2$). Types whose valuation falls into the first subinterval discard probability distributions once the price starts descending (see Figure 5), whereas types belonging to the second subinterval maintain the whole set \mathcal{F} . The reason for this is that types in the second subinterval face no time-inconsistency issues; the conditional expected payoff is evaluated with the distribution that maximizes the reversed hazard rate in the second subinterval, both at the start of the auction and at the time they exit. Hence, under the dynamically consistent updating rule, bidders in our auction behave as if the ambiguity they perceive diminishes over time, the extent to which depends on their (independent) payoff type. The latter feature might be difficult to reconcile if ambiguity is to be understood as an informational phenomenon (for a further discussion see Siniscalchi (2009)).

¹⁶As we mentioned in Section 2, if ambiguity attitudes are not to be restricted, any updating rule must dispense with one of two natural requirements: dynamic consistency or consequentialism. Dropping consequentialism means that updated preferences may depend on the outcomes of events that are no longer possible. Hence, a decision maker violating consequentialism can have opposing preferences over choices that are identical on all non-null events. Dominiak et al. (2012) provide evidence showing that ambiguity-averse subjects tend to drop dynamic consistency over consequentialism.

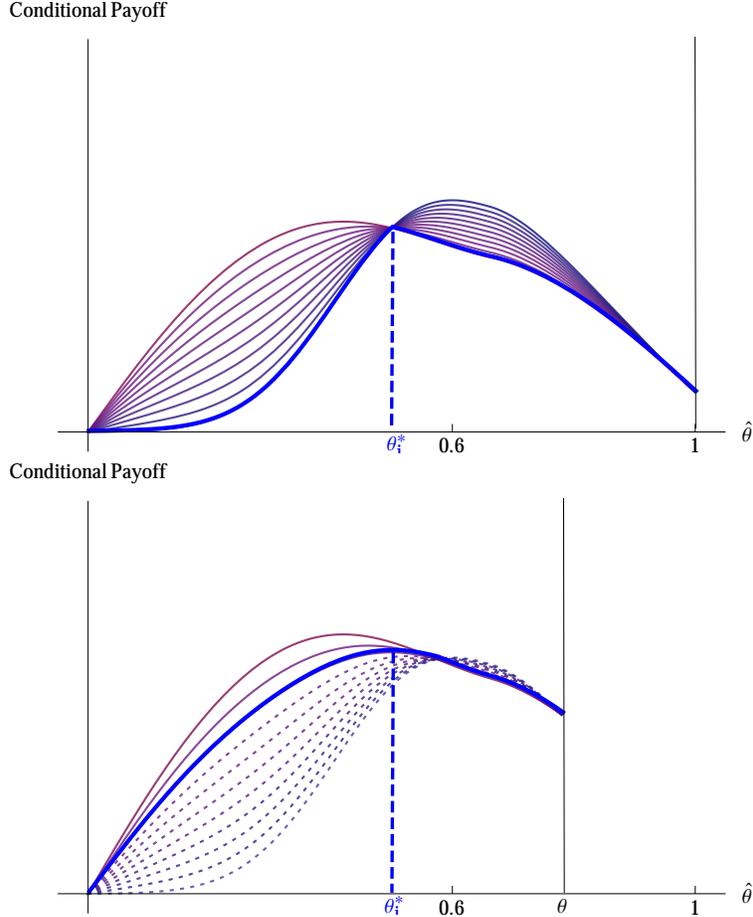


Figure 5: Dynamically consistent updating in the Dutch auction for $\mathcal{N} = \{1, 2\}$, $M = 2$, and $\theta_i = 0.6 < \bar{t}_1$. The upper figure shows the bidder’s expected payoff at time one, evaluated at F^1 , F^2 , and convex combinations thereof. The ex-ante optimal exit time is denoted by θ_i^* . The lower figure shows the bidder’s payoff at time $\theta < 1$ after updating. All probability distributions under which exiting earlier than θ_i^* is optimal (dotted curves) are discarded.

5.2 Relation to the Literature

There is a growing literature on ambiguity and auctions. While most work in that literature is focussed on static settings,¹⁷ some papers have considered dynamic auctions as well. Karni (1988) is the first to show that the equivalence of first-price and descending price auctions

¹⁷Early work includes Salo and Weber (1995) and Lo (1998), who show revenue non-equivalence between static auctions in different frameworks, while Bose et al. (2006) analyze the optimal static auction. Bodoh-Creed (2012), Wolitzky (2016), and De Castro and Yannelis (2018) consider the general problem of static mechanism design with maxmin agents and study different applications. In contrast to these papers, Bose and Renou (2014) and Di Tillio et al. (2016) show that even if type distributions are not ambiguous, the mechanism designer can benefit from introducing ambiguity to the mechanism that is carried out.

relies on bidders being dynamically consistent, however without providing a description of equilibrium bidding when dynamic consistency fails. Nakajima (2011) finds that if preferences exhibit the typical choice pattern of the Allais paradox, the Dutch auction will result in higher bids and thus higher revenue. Conversely, Weber (1998) studies a non-expected utility preference of a form that displays a reversal of the typical Allais choices, which could be more empirically relevant for small stakes. For such non-expected utility preferences, the first-price auction generates higher revenues. Closest to our paper is the work by Bose and Daripa (2009), who emphasize the benefits of using dynamic mechanisms when agents are ambiguity-averse. They show that when bidders' preferences are represented by the ε -contamination model—a special case of the maxmin expected utility model—a discretized Dutch auction with a particular tie-breaking rule can extract almost all of the surplus. As Bose and Daripa (2009) emphasize, the full surplus extraction result hinges on the assumed preference representation. In particular, the result relies on the property whereby bidders assign at least an ε weight to the event that the other bidder terminates the auction when the price decreases by a certain amount, no matter how small that amount is. This requirement is not only violated when the set of distributions functions is a singleton (the standard case) but, more generally, when bidders minimize over a finite set of smooth distribution functions (or the convex hull thereof). We view our work as complementary to Bose and Daripa (2009), in that we show the profitability of dynamic mechanisms even when full surplus extraction is not attainable.

The Dutch auction is an instance of a dynamic game of incomplete information with ambiguity-averse players. Some recent papers discuss general solution concepts and their properties. Battigalli et al. (2017) explore self-confirming equilibria in dynamic games with players whose preferences are represented by the smooth ambiguity model of Klibanoff et al. (2005). Like us, they permit dynamic inconsistency and assume that players are sophisticated. Also Hanany et al. (2018) consider multistage games with incomplete information and smoothly ambiguity-averse players. In contrast to Battigalli et al. (2017), they require equilibria to satisfy sequential optimality and show that such games are as if players update in a dynamically consistent way. In their framework, suitably extended to continuous time, the Dutch and first-price auctions thus remain strategically equivalent. Finally, Pahlke (2018) considers a setting in which players have maxmin expected utility and update their beliefs prior by prior, as do we. In contrast to our framework, she allows the set of probability distributions over which players minimize to depend on the game that is played in such a way that dynamic consistency is maintained within the game. For our setting, this

requires that players expand the initial set of probabilities \mathcal{F} so as to include the auxiliary distribution function \tilde{F} when faced with the Dutch auction but not when playing a first-price auction. The interpretation changes in this case, while the main insights do not.

There is also some work on dynamic games with ambiguity-averse preferences in applications unrelated to auctions. For instance, Kellner and Le Quement (2017, 2018) study the effects of exogenous and endogenous ambiguity in a cheap-talk setting; Beauchene, et al. (2019) consider the case of persuasion. As in our work, both consider consistent planning equilibria with prior-by-prior Bayesian updating.

6 Conclusion

The paper shows that the consideration of uncertainty over probability distributions of bidders' valuations has important implications for equilibrium bidding in dynamic auctions. The gradual revelation of information about the opponent's valuation in the Dutch auction can cause a bidder's worst-case scenario to shift over time. We show that this leads bidders to purchase the good at higher prices compared to the case when no information is transmitted, which then benefits the auctioneer.

Throughout the analysis, we maintain the assumption that bidders are symmetric with respect to the set of distributions over which they minimize. A question for future research is how the equilibrium outcome is affected when we allow for asymmetries across bidders. Of particular interest are situations where more experienced bidders compete against less experienced ones and this difference is reflected in the degree of ambiguity they perceive.

Another important question concerns the optimal design of dynamic mechanisms in environments with ambiguity-averse agents. While Bose and Daripa (2009) show the possibility of full surplus extraction for certain cases, a comprehensive answer to the question does not yet exist. We saw for our setting that different mechanisms induce different distributions that describe an agent's 'as if' beliefs. It would be interesting to explore whether this feature can be used to map the problem of designing an optimal dynamic mechanism into a simpler one over a set of such distributions.

A Appendix

A.1 Proof of Lemma 1

Proof. We start by proving that $\min_{F \in \mathcal{F}} F(\hat{\theta})/F(\theta) = \min_{n \in \mathcal{N}} F^n(\hat{\theta})/F^n(\theta)$. To see this, consider \bar{n} and n such that $F^{\bar{n}}(\hat{\theta})/F^{\bar{n}}(\theta) \leq F^n(\hat{\theta})/F^n(\theta)$, i.e. $F^n(\theta)F^{\bar{n}}(\hat{\theta})/F^{\bar{n}}(\theta) \leq F^n(\hat{\theta})$. Then, for any $\alpha(n) \in [0, 1]$, $\alpha(n)F^n(\theta)F^{\bar{n}}(\hat{\theta})/F^{\bar{n}}(\theta) \leq \alpha(n)F^n(\hat{\theta})$. Integrating over n , $\int_{n \in \mathcal{N}} \alpha(n)F^n(\theta)F^{\bar{n}}(\hat{\theta})/F^{\bar{n}}(\theta) \, dn \leq \alpha(n) \int_{n \in \mathcal{N}} F^n(\hat{\theta}) \, dn$, for any $\alpha : \mathcal{N} \rightarrow (0, 1)$ and thus $\frac{F^{\bar{n}}(\hat{\theta})}{F^{\bar{n}}(\theta)} \leq \frac{\int_{n \in \mathcal{N}} \alpha(n)F^n(\hat{\theta}) \, dn}{\int_{n \in \mathcal{N}} \alpha(n)F^n(\theta) \, dn}$ for any $\alpha : \mathcal{N} \rightarrow (0, 1]$.

Hence, if for all n , $F^{\bar{n}}(\hat{\theta})/F^{\bar{n}}(\theta) \leq F^n(\hat{\theta})/F^n(\theta)$, it follows that also for all $F \in \mathcal{F}$ it holds that $F^{\bar{n}}(\hat{\theta})/F^{\bar{n}}(\theta) \leq F(\hat{\theta})/F(\theta)$, as desired.

Note next that, for all $\theta > 0$, $f^n(\hat{\theta})/F^n(\theta) = d \frac{F^n(\hat{\theta})}{F^n(\theta)} / d\hat{\theta}$ is continuous in $(\hat{\theta}, n)$ and $F^n(\hat{\theta})/F^n(\theta)$ is continuous in n . Corollary 4(ii) in Milgrom and Segal (2002) applies to a minimisation problem by exchanging the “max” operator to a “min” (and vice versa). Letting $n^*(\hat{\theta}, \theta) \equiv \arg \min_{n \in \mathcal{N}} (F^n(\hat{\theta})/F^n(\theta))$, it establishes that for all $\theta \in (0, 1]$,

$$\frac{d^- \min_{n \in \mathcal{N}} (F^n(\hat{\theta})/F^n(\theta))}{d\hat{\theta}} = \max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\theta)}. \quad (7)$$

Since for $\hat{\theta} = \theta$, $F^n(\hat{\theta})/F^n(\theta) = 1$, irrespective of n , it is true that $n^*(\theta, \theta) = \mathcal{N}$, so that the desired result follows. □

A.2 Proof of Lemma 2

Proof. For $\theta > 0$, let $\rho(\theta) \equiv \max_{n \in \mathcal{N}} \frac{f^n(\theta)}{F^n(\theta)}$. By the Maximum Theorem, ρ is continuous. For any $n \in \mathcal{N}$ and $\theta > 0$, we can write $f^n(\theta)/F^n(\theta) = d \ln F^n(\theta) / d\theta$. An integration yields

$$\int_{\theta}^1 \frac{f^n(t)}{F^n(t)} \, dt = [\ln F^n(t)]_{\theta}^1,$$

and hence

$$F(\theta) = \exp \left(- \int_{\theta}^1 \frac{f(t)}{F(t)} \, dt \right), \quad (8)$$

so that $\int_{\theta}^1 \frac{f^n(t)}{F^n(t)} \, dt$ tends to infinity as θ approaches 0. Thus, $\int_{\theta}^1 \rho(t) \, dt$ tends to infinity, and \tilde{F} is continuous, even at $\theta = 0$. Given $\rho(\theta) > 0$ for all $\theta > 0$, $\tilde{F}(\theta)$ is non-decreasing. Recall

that

$$\tilde{F}(\theta) \equiv \exp\left(-\int_{\theta}^1 \rho(t)dt\right),$$

Since ρ is continuous, it follows from Leibnitz' Integral rule that

$$\tilde{f}(\theta) = \rho(\theta) \exp\left(-\int_{\theta}^1 \rho(t)dt\right) = \rho(\theta)\tilde{F}(\theta),$$

which establishes continuity of \tilde{f} and the desired equality. □

A.3 Proof of Lemma 3

Given that \tilde{F} is differentiable on $(0, \theta_i)$ we can integrate (4) to:

$$(b(\theta)\tilde{F}(\theta))\Big|_0^{\theta_i} = \int_0^{\theta_i} \theta \tilde{f}(\theta) d\theta.$$

A reformulation yields

$$b(\theta_i) = \frac{\int_{t(\theta_i)}^{\theta_i} \theta \tilde{f}(\theta) d\theta}{\tilde{F}(\theta_i)},$$

and integration by parts can be used to express b as in the lemma.

A.4 Proof of Proposition 4

Let $\mu(\theta)$ denote the correspondence of maximisers of the maximisation problem $\max_{n \in \mathcal{N}} (f^n(\theta)/F^n(\theta))$.

Outline

The proof begins with a lemma which will restrict the nature of possible preference reversals given any increasing bidding strategy. It then identifies four properties of the objective function of the bidders given the putative equilibrium strategy at various points in time, and finishes by using these properties to show that if a player bids according to b^* , there is a no-one-stage-deviation plan for the other player such that she exits at θ_i . Lastly, we show, using a backward-induction type argument, that given the equilibrium bidding strategy the no one-stage-deviation plan is unique, ensuring that it is also a consistent planning equilibrium.

For use below, for any $n \in N$, let $V_n(\theta_i, \hat{\theta}, \theta) \equiv (\theta_i - b(\hat{\theta}))F^n(\hat{\theta})/F^n(\theta)$.

Lemma 8. Fix arbitrary $\bar{\theta}, \theta \in [0, 1]$ such that $\bar{\theta} < \theta^* < \theta$. Then

(i) If there exists an $n \in N$ such that $V_n(\theta_i, \bar{\theta}, \theta) < V_n(\theta_i, \theta^*, \theta)$, then $V(\theta_i, \bar{\theta}, \theta^*) < V(\theta_i, \theta^*, \theta^*)$.

(ii) If $V(\theta_i, \bar{\theta}, \theta) < V(\theta_i, \theta^*, \theta)$, then $V(\theta_i, \bar{\theta}, \theta^*) < V(\theta_i, \theta^*, \theta^*)$.

Proof. To prove i), suppose indeed for some $n \in N$,

$$(\theta_i - b(\bar{\theta}))F^n(\bar{\theta})/F^n(\theta) < (\theta_i - b(\theta^*))F^n(\theta^*)/F^n(\theta).$$

Then,

$$(\theta_i - b(\bar{\theta}))F^n(\bar{\theta})/F^n(\theta^*) < \theta - b(\theta^*),$$

so that

$$(\theta_i - b(\bar{\theta})) \min_{n \in \mathcal{N}} F^n(\bar{\theta})/F^n(\theta^*) < \theta_i - b(\theta^*).$$

To prove ii), suppose $V(\theta_i, \bar{\theta}, \theta) < V(\theta_i, \theta^*, \theta)$, i.e.

$$\min_{n \in \mathcal{N}} (\theta_i - b(\bar{\theta}))F^n(\bar{\theta})/F^n(\theta) < \min_{n \in \mathcal{N}} (\theta_i - b(\theta^*))F^n(\theta^*)/F^n(\theta),$$

and since $n^*(\bar{\theta}, \theta_i)$ is one of the minimisers of the problem on the left hand side and in the choice set of the problem on the right hand side of the above equation,

$$(\theta_i - b(\bar{\theta}))F^{n^*(\bar{\theta}, \theta_i)}(\bar{\theta})/F^{n^*(\bar{\theta}, \theta_i)}(\theta) < (\theta_i - b(\theta^*))F^{n^*(\bar{\theta}, \theta_i)}(\theta^*)/F^{n^*(\bar{\theta}, \theta_i)}(\theta),$$

and thus $V_{n^*(\bar{\theta}, \theta_i)}(\theta_i, \bar{\theta}, \theta) < V_{n^*(\bar{\theta}, \theta_i)}(\theta_i, \theta^*, \theta)$. Point (i) can be used to establish the result. \square

In what follows, fix an arbitrary $\theta_i \in (0, 1]$ and assume $b(\theta) = \theta - \int_0^\theta \tilde{F}(t)dt/\tilde{F}(\theta)$.

The case of $\theta > \theta_i$

Lemma 9. For any θ such that $\theta > \theta_i$, the function $V(\theta_i, \hat{\theta}, \theta)$ is decreasing in $\hat{\theta}$ for all $\hat{\theta} \in (\theta_i, \theta]$.

Proof. Let $H(\hat{\theta}, \theta)$ denote the minimal probability that the other bidder's type is less than $\hat{\theta}$ at time θ , i.e. $H(\hat{\theta}, \theta) = \min_{n \in \mathcal{N}} \frac{F^n(\hat{\theta})}{F^n(\theta)}$.

$$\begin{aligned}
\frac{d^-V(\theta_i, \hat{\theta}, \theta)}{d\hat{\theta}} &= (\theta_i - b(\hat{\theta})) d^-H(\hat{\theta}, \theta)/d\hat{\theta} - b'(\hat{\theta})H(\hat{\theta}, \theta) \\
&= \left(\theta_i - \hat{\theta} + \frac{\int_0^{\hat{\theta}} \tilde{F}(t)dt}{\tilde{F}(\hat{\theta})} \right) dH(\hat{\theta}, \theta)/d\hat{\theta} - \frac{\tilde{f}(\hat{\theta}) \int_0^{\hat{\theta}} \tilde{F}(t)dt}{\tilde{F}^2(\hat{\theta})} H(\hat{\theta}, \theta) \quad (9) \\
&= (\theta_i - \hat{\theta}) \frac{d^-H(\hat{\theta}, \theta)}{d\hat{\theta}} - \frac{\int_0^{\hat{\theta}} \tilde{F}(t)dt}{\tilde{F}(\hat{\theta})} \left(H(\hat{\theta}, \theta) \frac{\tilde{f}(\hat{\theta})}{\tilde{F}(\hat{\theta})} - \frac{d^-H(\hat{\theta}, \theta)}{d\hat{\theta}} \right).
\end{aligned}$$

Using equation (7), we get

$$\begin{aligned}
\frac{d^-V(\theta_i, \hat{\theta}, \theta)}{d\hat{\theta}} &= (\theta_i - \hat{\theta}) \max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\theta)} - \frac{\int_0^{\hat{\theta}} \tilde{F}(t)dt}{\tilde{F}(\hat{\theta})} \left(H(\hat{\theta}, \theta) \frac{\tilde{f}(\hat{\theta})}{\tilde{F}(\hat{\theta})} - \max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\theta)} \right) \\
&= (\theta_i - \hat{\theta}) \max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\theta)} - \frac{\int_0^{\hat{\theta}} \tilde{F}(t)dt}{\tilde{F}(\hat{\theta})} H(\hat{\theta}, \theta) \left(\frac{\tilde{f}(\hat{\theta})}{\tilde{F}(\hat{\theta})} - \max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\theta)} \right) \quad (10)
\end{aligned}$$

Note that the last step follows from

$$\max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\theta)} / H(\hat{\theta}, \theta) = \max_{n \in n^*(\hat{\theta}, \theta)} \frac{f^n(\hat{\theta})}{F^n(\hat{\theta})},$$

which holds by definition of $H(\hat{\theta}, \theta)$ and $n^*(\hat{\theta}, \theta)$. The desired result follows from observing that $V(\theta_i, \hat{\theta}, \theta)$ is continuous in $\hat{\theta}$, that the first term of equation (10) is negative if $\theta > \theta_i$, and that the second term is always weakly negative by Lemma 2. \square

The case of $\theta < \theta_i$

Lemma 10. *Fix an arbitrary $\theta_0 \in (0, \theta_i)$. Then, there exists some $\bar{\theta}$ and $\underline{\theta}$ such that $\theta_i > \bar{\theta} > \theta_0 > \underline{\theta} \geq 0$ and $\forall \theta' \in (\underline{\theta}, \bar{\theta}]$ and $\forall \hat{\theta} \in [\underline{\theta}, \theta']$ it is true that $V(\theta_i, \hat{\theta}, \theta') < V(\theta_i, \theta', \theta')$.*

Proof. Fix $\theta_0 \in (0, \theta_i)$ and any $n_0 \in \mu(\theta_0)$. Similar to equation (9), one can derive:

$$\frac{\partial V_{n_0}(\theta_i, \hat{\theta}, \theta)}{\partial \hat{\theta}} = (\theta_i - \hat{\theta}) \frac{f^{n_0}(\hat{\theta})}{F^{n_0}(\theta)} - \frac{\int_0^{\hat{\theta}} \tilde{F}(t)dt}{\tilde{F}(\hat{\theta})} \frac{F^{n_0}(\hat{\theta})}{F^{n_0}(\theta)} \left(\frac{\tilde{f}(\hat{\theta})}{\tilde{F}(\hat{\theta})} - \frac{f^{n_0}(\hat{\theta})}{F^{n_0}(\hat{\theta})} \right). \quad (11)$$

Note that this function is continuous in $\hat{\theta}$ and θ . Since $\frac{F^{n_0}(\theta_0)}{F^{n_0}(\theta_0)} = 1$ and $\frac{f^{n_0}(\theta_0)}{F^{n_0}(\theta_0)} = \frac{\tilde{f}(\theta_0)}{\tilde{F}(\theta_0)}$, we

have:

$$\frac{\partial V_{n_0}(\theta_i, \theta_0, \theta_0)}{\partial \hat{\theta}} = (\theta_i - \theta_0) \max_{n \in \mathcal{N}} \frac{f^n(\theta_0)}{F^n(\theta_0)} > 0.$$

Hence, there exist $\theta^* > \theta_0$ such that it is true that $\left. \frac{\partial V_{n_0}(\theta_i, \hat{\theta}, \theta^*)}{\partial \hat{\theta}} \right|_{\hat{\theta}=\theta_0} > 0$. Since V_{n_0} is differentiable in its second argument, there exists an interval $[\underline{\theta}, \bar{\theta}]$ around θ_0 such that $\bar{\theta} \leq \theta^*$ and $V_{n_0}(\theta_i, \hat{\theta}, \theta^*)$ is increasing in $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$. Thus, $\forall \theta' \in (\underline{\theta}, \bar{\theta}]$ and $\forall \hat{\theta} \in [\underline{\theta}, \theta')$ it is true that $V_{n_0}(\theta_i, \hat{\theta}, \theta^*) < V_{n_0}(\theta_i, \theta', \theta^*)$. The application of Lemma 8(i) concludes the proof. \square

The case of $\theta = \theta_i$

Lemma 11. *There exists a $\underline{\theta} < \theta_i$ such that $\forall \theta' \in (\underline{\theta}, \theta_i]$ and $\hat{\theta} \in [\underline{\theta}, \theta')$, it is true that $V(\theta_i, \hat{\theta}, \theta') < V(\theta_i, \theta', \theta')$.*

Proof. Let $n_0 \in \mu(\theta_i)$ satisfy

$$d \frac{f^{n_0}(\theta_i)}{F^{n_0}(\theta_i)} / d \theta_i = \min_{n \in \mu(\theta_i)} d \frac{f^n(\theta_i)}{F^n(\theta_i)} / d \theta_i.$$

We show now that if $\theta_0 = \theta_i$, then both $\partial V_{n_0}(\theta_i, \theta_i, \theta_i) / \partial \hat{\theta} = 0$ and $\partial^{2-} V_{n_0}(\theta_i, \theta_i, \theta_i) / \partial^2 \hat{\theta} < 0$. Observe first that $\partial V_{n_0}(\theta_i, \theta_i, \theta_i) / \partial \hat{\theta} = 0$ simply follows from (11) since $\hat{\theta} = \theta_i$ so that the first term vanishes (while, as before, the second term vanishes by Lemma 2). Note that by Corollary 4.ii in Milgrom and Segal (2002),

$$\frac{d^- \frac{\tilde{f}(\theta_i)}{F(\theta_i)}}{d \theta_i} = \frac{d^- \left(\max_{n \in \mathcal{N}} \frac{f^n(\theta_i)}{F^n(\theta_i)} \right)}{d \theta_i} = \min_{n \in \mu(\theta_i)} \frac{d \left(\frac{f^n(\theta_i)}{F^n(\theta_i)} \right)}{d \theta_i} = \frac{d \left(\frac{f^{n_0}(\theta_i)}{F^{n_0}(\theta_i)} \right)}{d \theta_i}.$$

Thus, from equation (11) we conclude that $\partial^{2-} V_{n_0}(\theta_i, \theta_i, \theta_i) / \partial^2 \hat{\theta} = -\frac{f^{n_0}(\theta_i)}{F^{n_0}(\theta_i)} < 0$. Thus, there exists $\underline{\theta} \in (0, \theta_i)$ such that $\forall \theta' \in (\underline{\theta}, \theta_i]$ and $\hat{\theta} \in [\underline{\theta}, \theta')$, $V_{n_0}(\theta_i, \hat{\theta}, \theta_i) < V_{n_0}(\theta_i, \theta', \theta_i)$. The results follows from Lemma 8(i). \square

The case of θ close to zero

Lemma 12. *There is a $\bar{\theta} \in (0, \underline{\theta})$ such that $\forall \theta' \in (0, \bar{\theta}]$ and $\forall \hat{\theta} \in [0, \theta')$ it is true that $V(\theta_i, \hat{\theta}, \theta') < V(\theta_i, \theta', \theta')$.*

Proof. To show this, fix some arbitrary $\theta_i > 0$ and $n_0 \in N$. Note that $V_{n_0}(\theta_i, \hat{\theta}, \theta_i)$ is differentiable in $\hat{\theta}$, $V_{n_0}(\theta_i, 0, \theta_i) = 0$ and, since $\forall \theta > 0$ it is true that $b(\hat{\theta}) < \hat{\theta}$ we have $V_{n_0}(\theta_i, \hat{\theta}, \theta_i) > 0$ and there exists a $\bar{\theta} \in (0, \underline{\theta})$ such that $V_{n_0}(\theta_i, \hat{\theta}, \theta_i)$ is increasing in $\hat{\theta} \forall \hat{\theta}_i \in [0, \bar{\theta}]$. Thus, $\forall \theta' \in (0, \bar{\theta}]$ and $\forall \hat{\theta} \in [0, \theta']$, $V_{n_0}(\theta_i, \hat{\theta}, \theta_i) < V_{n_0}(\theta_i, \theta', \theta_i)$. The desired inequality now follows from Lemma 8(i). □

No one-shot deviations

Fix arbitrary $\theta_i > 0$. The proof proceeds by showing that the function $x(\cdot, \theta_i) \rightarrow \{0, 1\}$ satisfies conditions 1) and 2) of Definition 1 if $x(\theta, \theta_i) = 1$ if $\theta \leq \theta_i$ and $x(\theta, \theta_i) = 0$ otherwise. Since such a function x trivially satisfies the other requirements of this definition, this establishes equilibrium. Note that for any $\theta < \theta_i$, by the definition of x , $x(\theta, \theta_i) = 1$ and $T(\theta, \theta_i) = \theta$ and thus Lemma 10 establishes that case point (ii) of condition 1) obtains whenever $x(\theta, \theta_i) = 1$. Condition 2 holds vacuously. For $\theta = \theta_i$ the same argument is established via Lemma 11. Now consider $\theta > \theta_i$, $x(\theta, \theta_i) = 0$ and $T(\theta, \theta_i) = \theta_i$. Then condition 1) holds vacuously and Lemma 9 establishes that condition 2) is satisfied.

Uniqueness

To show uniqueness in symmetric, differentiable strategies, note that in any such equilibrium, for any θ_i it is true that $T(\theta', \theta_i) = \theta_i$ if $\theta' \geq \theta_i$. Fix an arbitrary $\theta_i > 0$. We will first show that in any equilibrium $\left. \frac{\partial V^-(\theta_i, \hat{\theta}, \theta_i)}{\partial \hat{\theta}} \right|_{\hat{\theta}=\theta_i} \geq 0$. Let $\bar{\theta} \equiv T(\theta_i, \theta_i)$. If $\bar{\theta} = \theta_i$ the claim follows from the definition of an equilibrium and the fact that, for differentiable strategies, $V(\theta_i, \hat{\theta}, \theta_i)$ is (left-) differentiable in the second argument. Else, suppose that $\bar{\theta} < \theta_i$ and, by contradiction, $\left. \frac{\partial V^-(\theta_i, \hat{\theta}, \theta_i)}{\partial \hat{\theta}} \right|_{\hat{\theta}=\theta_i} < 0$. By the definition of equilibrium, $V(\theta_i, \bar{\theta}, \theta_i) \leq V(\theta_i, \theta_i, \theta_i)$. Since $V(\theta_i, \hat{\theta}, \theta_i)$ is locally decreasing in $\hat{\theta}$, there is $\theta^* \in (\bar{\theta}, \theta_i)$ such that $V(\theta_i, \bar{\theta}, \theta_i) < V(\theta_i, \theta^*, \theta_i)$. By Lemma 8(ii), $V(\theta_i, \bar{\theta}, \theta^*) < V(\theta_i, \theta^*, \theta^*)$, contradicting the definition of $\bar{\theta}$.

We next show that the derivative in question cannot be strictly positive (and hence, equals zero). So suppose, by way of contradiction, that

$$\left. \frac{\partial V^-(\theta_i, \hat{\theta}, \theta_i)}{\partial \hat{\theta}} \right|_{\hat{\theta}=\theta_i} > 0. \tag{12}$$

This implies that

$$(\theta_i - b(\theta_i)) \max_{n \in \mathcal{N}} \frac{f^n(\theta)}{F^n(\theta_i)} - b'(\theta_i) > 0.$$

Thus, there exists a $\theta' > \theta_i$ such that

$$(\theta_i - b(\theta_i)) \max_{n \in \mathcal{N}} \frac{f^n(\theta)}{F^n(\theta')} - b'(\theta_i) > 0.$$

Thus, for some $\theta^* \in (\theta_i, \theta')$, $V(\theta_i, \theta_i, \theta') < V(\theta_i, \theta^*, \theta')$. By Lemma 8(ii), also $V(\theta_i, \theta_i, \theta^*) < V(\theta_i, \theta^*, \theta^*)$, so that given $T(\theta^*, \theta^i) = \theta_i$, $x(\theta^*, \theta_i) = 1$ in violation of Definition 1.

Consistent Planning

Finally, we show that the function $x(\cdot, \theta_i) \rightarrow \{0, 1\}$ satisfies conditions 1) and 2) of Definition 1 only if $x(\theta, \theta_i) = 1$ if $\theta \leq \theta_i$ and $x(\theta, \theta_i) = 0$ otherwise and, hence, that the consistent planning property is satisfied.

Assume conditions 1) and 2) of Definition 1 are met. Lemma 12 says that there is a $\bar{\theta} \in (0, \underline{\theta})$ so that $x(\theta, \theta_i) = 1$ if $\theta \in (0, \bar{\theta})$, i.e. for low enough values of θ , the bidder will find it optimal to exit irrespective of the behavior of future selves.

Now, consider any θ in $(\bar{\theta}, \underline{\theta})$. We will argue that $x(\theta, \theta_i) = 1$. Note that Lemma 10 assigns to any point $\theta_0 \in [\bar{\theta}, \underline{\theta}]$ two values $\underline{\theta}(\theta_0)$ and $\bar{\theta}(\theta_0)$ such that $\underline{\theta}(\theta_0) < \theta_0 < \bar{\theta}(\theta_0)$. Hence, the set of intervals

$$\{(\underline{\theta}(\theta_0), \bar{\theta}(\theta_0))\}_{\theta_0 \in [\bar{\theta}, \underline{\theta}]}$$

form an open cover of the compact set $[\bar{\theta}, \underline{\theta}]$, and thus have a finite subcover. Consider now this subcover, a finite collection of (overlapping, open) intervals.

Lemma 10 also says that for any $\theta \in (\underline{\theta}(\theta_0), \bar{\theta}(\theta_0))$, provided that $T(\underline{\theta}(\theta_0), \theta_i) \geq \underline{\theta}(\theta_0)$ it holds that $x(\theta, \theta_i) = 1$. Since, $x(\theta, \theta_i) = 1$ for any $\theta \leq \bar{\theta}$ it is true that for any θ in the first interval, if $\underline{\theta}$ denotes its lower bound, $T(\underline{\theta}, \theta_i) \geq \underline{\theta}$. Thus for all θ within the first interval, $x(\theta, \theta_i) = 1$. Also, given $x(\theta, \theta_i) = 1$ for all values weakly below an interval, the same argument establishes that $x(\theta, \theta_i) = 1$ also within this interval. Thus, by induction, $\forall \theta \leq \underline{\theta}, x(\theta, \theta_i) = 1$.

Since $\forall \theta \in [\underline{\theta}, \theta_i]$, $T(\theta, \theta_i) \geq \underline{\theta}$, Lemma 11 then implies that $\forall \theta \in [\underline{\theta}, \theta_i], x(\theta, \theta_i) = 1$. Lemma 9 then implies that for all $\theta > \theta_i$, since $T(\theta, \theta_i) \geq \underline{\theta}$, it is true that $x(\theta, \theta) = 0$ as desired.

□

A.5 Proof of Proposition 5

The condition $b^*(\theta_i) \geq b^{FPA}(\theta_i)$ for all $\theta_i \in [0, 1]$ is equivalent to $\int_0^{\theta_i} F^{min}(\theta) d\theta / F^{min}(\theta_i) \geq \int_0^{\theta_i} \tilde{F}(\theta) d\theta / \tilde{F}(\theta_i)$ for all $\theta_i \in [0, 1]$. A sufficient condition for this inequality to be satisfied is:

$$\frac{F^{min}(\theta)}{F^{min}(\theta_i)} \geq \frac{\tilde{F}(\theta)}{\tilde{F}(\theta_i)}, \quad (13)$$

for all θ_i and all $\theta \leq \theta_i$. Since dominance in the reversed hazard rate implies conditional stochastic dominance (see Shaked and Shanthikumar, 1994), this condition is satisfied.

As we have shown in Appendix A.2, a cumulative distribution function $F(\theta)$ is uniquely determined by its reversed hazard rate $\rho_F(\theta)$:

$$F(\theta) = \exp\left(-\int_{\theta}^1 \rho_F(t) dt\right).$$

If there is a distribution function F belonging to \mathcal{F} such that $\rho_F(\theta) = \max_{n \in \mathcal{N}} f^n(\theta) / F^n(\theta)$, then this distribution function is \tilde{F} . Hence, $F^{min}(\theta) \leq \tilde{F}(\theta), \forall \theta$. By (13), we have $F^{min}(\theta) \geq \tilde{F}(\theta), \forall \theta$ and hence $F^{min} = \tilde{F}$.

Now suppose there is no single distribution in \mathcal{F} that maximizes the reversed hazard rate on the interval $[0, 1]$. We can then find some $\theta_0 \in (0, 1)$ such that the same property holds on $[\theta_0, 1]$. Let $n \in \mathcal{N}$ be such that $F^n(\theta_0) = F^{min}(\theta_0)$. Then there exists an open interval $(\underline{\theta}, \bar{\theta})$ with $\theta_0 \leq \underline{\theta} < \bar{\theta}$ such that $n \neq \arg \max_{n' \in \mathcal{N}} f^{n'}(\theta) / F^{n'}(\theta)$. For any $\theta \in (\underline{\theta}, \bar{\theta})$, we thus have $f^n(\theta) / F^n(\theta) < \tilde{f}(\theta) / \tilde{F}(\theta)$ (while this inequality always holds as a weak inequality). Hence,

$$\begin{aligned} F^n(\theta_0) &= \exp\left(-\int_{\theta_0}^1 f^n(\theta) / F^n(\theta) d\theta\right) \\ &> \exp\left(-\int_{\theta_0}^1 \tilde{f}(\theta) / \tilde{F}(\theta) d\theta\right) \\ &= \tilde{F}(\theta_0), \end{aligned}$$

so that $F^{min}(\theta_0) > \tilde{F}(\theta_0)$. It follows that $b^*(1) > b^{FPA}(1)$. By continuity of b^* and b^{FPA} , the same inequality holds on a left neighbourhood of 1. \square

A.6 Proof of Corollary 6

To show that b_r^* constitutes an equilibrium, one follows the steps of the proof of Proposition 4 in Appendix A.4, simply replacing zero with r as the lower bound of the domain of integration in each of the shown integrals.

A.7 Proof of Proposition 7

We start by computing:

$$\begin{aligned}\Pi'(r) &= 2 \left(-F^S(r) f^S(r) r + \tilde{F}(r) \int_r^1 F^S(\theta_i) / \tilde{F}(\theta_i) dF^S(\theta_i) \right) \\ \Pi''(r) &= 2 \left(-f^S(r)^2 r - F^S(r) f^{S'}(r) r - 2F^S(r) f^S(r) + \tilde{f}(r) \int_r^1 F^S(\theta_i) / \tilde{F}(\theta_i) dF^S(\theta_i) \right).\end{aligned}$$

Since $\Pi(r)$ is continuous on the compact domain $[0, 1]$, a maximum exists. We can first show that Π is maximized on the interior of the interval $[0, 1]$. Noticing that $\Pi''(0)$ (to be seen as the right derivative) is strictly positive, we know that $r = 0$ does not constitute a maximum point. Similarly, $\Pi'(1)$ (to be seen as the left derivative) is strictly negative, which implies that $r = 1$ does not maximize Π either. Differentiability of Π then implies that the maximum point is characterized by the first-order condition. Setting $\Pi'(r)$ equal to zero, a simple manipulation yields condition (6).

Next, we want to compare the optimal reserve price in the Dutch auction with the optimal reserve price in the first-price auction. Lo (1998) shows that the equilibrium bidding function in the first-price auction with reserve price r is $b_r^{FPA}(\theta) = \theta - \frac{\int_r^\theta F^{min}(\theta) d\theta}{F^{min}(\theta)}$. Analogous to above, we can write the seller's revenue in the first-price auction as a function of the reserve price r :

$$\Pi_{FPA}(r) = 2 \int_r^1 F^S(\theta_i) \left(\theta_i - \frac{\int_r^{\theta_i} F^{min}(\theta) d\theta}{F^{min}(\theta_i)} \right) dF^S(\theta_i).$$

If we consider the difference between the two revenues and differentiate with respect to r , we obtain

$$\Pi'(r) - \Pi'_{FPA}(r) = \int_r^1 F^S(\theta_i) \left(\frac{F^{min}(r)}{F^{min}(\theta_i)} - \frac{\tilde{F}(r)}{\tilde{F}(\theta_i)} \right) dF^S(\theta_i). \quad (14)$$

We know that reversed hazard rate dominance implies conditional stochastic dominance (see,

for instance, Shaked and Shanthikumar, 1994). The fact that $\tilde{f}(\theta_i)/\tilde{F}(\theta_i) \geq f^{min}(\theta_i)/F^{min}(\theta_i)$ holds for all $\theta_i > r$ thus implies

$$\frac{\tilde{F}(r)}{\tilde{F}(\theta_i)} \geq \frac{F^{min}(r)}{F^{min}(\theta_i)}, \quad \text{for all } \theta_i > r. \quad (15)$$

It follows that the difference of the first derivatives in (14) is negative. Since

$$\Pi(r^*) - \Pi_{FPA}(r^*) \geq \Pi(r_{FPA}^*) - \Pi_{FPA}(r^*) \geq \Pi(r_{FPA}^*) - \Pi_{FPA}(r_{FPA}^*),$$

and the difference between $\Pi(r)$ and $\Pi_{FPA}(r)$ is weakly decreasing in r , it follows that r^* must be weakly smaller than r_{FPA}^* . Moreover, if $F^{min}(r^*) < \tilde{F}(r^*)$, the difference in (14) is strictly negative, implying that $\Pi'_{FPA}(r^*)$ is strictly positive. In that case the optimal reserve price in the first-price auction, r_{FPA}^* is strictly larger than r^* . \square

A.8 Further details for Example 1

We start by showing that $\tilde{F}(\theta) = \alpha_m F^m(\theta)$, where $\alpha_{m(\theta)}$ is defined in Example 1. Since $\alpha_M = 1$ and $F^{m(0)}(0) = 0$, we have $\tilde{F}(1) = 1$ and $\tilde{F}(0) = 0$. Let θ belong to the interval indexed by m and consider:

$$\begin{aligned} \lim_{\theta \rightarrow \bar{t}_{m-1}^+} \alpha_m F^m(\theta) &= \prod_{i=m}^{M-1} \frac{F^{i+1}(\bar{t}_i)}{F^i(\bar{t}_i)} F^m(\bar{t}_{m-1}), \\ &= \prod_{i=m}^{M-1} \frac{F^m(\bar{t}_{m-1})}{F^{m-1}(\bar{t}_{m-1})} \frac{F^{i+1}(\bar{t}_i)}{F^i(\bar{t}_i)} F^{m-1}(\bar{t}_{m-1}), \\ &= \prod_{i=m-1}^{M-1} \frac{F^{i+1}(\bar{t}_i)}{F^i(\bar{t}_i)} F^{m-1}(\bar{t}_{m-1}), \\ &= \alpha_{m-1} F^{m-1}(\bar{t}_{m-1}). \end{aligned}$$

Since α_m is constant on θ in each subinterval, \tilde{F} is differentiable on the interior of each subinterval. Differentiability at the boundaries of the subintervals requires that for each $\bar{t}_m, m = 1, \dots, M-1$ the derivative of $\alpha_m F^m(\theta)$ with respect to θ is equal to that of $\alpha_{m+1} F^{m+1}(\theta)$. That is,

$$\alpha_m f^m(\bar{t}_m) = \alpha_{m+1} f^{m+1}(\bar{t}_m).$$

By definition of \bar{t}_m , we have $f^m(\bar{t}_m)/F^m(\bar{t}_m) = f^{m+1}(\bar{t}_m)/F^{m+1}(\bar{t}_m)$. Together with $\alpha_m = F^{m+1}(\bar{t}_m)/F^m(\bar{t}_m)\alpha_{m+1}$, this implies that the above equality holds for all $\bar{t}_m, m = 1, \dots, M-1$. Further, since all functions in $\{F^m\}_{m=1}^M$ are strictly increasing in θ and α_m is strictly positive for all θ , \tilde{F} is strictly increasing on its domain. Finally, given that for each subinterval α_m is just a scaling factor, we clearly have $\frac{\tilde{f}(\theta)}{\tilde{F}(\theta)} = \frac{f^m(\theta)}{F^m(\theta)}, \forall \theta \in [0, 1]$.

Next, we will show that bidders evaluate exit times in the current interval with the auxiliary distribution function \tilde{F} :

$$\min_{n \in \mathcal{N}} \frac{F^n(\hat{\theta})}{F^n(\theta)} = \frac{F^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\theta)} = \frac{\tilde{F}(\hat{\theta})}{\tilde{F}(\theta)}, \quad \forall \hat{\theta} \in [\bar{t}_{m(\theta)-1}, \theta].$$

To prove the first equality, suppose by contradiction that there exists an $n \in \mathcal{N}$ such that $\frac{F^n(\hat{\theta})}{F^n(\theta)} < \frac{F^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\theta)}$ for some $\hat{\theta} \in [\bar{t}_{m(\theta)-1}, \theta]$. Define the difference function

$$D_{\theta,n}(\hat{\theta}) := \frac{F^n(\hat{\theta})}{F^n(\theta)} - \frac{F^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\theta)}$$

on $[\bar{t}_{m(\theta)-1}, \theta]$ with

$$D'_{\theta,n}(\hat{\theta}) = \frac{f^n(\hat{\theta})}{F^n(\theta)} - \frac{f^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\theta)}.$$

Notice that $D_{\theta,n}(\hat{\theta})$ is continuous on $[\bar{t}_{m(\theta)-1}, \theta]$ and differentiable on $(\bar{t}_{m(\theta)-1}, \theta)$. By assumption, we have $D_{\theta,n}(\hat{\theta}) < 0$. This, together with $D_{\theta,n}(\theta) = 0$, implies that we can then find some $\hat{\theta}_0 \in (\hat{\theta}, \theta)$ such that $D_{\theta,n}(\hat{\theta}_0) < 0$ and

$$D'_{\theta,n}(\hat{\theta}_0) = \frac{f^n(\hat{\theta}_0)}{F^n(\hat{\theta}_0)} \frac{F^n(\hat{\theta}_0)}{F^n(\theta)} - \frac{f^{m(\theta)}(\hat{\theta}_0)}{F^{m(\theta)}(\hat{\theta}_0)} \frac{F^{m(\theta)}(\hat{\theta}_0)}{F^{m(\theta)}(\theta)} > 0.$$

However, since $\frac{f^n(\hat{\theta})}{F^n(\hat{\theta})} \leq \frac{f^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\hat{\theta})}$ for all $\hat{\theta} \in (\bar{t}_{m(\theta)-1}, \theta)$, by definition of $F^{m(\theta)}$, this requires $\frac{F^n(\hat{\theta}_0)}{F^n(\theta)} > \frac{F^{m(\theta)}(\hat{\theta}_0)}{F^{m(\theta)}(\theta)}$ and therefore $D_{\theta,n}(\hat{\theta}_0) > 0$. A contradiction. Hence, $\min_n \frac{F^n(\hat{\theta})}{F^n(\theta)} = \frac{F^{m(\theta)}(\hat{\theta})}{F^{m(\theta)}(\theta)}$ for all $\hat{\theta} \in [\bar{t}_{m(\theta)-1}, \theta]$.

References

- [1] Bade, Sophie (2008): “Stochastic independence with maxmin expected utilities,” Working Paper.
- [2] Battigalli, P., Catonini, E., Lanzani, G. and M. Marinacci (2017): “Ambiguity attitudes and self-confirming equilibrium in sequential games,” Working Paper.
- [3] Beauchene, D., J. Li and M. Li (2019): “Ambiguous persuasion,” *Journal of Economic Theory*, 179, Pages 312-365.
- [4] Bodoh-Creed, A. (2012): “Ambiguous beliefs and mechanism design,” *Games and Economic Behavior*, 75, 518-537.
- [5] Bose S. and A. Daripa (2009): “A dynamic mechanism and surplus extraction under ambiguity,” *Journal of Economic Theory*, 144(5), 2084-2114.
- [6] Bose S., Ozdenoren E. and A. Pape (2006): “Optimal auctions with ambiguity,” *Theoretical Economics*, 1(4), 411-438.
- [7] Bose, S. and L. Renou (2014): ”Mechanism design with ambiguous communication devices,” *Econometrica*, 82: 1853-1872.
- [8] Carrasco, V., Farinha Luz, V., Kos, N., Messner, M., Monteiro, P. and H. Moreira (2018): “Optimal selling mechanisms under moment conditions,” *Journal of Economic Theory*, 177, 245-279.
- [9] De Castro, L. and N.C. Yannelis (2018): “Uncertainty, efficiency and incentive compatibility: Ambiguity solves the conflict between efficiency and incentive compatibility,” *Journal of Economic Theory*, 177, 678-707.
- [10] Di Tillio, A., Kos, N. and M. Messner (2017): “The design of ambiguous mechanisms,” *The Review of Economic Studies*, 84, 1, Pages 237-276.
- [11] Dominiak, A., Duersch, P. and J.P. Lefort (2012): “A dynamic Ellsberg urn experiment,” *Games and Economic Behavior*, 75, 625-638.
- [12] Epstein, L. G., and Schneider, M. (2003): “Recursive multiple-priors,” *Journal of Economic Theory*, 113(1), 1-31.

- [13] Eeckhoudt, L. and C. Gollier (1995): “Demand for risky assets and the monotone probability ratio order,” *Journal of Risk and Uncertainty*, 11(2), 113-122.
- [14] Ghirardato, P. (2002): “Revisiting Savage in a conditional world,” *Economic Theory*, 20, 83-92.
- [15] Gilboa, I. and D. Schmeidler (1989): “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18(2), 141-153.
- [16] Gilboa, I. and D. Schmeidler (1993): “Updating ambiguous beliefs,” *Journal of Economic Theory*, 59(1), 33-49.
- [17] Hanany, E. and P. Klibanoff (2007): “Updating preferences with multiple priors,” *Theoretical Economics*, 2(3), 261-298.
- [18] Hanany, E., P. Klibanoff and S. Mukerji (2018): “Incomplete information games with ambiguity averse players”, Working Paper.
- [19] Kellner, C. and M. Le Quement (2017): “Modes of ambiguous communication”, *Games and Economic Behavior*, 104, 271-292.
- [20] Kellner, C. and M. Le Quement (2018): “Endogenous ambiguity in cheap talk”, *Journal of Economic Theory*, 173, 1-17.
- [21] Klibanoff P., M. Marinacci and S. Mukerji (2005): “A smooth model of decision making under ambiguity,” *Econometrica*, 73, 1849-1892.
- [22] Karni, E. (1988): “On the equivalence between descending bid auctions and first price sealed bid auctions,” *Theory and Decision*, 25, 211-217.
- [23] Karni, E. and Z. Safra (1989): “Dynamic consistency, revelations in auctions and the structure of preferences,” *Review of Economic Studies*, 56, 421-433.
- [24] Lo, K.C. (1998): “Sealed bid auctions with uncertainty averse bidders,” *Economic Theory*, 12(1), 1-20.
- [25] Lucking-Reiley, D. (1999): “Using Field Experiments to Test Equivalence between Auction Formats: Magic on the Internet,” *American Economic Review*, 89(5), 1063-1080.
- [26] Milgrom, P.R. (1981): “Good news and bad news: Representation theorems and applications,” *Bell Journal of Economics*, 12, 380-391.

- [27] Milgrom, P. and I. Segal (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70(2), 583-601.
- [28] Nakajima, D. (2011): "First-price auctions, Dutch auctions, and buy-it-now prices with Allais paradox bidders," *Theoretical Economics*, 6: 473-498.
- [29] Pahlke, M. (2018): "Dynamic consistency in incomplete information games with multiple priors", Working Paper.
- [30] Pires, C.P. (2002): "A rule for updating ambiguous beliefs," *Theory and Decision*, 53(2), 137-152.
- [31] Riedel, F. (2009): "Optimal stopping with multiple priors," *Econometrica*, 77, 857-908.
- [32] Salo, A. and M. Weber (1995): "Ambiguity aversion in first-price sealed-bid auctions," *Journal of Risk and Uncertainty*, 11(2), 123-37.
- [33] Shaked, M. and J.G. Shanthikumar (1994): "Stochastic orders and their applications," San Diego: Academic Press.
- [34] Siniscalchi, M. (2009): "Two out of three ain't bad: A comment on 'The ambiguity aversion literature: A critical assessment'," *Economics and Philosophy*, 25, 335-356.
- [35] Siniscalchi, M. (2011): "Dynamic choice under ambiguity," *Theoretical Economics*, 6, 379-421.
- [36] Strotz, R.H. (1955-56): "Myopia and inconsistency in dynamic utility maximization," *The Review of Economic Studies*, 23(3), 165-180.
- [37] Weber, R. J. (1982): "The Allais paradox, Dutch auctions, and alpha-utility theory", Working paper.
- [38] Wolitzky, A. (2016): "Mechanism design with maxmin agents: Theory and an application to bilateral trade," *Theoretical Economics*, 11, 971-1004.