

Three essays on auctions



Itzhak Rasooly

Exeter College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Trinity 2022

Acknowledgements

I could not have completed this thesis without the support of very many people. First, I am grateful to my supervisor Miguel Ballester for his constant guidance and advice. He has provided insightful comments on nearly every aspect of this thesis, and many of these comments shine through in the current exposition. I am also especially grateful to both Jasmine Theilgaard and Adam Brzezinski for spending a quite unreasonable amount of time helping me to write the three papers that form this DPhil thesis.

For useful comments and suggestions, I am grateful to many members of the Oxford economics department, including Aiden Smith, Alan Beggs, Alex Teytelboym, Elizabeth Baldwin, Inés Moreno de Barreda, Jasmin Droege, Johannes Abeler, Luke Milsom, Meg Meyer, Paul Klemperer, Peter Eso, Rossa O’Keeffe-O’Donovan, Sam Altmann, Séverine Toussaert, Simon Cowan, Tommaso Batistoni and Vince Crawford. I am also grateful to many outside Oxford, including Bernhard Kasberger, Caspar Jacobs, Charlie Plott, Dan Friedman, David Porter, David Van Dijke, Erik Kimbrough, Frieder Neunhoeffler, Heinrich Nax, Jesper Åkesson, John Ledyard, Mikhail Drugov, Nat Levine, Paul Brewer, Riccardo Camboni, Simon Finster and Steve Gjerstad. Finally, I would like to thank Darija Halatova and Tom Batistoni for helping me acquire the (admittedly limited) set of coding skills I now possess.

Contents

Introduction	4
1 The importance of being discrete: on the (in-)accuracy of continuous approximations in auction theory	7
1 Introduction	8
2 The model	11
3 The second-price auction	12
4 The first-price auction	12
5 The all-pay auction	16
6 Asymmetric equilibria	18
7 Mixed equilibria	20
8 Concluding remarks	22
9 Appendices	29
2 Going... going... wrong: a test of the level-k (and cognitive hierarchy) models of bidding behaviour	49
1 Introduction	50
2 Separating level- k from equilibrium	53
2.1 The all-pay auction	53
2.1 The first-price auction	58
3 Experimental design	63
4 Results	65
4.1 Pooling the data	65
4.2 Structural estimates	68
4.3 Correlating levels	70

4.3	Subject reports	71
5	Robustness checks	72
6	Concluding remarks	77
7	Appendices	82
3	Competitive equilibrium and the double auction	99
1	Introduction	100
2	CE preserving shifts	103
3	Experimental design	107
4	Results	109
5	Understanding the monotonicity	113
6	The Marshallian path	114
7	Concluding remarks	123
8	Appendices	129

Introduction

In this introduction, I begin with a brief summary of the three chapters that jointly form this thesis. Since each chapter is already preceded by an abstract summarising the main ideas which it contains, my summaries here will be rather informal. After providing the summaries, I will then take the opportunity to discuss some connections (and even apparent tensions!) between the chapters.

Chapter 1. My first chapter begins by noting that, although auction theory traditionally treats valuations and bids as continuous variables, in reality both are discrete. This then raises the question of whether anything interesting changes once we incorporate discreteness into standard auction models. The main finding is that introducing arbitrarily fine discreteness can rob auctions of pure strategy equilibria (symmetric or otherwise). As a result, the only equilibria are in mixed strategies, meaning that auction participants are supposed to randomise between different bids (despite having no strict incentive to do so). Whether one will find this interesting obviously depends on whether one finds the idea of this kind of randomisation plausible or not. For my part, however, I find it rather implausible; and indeed I make various arguments as to why such randomisation is unlikely to arise. As a result, my findings provide one with (another) reason to be sceptical of equilibrium-based auction models.

Chapter 2. My second chapter builds on the first, but has a rather different focus. Instead of being concerned with the accuracy of equilibrium, it is principally concerned with the accuracy of the level- k model of auctions introduced by Crawford and Iriberri (2007) (although a comparative exercise in which this is compared to equilibrium is also conducted). One central point of the paper is that the level- k model makes extremely unrealistic (but previously unnoticed) predictions when applied to a variety of auction structures. Unsurprisingly, these predictions are strongly falsified by the data, suggesting that the level- k model is not suited to explain auction behaviour (despite its noted successes in capturing aspects of strategic behaviour in various other settings).

Chapter 3. While again concerned with auctions, the third chapter has a rather different focus than the first two. The point of departure is a very large literature documenting rapid convergence to competitive equilibrium in double auction settings. In the paper, I revisit this conclusion by using a class of ‘equilibrium preserving shifts’ to provide a very challenging test of competitive equilibrium theory. I make two central claims. First, I argue that if traders are not permitted to leave the market without replacement as time progresses,

then competitive equilibrium actually provides a rather poor description of double auctions. Second, however, I also argue that double auctions have strong equilibrating tendencies once traders are allowed to drop out (without replacement). Taken together, this suggests that a dropping out dynamic (which I formalise) is crucial for explaining equilibration in double auctions (when it occurs), something which helps solve a ‘scientific mystery’ introduced by Smith (1962) over 60 years ago.

Connections. Although each chapter is self-contained, some interesting connections appear between them once they are presented as a single work. First, chapters 1 and 2 both share a preoccupation with how economic models change once crucial variables are discretised. Some of the results in chapter 2 rely on results established in chapter 1; and one could conceivably use some of the ideas developed in chapter 2 to revisit some of the questions in chapter 1 (although I have not attempted to do this).¹ Despite this interesting overlap, however, the focus of the two papers is rather different. As discussed above, chapter 1 is concerned with the question of whether the equilibria of discretised auctions are in pure strategies. Chapter 2, however, is concerned with the more practical question of how one can calculate the equilibria of discrete auctions (whether they are in pure or mixed strategies) for the purposes of testing equilibrium predictions against empirical observation.

In addition to this connection, we should also note two apparent tensions between the papers. In chapter 1, it is suggested that introducing discreteness makes an interesting qualitative difference; and that it can also change the quantitative character of equilibrium predictions. In chapter 2, however, it is suggested that the equilibrium predictions for the relevant auction experiments are fairly well approximated by the continuous theory. So what explains the apparent tension — and are the continuous approximations accurate or not? The answer, of course, is that it depends. One can certainly find cases where introducing discreteness makes a large quantitative difference to the theory’s predictions (see Figure 2 in chapter 1 for an illustration). However, one can also find cases where it does not make a large quantitative difference, and the experiments in chapter 2 turn out to be such cases (although they do still involve randomisation, thereby illustrating the central point of chapter 1).

Another apparent contrast between the chapters lies in their implicit attitudes towards Bayes Nash equilibrium. In chapter 1, the overall view is fairly negative; and it is suggested that the lack of pure equilibria in discretised auctions provides us with a reason to be sceptical of the

¹More concretely, one can address the question of whether pure equilibria exist (the focus of chapter 1) by computing the equilibrium using the algorithm in chapter 2 and checking if all ‘jump points’ are integers. For example, in an all-pay auction with 2 bidders and uniformly distributed values, this involves repeatedly searching for the smallest real number x that satisfies inequalities of the form $x[x] \geq S$ (where $S > 0$ is a parameter) and checking if one only obtains integer solutions.

theory. In chapter 2, however, it is suggested that equilibrium fares substantially better than a rival behavioural theory, namely the level- k model. Of course, these claims are entirely consistent: it is perfectly possible that Bayes-Nash equilibrium is both an unsuccessful theory (when applied for auctions) but also that alternative behavioural theories (e.g., level- k) are even worse.

A final tension we should note is that despite all that we say in chapter 1, chapter 3 uses a continuum model both to describe the idea of equilibrium preserving shifts as well as to formalise the Marshallian path notion that we claim is so crucial for generating equilibration. This is of course rather unrealistic: experimental double auctions do not literally have an infinite number of participants. As it turns out, however, exactly analogous results are available in the finite case – so the use of the continuum model seems rather harmless. This is quite a contrast to the setting of chapter 1, where employing continuous approximations can make a real difference (or so I argue).

I conclude with a note on terminology. Despite the fact that two of the three chapters were single authored, I use the pronoun ‘we’ throughout. As remarked by Krishna (2009) in a related context, this is not to indicate any royal lineage; and nor is it to indicate a strange degree of reliance on my academic supervisor. It is instead a stylistic approach that I preferred at the beginning of my (our?) DPhil thesis; and one which despite some current misgivings, I have maintained for purposes of stylistic consistency.

References

- Crawford, V. P. and Iriberri, N. (2007). Level- k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770.
- Krishna, V. (2009). Auction theory. *Elsevier Monographs*.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–137.

The importance of being discrete: on the (in-)accuracy of continuous approximations in auction theory*

Abstract

While auction theory views bids and valuations as continuous variables, real-world auctions are necessarily discrete. In this paper, we use a combination of analytical and computational methods to investigate whether incorporating discreteness substantially changes the predictions of the theory. In some cases, we find that introducing discreteness leaves auction theoretic predictions almost entirely unaffected. In other cases, however, we show that introducing discreteness can rob auctions of their usual symmetric, pure-strategy equilibria; and can even rob them of pure-strategy equilibria altogether (symmetric or otherwise). We discuss the implications of our findings for both theoretical and experimental work on auctions.

KEYWORDS: auction theory, equilibrium, discrete, experiments

JEL CODES: C70, C72, C91, D44

*This chapter was co-authored with Carlos Gavidia-Calderon.

1 Introduction

In auction theory, bids and valuations are typically viewed as continuous variables. That is to say, the theory assumes that players may bid any non-negative real number; and that valuations are continuously distributed over some subset of the real line. In actual auctions, however, both valuations and bids are discrete. For example, a participant in an auction would never be able to submit a bid of $\sqrt{2}$ pounds; and must inevitably choose from some finite set of possible bids. Similarly, since individual valuations are defined as the maximum that a player is willing to pay — and therefore measured in currency units — it follows that $\sqrt{2}$ pounds could never be a player’s valuation. This raises a simple question: does introducing discreteness into standard game theoretic models of auctions substantially change their predictions?

Surprisingly, this question has received very little attention from economists. In the theoretical literature, there are some papers that examine the case of discrete valuations but continuous bids; and some others that examine the case of discrete bids but continuous valuations.¹ However, there are far fewer that consider the case of discrete bids and discrete valuations. Moreover, those papers that do are interested in questions that are very different from ours. For instance, Dekel and Wolinsky (2003) introduce discrete values and bids in order to study the implications of rationalisability (whereas we focus on equilibrium); and in any case their results rest on the assumption that the number of bidders is very large.²

The question is also almost completely ignored in the experimental literature. In this literature, it is standard for experimenters to begin by writing down a continuous auction model and to then ‘test’ the predictions of this model using a discrete laboratory experiment. It is occasionally asserted, without any real justification, that the discreteness makes little difference if the number of possible valuations is large (see, e.g., Noussair and Silver (2006), footnote 3). More often, however, the fact that discrete games are being used to test a continuous theory is not even discussed (see, e.g., Coppinger et al. (1980); Cox et al. (1982); Kagel and Levin (1985); Dyer et al. (1989); Chen and Plott (1998); Filiz-Ozbay and Ozbay (2007); Kirchkamp and Reiß (2011)).³ This is potentially concerning: if discrete equilib-

¹For discrete valuations and continuous bids, see Maskin and Riley (1985), Riley (1989), Wang (1991), Che and Gale (2006) and Bergemann et al. (2017). For discrete bids and continuous valuations, see Chwe (1989), Rothkopf and Harstad (1994), David et al. (2007), Cai et al. (2010) and Gonçalves and Ray (2017).

²There are also some papers which analyse the equilibria of auctions with discrete bids and known valuations (Boudreau, 2011; Li, 2017). Our interest, however, lies in the case where valuations are uncertain.

³An important exception is a paper by Goeree et al. (2002) that carefully analyses the equilibria of their experimental auction game. While their analysis is insightful, the case they consider (2 bidders, 6 possible valuations, 1 auction structure) is obviously a rather special one and could not (and is not intended to) shed very much light on the general case.

ria are importantly different from continuous equilibria, then these experimental tests are invalid.

In this paper, we make substantial progress in answering this question through study of a discretised version of the canonical independent private values model. We assume that values and bids are restricted to the same finite and evenly spaced set (a fairly common set-up in the experimental literature, although not one that may be universally applied). We allow the discretisation to be arbitrarily fine (so that the difference between the bids can be as small as a penny). We then study the equilibria of three standard auction structures: the first-price, second-price and all-pay auctions. We also consider two different tie-breaking rules: fair tie-breaking (in which tying winners are chosen randomly) and the case in which players cannot win through a tie. While the former assumption is more popular in auction experiments, the latter is easier to study analytically. Ultimately, though, our results do not substantially hinge on which tie-breaking rule we assume.

We begin the analysis by considering the second-price auction. The central insight from classical auction theory – that bidders should submit bids equal to their valuation – applies in our discrete setting. Thus, introducing discreteness makes essentially no difference in the second-price auction, a fact which confirms the intuition of the many experimenters who have attempted to test the continuous theory using discrete laboratory auctions (Coppinger et al. (1980); Cox et al. (1982); Kagel and Levin (1993); Aseff (2004); Andreoni et al. (2007); Cooper and Fang (2008); Drichoutis et al. (2015); Georganas et al. (2017)).

We then turn to the first-price auction and begin by proving some useful lemmas that hold regardless of the distribution that generates each player’s valuation. First, we show that bids are weakly increasing in values. Second, we show that symmetric equilibria cannot have jumps (for example, a situation in which bids jump from 3 to 5 when the valuation increases from 10 to 11). Finally, we show that if a discrete auction has a symmetric pure-strategy equilibrium, then it is (essentially) unique.

Using these results, we then study the equilibria of the first-price auction, focusing on the case of uniformly distributed valuations so that our results speak to the majority of auction experiments (see Kagel (1995) for a review of such experiments). In some cases (for example, with two bidders and an even number of values), we find that the auction can possess an equilibrium which is almost identical to the equilibrium in the continuous case and converges to the continuous equilibrium as the discreteness vanishes. In other cases, however, we uncover discontinuities. For example, with three or more bidders and a suitable tie breaking rule, the auction has no symmetric, pure-strategy equilibria – in sharp contrast

to the continuous case.

We then conduct a similar analysis of the all-pay auction. In this case, we once again find dramatic contrasts between our discrete games and their continuous counterparts (which are characterised by symmetric pure-strategy equilibria). Indeed, there are now no symmetric, pure-strategy equilibria at all – a result that holds regardless of the number of bidders, tie-breaking rule and the fineness of the discretisation. Like our results for the first-price auction, this reveals a discontinuity in the set of equilibria (as the discretisation goes to 0) and contradicts assertions by experimenters about the equilibria of the games they study.

In principle, discrete auctions could have asymmetric equilibria that closely resemble the symmetric equilibria of their continuous counterparts. To investigate this possibility, we take a computational approach, writing down our games in normal form and computing their equilibria using standard solvers. Doing so is computationally challenging since the number of possible strategies quickly becomes very large (e.g. in one game whose equilibria we ultimately compute, each player has ten billion possible bidding functions to choose from). Nonetheless, we are able to significantly reduce the dimensionality of our problem through iterative deletion of dominated bids and thus compute the equilibria of a series of representative auction games.

In some cases (e.g. the first-price auction with 3 bidders), we do find asymmetric equilibria that closely resemble the symmetric equilibria of continuous theory. In general, however, the message of our computational results reinforces our analytical findings about symmetric equilibria. For almost all types of auction game that we consider, we find examples in which the auction has no pure-strategy equilibria (symmetric or otherwise). This further reinforces the contrast between the discrete and the continuous case.

In cases where discrete auctions can lack pure strategy equilibria, it follows from Nash (1950) that their equilibria must be in mixed strategies. In our final section, we study such equilibria, paying particular attention to the question of whether the randomisation is ‘local’ in the sense of occurring between at most two neighbouring bids. We show that this is emphatically *not* the case: indeed, it is possible to construct discrete auctions in which one type randomises over nearly the entire range of feasible bids. Thus, not only can introducing discreteness make an interesting qualitative difference (by robbing auctions of pure strategy equilibria), but it can lead to equilibria which are quantitatively rather dissimilar from the pure strategy equilibria of textbook theory.

The remainder of this article is structured as follows. Section 2 presents the (family of) models under consideration. Sections 3 – 6 present the results for our three auction structures.

Section 7 studies mixed equilibria and shows that they can involve randomisation over an arbitrarily large range. Finally, Section 8 concludes by discussing the implications of our findings for experimental and theoretical research.

2 The model

We consider $n \geq 2$ bidders who compete for a single and indivisible object. Bidder i 's value for this object v_i is drawn from the finite set $X = \{0, \delta, 2\delta, \dots, x\delta\}$ for some $\delta > 0$ and $x \in \mathbb{N}$. Without loss of generality, assume that every element in X is drawn with strictly positive probability; and (with some redundancy) let $S \equiv x + 1$ denote the number of elements in X . As usual, bidders know their own valuation but only the joint distribution of the values of the other bidders. Valuations are private and independently drawn. In addition, we assume that player i 's bid b_i must belong to X (a common situation in the experimental literature, and one which plausibly applies to many non-experimental auctions).⁴

In all auction structures that we consider, a player who submits a bid that is strictly higher than all opponent bids wins the auction. We consider two different rules that come into play in the event that two or more bids tie for highest. In the *model without ties*, players win the auction if and only if their bid is strictly higher than all the bids submitted by their opponents. Thus, $P(i \text{ wins}) = P(b_i > b_j \ \forall j \neq i)$. In the *model with ties*, a player can win if they tie with others who also submit the highest bid. Specifically, if m players submit the highest bid, they all have an equal chance ($1/m$) of being selected as the winner. This means, for example, that $P(i \text{ wins}) = P(b_i > b_j) + 0.5P(b_i = b_j)$ in the case of $n = 2$.

Of course, bidder payoffs do depend on the auction structure. In the first-price auction, the player who submits the highest bid wins the object and pays a price equal to their bid. Thus, the expected payoff of a bidder with a value v_i and submits a bid of b_i is $\pi(v_i, b_i) = (v_i - b_i)P(i \text{ wins})$. In the all-pay auction, players pay their bid even if they lose. Hence, $\pi(v_i, b_i) = v_iP(i \text{ wins}) - b_i$. In the second-price auction, the player who submits the highest bid wins the object but pays a price equal to second highest bid. We assume throughout that bidders are risk-neutral and therefore maximise their expected payoff.

Bidder's i 's (ex-ante) strategy is a function $\beta_i: X \rightarrow X$ mapping from each valuation they might have to a (permissible) bid. We define a (pure) *equilibrium* as a set of bidding functions, one for each player, such that each player's bidding function maximises their expected payoff

⁴As in the continuous model, analogous results are available if one defines the value/bid space as $X = \{c, c + \delta, c + 2\delta, \dots, c + x\delta\}$ for some constant $c > 0$. Thus, the decision to set the minimum value/bid at zero should be viewed simply as a normalisation.

given the bidding functions of the other players.⁵ (This is equivalent to a pure-strategy Bayes-Nash equilibrium with risk-neutral bidders.) When all bidders choose the same bidding function, we say that the equilibrium is *symmetric*. We focus on symmetric, pure-strategy equilibria since these equilibria characterise the continuous case – and the point of the exercise is to see how discrete equilibria compare to this benchmark.

Finally, we restrict attention to equilibria in undominated strategies (i.e. equilibria in which players never choose strategies which are weakly dominated). While this might seem restrictive, this cannot rule out any of the equilibria considered in the standard analysis of continuous auctions (since these are in undominated strategies). In addition, the restriction is easy to motivate: for example, players will never choose dominated strategies if their opponents ‘tremble’ with arbitrarily low probability (Selten, 1975). Since we only consider equilibria in undominated strategies, we omit the phrase ‘in undominated strategies’ when stating our results.

3 The second-price auction

To begin with, we consider the second-price auction. Here the analysis is fairly trivial, as the following proposition indicates.⁶

Proposition 1 *In the model with ties, each player sets $\beta(v) = v$ in the unique equilibrium. In the model without ties, the set of equilibria is precisely the set of strategy profiles in which each player sets either $\beta(v) = v$ or $\beta(v) = v + \delta$.*

The point here is that in any equilibrium (in undominated strategies), players bid their valuation (or perhaps their valuation plus δ). Of course, players also bid their valuation in the unique equilibrium in undominated strategies of the continuous model. Thus, introducing discreteness makes essentially no difference in the second-price auction, a fact which confirms the intuition of the many experimenters who have attempted to test the continuous theory using a discrete laboratory auction. Since the analysis of the second-price auction is so straightforward, we quickly move on to our next auction structure.

4 The first-price auction

We begin our analysis of the first-price auction with some useful lemmas that hold independently of the distribution that generates the players’ valuations.

⁵We defer all discussion of mixed equilibria to Section 7.

⁶All proofs are collected in the appendix.

Lemma 1 *The set of equilibria does not depend on δ .*

Since our interest is in equilibria, Lemma 1 allows us normalise $\delta = 1$ without loss of generality. In other words, we may view the bids and values as integers and strategies as a map from some set of consecutive integers to itself. This is not very surprising: obviously there is no important distinction between the game in which players bid 2δ when their value is 3δ , and the game in which they bid 2 when their value is 3. Holding the number of valuations and bids fixed, the only effect of varying δ is to induce a (positive) linear transformation in the payoffs. Throughout the following, then, we will set $\delta = 1$.

Lemma 2 (No Dominated Strategies.) *Suppose that a bidding function β is chosen in an equilibrium. In the model with ties, $\beta(0) = 0$ and $\beta(v) \leq v - 1$ for all $v \geq 1$. In the model without ties, $\beta(0) = 0$, $\beta(1) \leq 1$ and $\beta(v) \leq v - 1$ for all $v \geq 2$.*

Roughly speaking, Lemma 2 rules out bids that exceed valuations: such bids are weakly dominated. As a result, it pins down bidding behaviour at the ‘start’ of the bidding function (e.g., when $v = 0$) in a similar way to the boundary condition of continuous theory. As it will turn out, these low bids then pin down bids for higher valuations by inductive arguments. Thus, the assumption that players avoid dominated strategies will prove crucial for the analysis.

Lemma 3 (Monotonicity.) *Suppose that a bidding function β is chosen in an equilibrium. Then for any two values $v, v' \in X$, $v' > v$ implies that $\beta(v') \geq \beta(v)$.*

Lemma 3 states that equilibrium bids must be (weakly) increasing in valuations. While this must hold in equilibrium, it is perhaps worth noting that it holds more generally. To be more precise, suppose that the player believes that there is a strictly positive probability that they will win by submitting a bid of $\beta(v')$. Then, regardless of whether they expect their opponents to play equilibrium strategies, their optimal bid cannot decrease when their valuation increases from v to v' . In other words, the monotonicity property arises more from the logic of best responses than from equilibrium reasoning.

We now examine some of the properties of *symmetric* equilibria (henceforth, SE), which are characterised by a single bidding function β . To begin with, note that, in the model without ties, an SE cannot have jumps (i.e. a situation in which the bid jumps from b to some $b' \geq b + 2$ when the valuation increases from v to $v + 1$). To see this, suppose (for contradiction) that there were such an SE and consider what would happen if a player with a value of $v + 1$ were to deviate by reducing their bid to $b + 1$. The deviation would not reduce their probability of winning: both before the deviation and after it, they win the auction if and only if the highest bid of their opponents is less than or equal to b . Obviously,

however, it would increase their payoff in the event that they did win. Finally, since they had a strictly positive probability of winning in an SE (e.g. whenever their opponents all bid b), the deviation would strictly increase their expected payoff, contradicting the supposed optimality of their strategy in the candidate SE. We thus have the following result.

Lemma 4 (No jumps) *Suppose that β is an SE in the model without ties. Then for every $v, v + 1 \in X$, $\beta(v + 1) \leq \beta(v) + 1$.*

It is not as straightforward to prove a corresponding lemma in a model with ties. The reason is that the reduction in a player's bid from b' to $b + 1$ now does reduce the player's probability of winning: for instance, it eliminates their chance of winning when their opponents all bid b' . Nonetheless, we are able to prove qualified versions of this lemma for particular distributions of player valuations (see Lemma 13 for the details).

Finally, we prove a uniqueness result.

Lemma 5 *In the model without ties, the SE (if it exists) is unique up to a choice of $\beta(1)$.*

Lemma 5 states that, in the model without ties, the SE (if it exists) is uniquely implied by the choice of $\beta(1)$. Since players avoid dominated strategies, either $\beta(1) = 0$ or $\beta(1) = 1$. Hence, there are at most two SE in the first-price auction. In some cases, the choice of $\beta(1)$ makes no difference to higher bids, so the SE is essentially unique. In others, however, changing $\beta(1)$ has implications for bids further up the chain – a fact that will be illustrated by the next proposition.

Using these results, we now examine the equilibria of the first-price auction. To do so, we will now assume that the values are uniformly distributed. The reasons are twofold. First, this distribution is by far most the popular choice by auction experimenters (Kagel, 1995; Kagel and Levin, 2011); so its analysis is of particular importance if we wish to interpret their results. Second, it turns out that it is impossible to say very much about the existence of SE for arbitrary distributions. Thus, some distributional choice is necessary.

We begin with the case of 2 bidders in the model without ties.

Proposition 2 *In the first-price auction with two bidders and without ties, there are exactly two SE: $\beta(v) = \lfloor v/2 \rfloor$ and $\beta(v) = \lceil v/2 \rceil$.*

Proposition 2 demonstrates that continuous auction theory can be an extremely good approximation to discrete reality (see Chwe (1989) for a similar result in the context of discrete bids but continuous valuations). In the continuous case, players set $\beta(v) = v/2$. In the discrete case, they also bid half their value when possible and they round their bid to the

one of the two closest integers otherwise. (In one SE, they round up; in the other, they round down.) Thus, the difference between continuous and discrete case can be extremely small (e.g. if the bids and values are discretised in pennies) and the difference vanishes as the discretisation goes to 0. That is,

Corollary 6 *In any SE with two bidders and without ties, bidding functions and expected revenues converge to their counterparts in the continuous auction as $\delta \rightarrow 0$.*

Unfortunately, however, this happy state of affairs ceases to be once we introduce additional bidders (while maintaining our previous assumption about ties). In fact, we have the following rather striking result:

Proposition 3 *Suppose that $n \geq 3$ and that the maximum x satisfies*

$$x > \frac{2^{\frac{1}{n-1}}}{2^{\frac{1}{n-1}} - 1}. \quad (1)$$

Then there are no SE in the first-price auction without ties.

To illustrate the argument, we sketch a proof for the case of $n = 3$. To begin, we study the start of the bidding function. By Lemma 2, $\beta(0) = 0$, $\beta(1) \leq 1$, and $\beta(2) = 1$. Furthermore, it is easy to check that $\beta(3) = 2$ in any SE (if this were not true, i.e. $\beta(3) = 1$, it would then pay for a bidder with a value of 3 to deviate to a bid of 2.)

Of particular interest is the bid when $v = 4$. By monotonicity and no jumps, there are only two possibilities: $\beta(4) = 2$ and $\beta(4) = 3$. Let us consider these possibilities in turn. In the first instance, suppose that $\beta(4) = 2$. In that case, a player with a value of 4 would win in equilibrium whenever both of their opponents' values were 2 or lower. Since there are 3 such values, their equilibrium (expected) pay-off would be $2(3/S)^2$ where S is the number of possible valuations. However, if the player were to deviate by bidding 3, they would win (at least) whenever their opponents' values were 4 or lower. Since there are 5 such values, their payoff would be at least $(5/S)^2 > 2(3/S)^2$. Since the deviation would be strictly profitable, we cannot have $\beta(4) = 2$ in an SE.

Let us then consider the second possibility: $\beta(4) = 3$. In that case, a player with a value of 4 would win in equilibrium whenever their opponents' values were both 3 or lower. Since there are 4 such values, their equilibrium pay-off would be $(4/S)^2$. However, if the player were to deviate by bidding 2, they would win whenever their opponents' values were both 2 or lower and therefore get an expected pay-off of $2(3/S)^2 > (4/S)^2$. Thus, we also cannot have $\beta(4) = 3$ in an SE. Since these were the only 2 possibilities, this means that there is no

SE in the case of $n = 3$.

We now briefly remark on the assumption about the number of possible values. As noted, the argument assumes that inequality (1) holds; i.e., that the number of valuations is large relative to the number of bidders. In our view, this is a reasonable requirement that is satisfied by virtually all real world and experimental auctions. If there are 1000 possible values (a fairly common strategy space in experiments), all this requires is that the number of bidders is fewer than 693. Even if the number of possible values is just 100, all this requires is that the number of bidders is fewer than 70. As an aside, we note that there *is* an SE in the very high bidder case in which each player bids 0 when $v = 0$ and $\beta(v) = v - 1$ otherwise. This last observation, incidentally, is consistent with the finding of Dekel and Wolinsky (2003) that an extremely large number of bidders can induce players to bid almost entirely their valuation.

It might be wondered whether this non-existence result survives as soon as ties are introduced. In fact, non-existence is common under either tie breaking rule (see the appendix for our full analysis of this case). Indeed, in some cases, allowing for ties can make things even worse. For example, in the case of 2 bidders (which had an SE in the no-ties case), SE now fail to exist if the number of possible values is odd.

In general, these results contradict assertions found in a number of experimental papers. For example, Cox et al. (1982) experimentally simulate a first-price auction with two bidders, an odd number of values and fair tie-breaking; and assert (based on a lengthy discussion of the continuous case) that has a symmetric, pure strategy equilibrium (see their equation 3.5). As it turns out, however, their games have no pure SE at all — and therefore cannot have the particular pure SE that they identify.

5 The all-pay auction

We now turn to the all-pay auction. As before, we begin with some useful lemmas.

Lemma 7 *The set of equilibria is invariant to δ .*

Lemma 8 (No Dominated Strategies.) *Suppose that a bidding function β is chosen in an equilibrium. Then $\beta(0) = 0$ and $\beta(v) \leq v - 1$ for all $v \geq 1$.*

Lemma 9 (Monotonicity.) *Suppose that a bidding function β is chosen in an equilibrium. Then for any two values $v, v' \in X$, $v' > v$ implies that $\beta(v') \geq \beta(v)$.*

Lemma 10 (No jumps) *Suppose that β is an SE in the model without ties. Then for every*

$v, v + 1 \in X$, $\beta(v + 1) \leq \beta(v) + 1$.

Lemma 11 *In the model without ties, the SE (if it exists) is unique.*

In general, the logic behind these results is very similar to that in the first-price auction. We pause, however, to note one minor contrast. In the all-pay auction, restricting attention to undominated strategies implies that $\beta(0) = \beta(1) = 0$. Thus there is no longer a choice to be made about the bid when $v = 1$. When we combine this with the inductive argument used to establish Lemma 5, we see that the all-pay auction has at most one SE (in undominated strategies) while the first-price auction had at most two.

These preliminaries out of the way, we examine the existence of symmetric equilibria, again assuming uniformly distributed valuations so that our results speak to the majority of experimental auctions. In this case, we find an even more dramatic contrast between discrete auctions and their continuous counterparts:

Proposition 4 *Suppose that $x \geq 10$. Then, regardless of the number of bidders and tie-breaking rule, there are no SE in the all-pay auction.*

To (very partially) illustrate the logic of the argument, consider the case of the model without ties and let $n = 2$. Where previously we studied the start of the bidding function, now we study the end. By Lemmas 9 and 10, there are only a limited number of ways in which the bidding function can end. For example, it might be flat over the last three values or it might twice increase. It turns out, however, that no matter how it ends, no bidding function is a best response to itself. For example, if one bidder increased their bid twice, then the other should reduce their bid by 1 when their value is the maximum.

As before, Proposition 4 contradicts assertions by experimenters that discrete all-pay auctions are characterised by pure-strategy SE. For example, Noussair and Silver (2006) conduct an experimental all-pay auction with $n = 6$, 1000 possible bids/values and a fair tie breaking rule. While recognising that their game is discrete, they claim that it possesses the pure-strategy SE of continuous theory since the number of possible bids/values is large. In fact, we see that the situation is the opposite: it is precisely because the number of values is large that a pure-strategy SE does not exist.⁷ Moreover, Proposition 4 corrects their claim (see footnote 3) that discrete auction games differ from their continuous counterparts because in the discrete case there is a strictly positive probability of winning through a tie. In fact, our difference arises even in the model without ties (when winning through a tie is impossible).

⁷In cases with a handful of possible bids and valuations, it is easy to verify that SE do exist. For instance, if $x = 2$, then (regardless of the number of bidders) there is an SE in which $\beta(0) = 0$, $\beta(1) = 0$ and $\beta(2) = 1$.

6 Asymmetric equilibria

In principle, discrete auctions could have asymmetric equilibria that closely resemble the symmetric equilibria of their continuous counterparts. To investigate this possibility, we take a computational approach, first representing the games in normal form and then finding their equilibria using standard solvers (McKelvey et al., 2014). Since Bayesian games can be represented in sequence form, this approach is not computationally efficient (von Stengel, 1996). However, the normal form representation allows us to apply standard algorithms which are guaranteed to find all pure-strategy equilibria and therefore establish the non-existence of pure-strategy equilibria in cases where the algorithmic search is unsuccessful.

Computing the equilibria of discrete auctions is challenging since the number of possible strategies is so large. Indeed, in a game with k possible bids and valuations, each player has k^k possible pure strategies (so the number of strategies grows super-exponentially). Nonetheless, we are able to significantly reduce the dimensionality of our problem through iterated deletion of dominated strategies. In the first stage, we delete strategies that involve bids above valuations since these strategies are weakly dominated. In subsequent stages, we iteratively delete strategies that are now strictly dominated.⁸ In other words, we delete strategies using the Dekel-Fudenberg procedure (Dekel and Fudenberg, 1990). While not designed for this purpose, this procedure is guaranteed to preserve all equilibria in undominated strategies.⁹ Furthermore, it reduces the number of strategies quite dramatically and therefore allows us to construct payoff matrices with hundreds of thousands of entries where previously the number of entries would have run into the billions.

We then compute the equilibria of a variety of auction games. We vary the auction structure (first-price vs. all-pay), number of bidders (two vs. three), tie breaking rule and number of possible valuations. Throughout, we assume uniformly distributed values. In general, we avoid searching for equilibria in cases where we have already identified symmetric equilibria analytically (see Sections 4 – 5).¹⁰

In some cases, we do find asymmetric equilibria that closely resemble the symmetric equilibria of continuous theory. For example, consider a first-price auction with three bidders and no-ties; and suppose that the number of valuations is *not* a multiple of three. Then (generalising one of our computational results) it is easy to check that there is an asymmetric

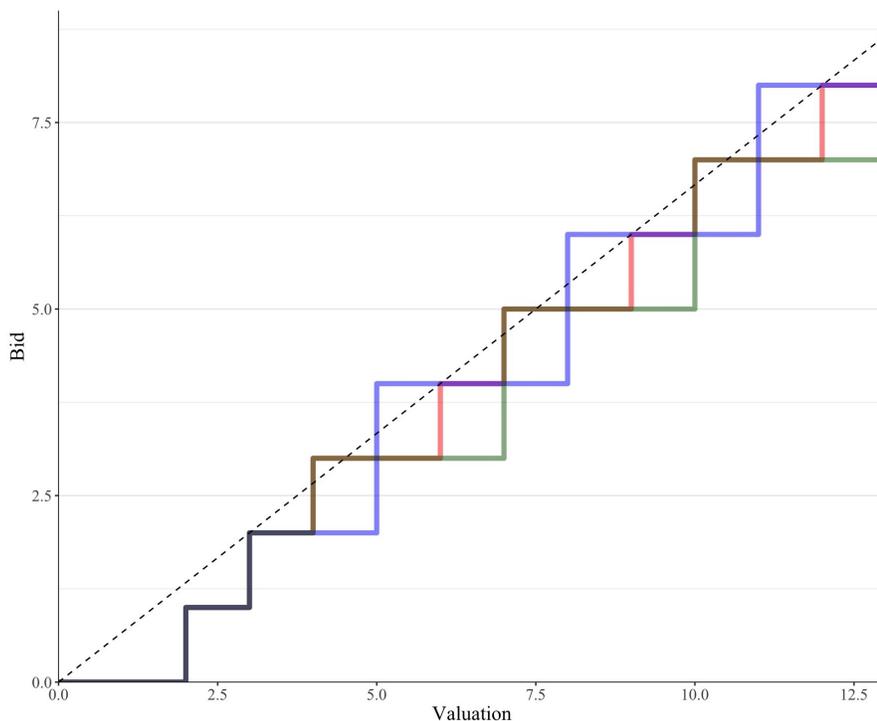
⁸For the details, see: <https://github.com/cptanalatriste/discrete-auction-builder/tree/master/docs>

⁹Clearly, the first round must preserve all equilibria in undominated strategies. Moreover, the subsequent rounds (in which strictly dominated strategies are deleted) are guaranteed to preserve all equilibria (in undominated strategies or otherwise).

¹⁰Our code is available here: <https://github.com/cptanalatriste/bayesian-game-builder>

equilibrium in which players submit bids that are very close to the continuous equilibrium of $\beta(v) = 2v/3$ (see Figure 1 for an illustration).¹¹ In this case, the discrete auction lacks symmetric equilibria (by Proposition 3), but allowing for asymmetry provides a way of recovering continuous results.

Figure 1: An asymmetric equilibrium with $n = 3$



Notes: Figure 1 plots the asymmetric equilibrium discussed in the text for the case of $x = 13$. The dotted line depicts the continuous equilibrium of $\beta(v) = 2v/3$.

In general, however, the analysis further emphasises the discontinuities highlighted in the previous section. Indeed, in six out of the seven types of auction game that we consider, we find examples in which the auction has no strategy equilibria whatsoever (see appendix Tables 1 and 2 for the full results). While we were only able to solve relatively small auctions, we conjecture (based on the results of Sections 4 – 5) that non-existence would also be common for auctions with larger numbers of values.

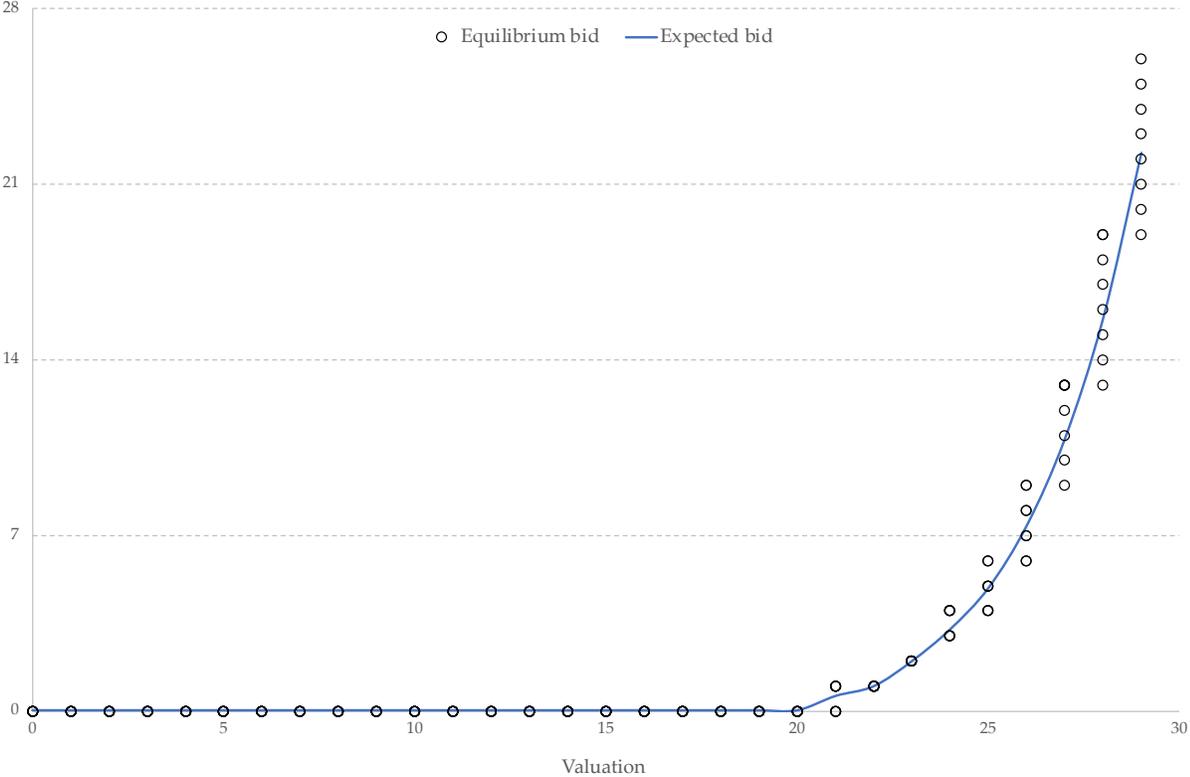
¹¹Let m denote an arbitrary multiple of 3. In this equilibrium, the first bidder sets $\beta(0) = \beta(1) = 0$ and then $\beta(m) = 2m/3$, $\beta(m-1) = \beta(m) - 1$ and $\beta(m+1) = \beta(m) + 1$ (for all $m \geq 3$). The second bidder sets $\beta(0) = \beta(1) = 0$, $\beta(2) = 1$, $\beta(3) = \beta(4) = 2$; and then $\beta(m-1) = \beta(m) = \beta(m+1) = 2m/3$. The third bidder sets $\beta(0) = \beta(1) = 0$, $\beta(2) = 1$, $\beta(3) = 2$; and then sets $\beta(m) = \beta(m-1) = \beta(m-2) = 2m/3 - 1$.

7 Mixed equilibria

The previous analysis demonstrates that discrete auctions can lack pure strategy equilibria, symmetric or otherwise. In such cases, it follows from Nash (1950) that their only equilibria are in mixed strategies. It is natural to conjecture that, although these mixed equilibria differ qualitatively from the pure equilibria of textbook theory (insofar as they involve randomisation), they might match textbook theory quantitatively since the randomisation might be quite local. For example, if the textbook (continuous) theory predicts that a player should submit a bid of 6.5, it is reasonable to guess that the discrete model with integer bids will predict randomisation between bids of 6 and 7. In this section, we show that this is emphatically *not* the case. In fact, as the next result indicates, randomisation can take place over almost the entire feasible range of bids.

Proposition 5 *Consider the all-pay auction without ties. If the number of bidders is sufficiently large, players with a value of x must submit all bids between 0 and $x - 1$ with positive probability in any symmetric equilibrium.*

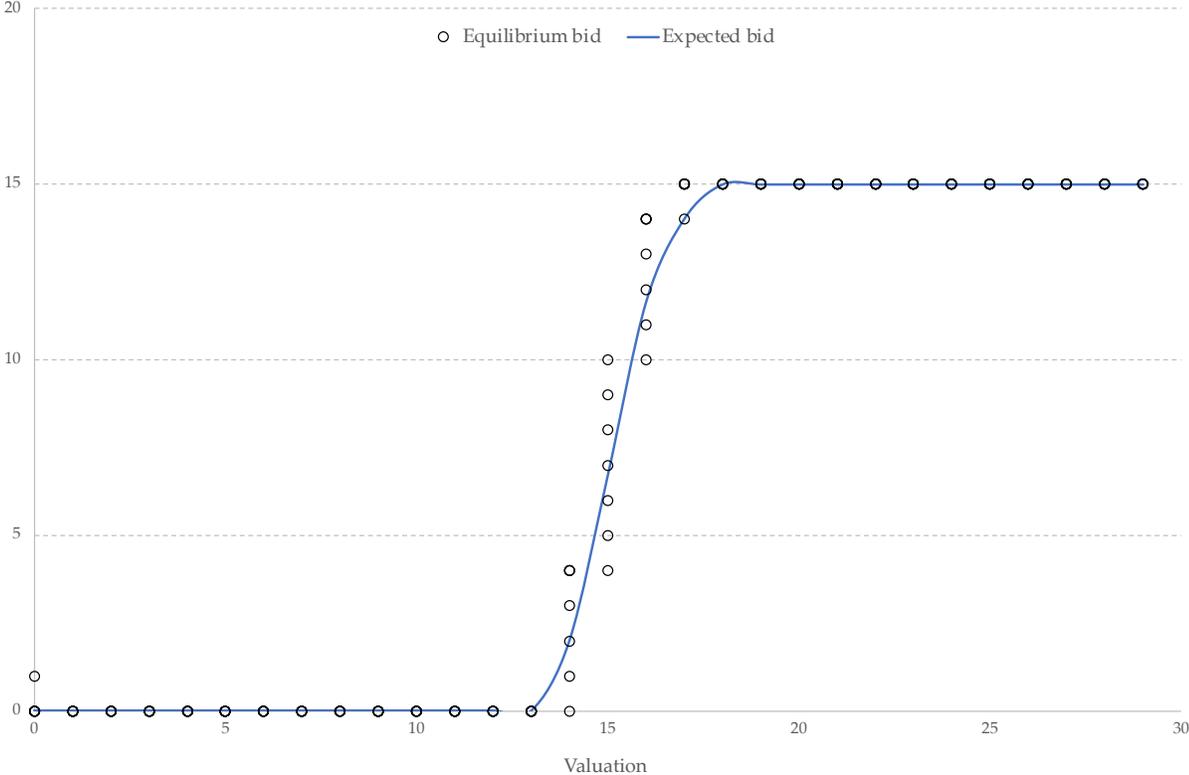
Figure 2: A mixed equilibrium with $n = 10$



Notes. This figure plots the symmetric equilibrium of the all-pay auction when values are uniform, $n = 10$, $x = 29$, and players cannot win through a tie. A dot represents a bid that is submitted with positive probability; so the set of dots at a particular value represents the equilibrium bid support at that value.

To understand the logic of Proposition 5, first consider players who do not have the maximum valuation, i.e. those for whom $v \leq x - 1$. As the number of players becomes large, such types become very unlikely to win in equilibrium since it becomes almost certain that one of their opponents will have drawn the maximum valuation. Thus, if the number of players is ‘sufficiently large’ (a sufficient condition is implicit in our proof), such types must bid 0 in any symmetric equilibrium. Given the logic of Lemmas 7–10, it then follows that players with the highest valuation will randomise over some interval of bids 0, 1, 2, ..., up to some maximum bid k . If the number of players is large, such types will be tempted to win the auction for sure by ‘over-cutting’ the maximum, i.e. submitting a bid of $k + 1$. The only way to prevent this (in the large n limit) is by setting k at $x - 1$, which means that no profitable overcutting can occur.

Figure 3: An equilibrium with normally distributed values



Notes. This figure plots the symmetric equilibrium of the all-pay auction when valuations are normally distributed ($\mu = 15, \sigma = 1$), $n = 2, x = 29$, and players cannot win through a tie. A dot represents a bid that is submitted with positive probability; so the set of dots at a particular value represents the equilibrium bid support at that value.

To illustrate the result, consider an all-pay auction with 30 bids and valuations (i.e. a maximum of $x = 29$). One can then check that if $n \geq 101$, we do indeed obtain maximal randomisation in line with Proposition 5. Of course, the idea of 101 bidders may seem rather

extreme; but it is important to emphasise that large amounts of randomisation can still be generated with fewer bidders. For example, if there are just 10 bidders, then players with the maximum valuation (29) will randomise between bids of 19, 20, ..., 26. Meanwhile, players with a valuation of 28 will randomise between bids of 13, 14, ..., 19 and so forth (see Figure 2 for illustration). Thus, randomisation need not be local; and one need not assume implausibly large numbers of bidders to obtain this result.

We should also emphasise that instead of assuming a large number of bidders, one could equally appeal to some substantial non-uniformity in the distribution of valuations. This should not come as a surprise: as far as the distribution of the maximum valuation goes, there is no difference between a model with two bidders with valuations drawn from distribution F^{n-1} , and a model with n bidders with valuations drawn from distribution F . Moreover, the large scale randomisation need not arise at the maximum valuation. To illustrate this, consider an all-pay auction with just two bidders; and suppose that the valuations are normally distributed with a mean of 15 and a standard deviation of 1 (we choose this example since the normal distribution might be a sensible choice for applications). Such a distribution concentrates probability mass in the 14 to 16 range, so one would expect it to produce randomisation at such valuations. Consistent with this, we find that such types randomise over a number of bids: for example, players with a value of 15 are predicted to randomise over all integers between 4 and 10 inclusive (see Figure 3 for illustration).

8 Concluding remarks

In this paper, we have made some progress towards understanding how auction theory changes once bids and valuations made discrete. In some cases, we find that introducing discreteness changes little. In others, however, we uncover discontinuity results. For example, in the all-pay auction with uniform values, introducing an arbitrarily small amount of discreteness makes the symmetric, pure-strategy equilibria of continuous theory disappear; and can make the pure-strategy equilibria disappear altogether. We have also shown that the remaining mixed equilibria can involve randomisation over the full range of non-dominated bids. In closing, we discuss the implications of these discontinuity results for theoretical and experimental research.

One conclusion for auction experimenters is that auction theory cannot be straightforwardly applied to the games they study. Indeed, ignoring the discreteness of laboratory games has led experimenters to make a series of errors (as documented in Sections 4 - 5). In future, experimenters may wish calculate the equilibria of their games directly, perhaps taking a

computational approach (along the lines of Section 6) or otherwise using our monotonicity, no-jumps and uniqueness lemmas to identify all *symmetric* equilibria in auction games without ties.¹²

For theorists, the implications are rather different. Over the last few decades, considerable effort has been expended in attempting to establish the existence of pure-strategy equilibria in successively more general auction models (see Maskin and Riley (2000), Athey (2001) and Reny and Zamir (2004) for notable examples). As our findings make clear, however, the existence of pure-strategy equilibria is far from guaranteed as soon as bids and valuations are made discrete. This is an issue since, as noted, discreteness is an inevitable feature of any real world or experimental auction and continuity approximations can be justified (only) as harmless assumptions that ease the analysis without substantially changing the results.¹³

As discussed, the games we study do have equilibria: this is an immediate consequence of Harsanyi (1967) and Nash (1950). The upshot, however, is that these equilibria are often in mixed strategies. The question then arises of whether it is plausible to suppose that players really do randomise when bidding in auctions (and in just the way required by auction theory) even though they have no strict incentive to do so. While we cannot address this question in full, we briefly note three points that may be relevant.

First, the ‘classical’ justification of mixed strategies — that are necessary to keep opponents guessing about what a player will do (von Neumann and Morgenstern, 1947) — would appear to have even less force than usual in this context (see Rubinstein (1991) for a general criticism of this type of justification). In auction games with plausible type spaces, individuals will already have very little idea about their opponent’s valuation. Thus, even if they somehow knew their opponents’ bidding functions, they would already have substantial uncertainty about their opponents’ bids. That is to say, it is not clear that people would ever feel the need to submit random bids in order to ‘keep their opponents guessing’.¹⁴

Second, even if individuals do choose to randomise in some environments — as suggested, for instance, by Chiappori et al. (2002) — it is not clear that they do so in auction contexts. Of course this is an empirical question and one that should be tested by empirical studies of auction behaviour. Nonetheless, we conjecture that the results of such studies will find little

¹²This algorithm is implemented in a companion paper Rasooly (2022) and is available here: <https://auctionsolver.herokuapp.com/>

¹³In fairness, we should point out that continuous equilibria (rounded appropriately to make the bids feasible) do remain ϵ -equilibria in our model. However, one should not make too much of this observation unless one can point to a compelling reason to think that ϵ -equilibria are likely to arise in auction settings.

¹⁴An alternative (and rather intriguing) rationale for randomised decision rules is that they can guarantee outcomes that are ‘not too bad’. Whatever the merits of this suggestion, however, it does not give one any reason to think that subjects will randomise with equilibrium probabilities over the equilibrium bid support.

evidence of random bidding by auction participants.

Third, introducing randomisation would appear to increase the cognitive requirements presupposed by equilibrium models. In a pure-strategy Bayes-Nash equilibrium, players are assumed to (i) know the joint distribution of opponent valuations, (ii) know which bidding functions their opponents have chosen, (iii) infer the joint distribution of opponent bids, (iv) infer the distribution of the highest opponent bid, and finally (v) calculate the optimal bid taking this distribution as given.¹⁵ Once we introduce randomisation, this already difficult task becomes even harder. For now, players need to combine their opponents' probability distribution over bidding functions with their probability distribution over valuations in order to arrive at the distribution over bids (in addition to all the other steps). In other words, players now best respond to two sources of randomness: random valuations and random bidding functions. It is unclear whether we can expect this from even relatively sophisticated auction participants.

It seems, then, that our results should push one in the direction of greater scepticism towards auction theory. Less controversially, however, they highlight the need for further research on discrete auctions. Despite the large number of papers on auction theory, the world of discrete auctions remains almost completely unexplored. We hope that the results of this paper will stimulate and provide the basis for further work in this area.

¹⁵Technically speaking, these steps are sufficient but not necessary for individuals to play a Bayes-Nash equilibrium. For instance, they might hold incorrect beliefs about their opponents' decision rules, fail to optimise correctly and then submit equilibrium bids because these mistakes end up cancelling out. In practice, however, it is very hard to see why anyone would ever play a Bayes-Nash equilibrium without proceeding through these steps.

References

- Andreoni, J., Che, Y.-K., and Kim, J. (2007). Asymmetric information about rivals' types in standard auctions: An experiment. *Games and Economic Behavior*, 59(2):240–259.
- Aseff, J. G. (2004). Learning to play second-price auctions, an experimental study. *Economics Letters*, 85(2):279–286.
- Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889.
- Bergemann, D., Brooks, B., and Morris, S. (2017). First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143.
- Boudreau, J. W. (2011). All-pay auctions with discrete action spaces. *Applied Economics Letters*, 18(2):161–162.
- Cai, G., Wurman, P. R., and Gong, X. (2010). A note on discrete bid first-price auction with general value distribution. *International Game Theory Review*, 12(01):75–81.
- Che, Y.-K. and Gale, I. (2006). Revenue comparisons for auctions when bidders have arbitrary types. *Theoretical Economics*, 1(1):95–118.
- Chen, K.-Y. and Plott, C. R. (1998). Nonlinear behavior in sealed bid first price auctions. *Games and Economic Behavior*, 25(1):34–78.
- Chiappori, P.-A., Levitt, S., and Groseclose, T. (2002). Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, 92(4):1138–1151.
- Chwe, M. S.-Y. (1989). The discrete bid first auction. *Economics Letters*, 31(4):303–306.
- Cooper, D. J. and Fang, H. (2008). Understanding overbidding in second price auctions: An experimental study. *The Economic Journal*, 118(532):1572–1595.
- Coppinger, V. M., Smith, V. L., and Titus, J. A. (1980). Incentives and behavior in english, dutch and sealed-bid auctions. *Economic inquiry*, 18(1):1–22.
- Cox, J. C., Roberson, B., and Smith, V. L. (1982). Theory and behavior of single object auctions. *Research in experimental economics*, 2(1):1–43.
- David, E., Rogers, A., Jennings, N. R., Schiff, J., Kraus, S., and Rothkopf, M. H. (2007). Optimal design of english auctions with discrete bid levels. *ACM Transactions on Internet Technology (TOIT)*, 7(2):12–es.

- Dekel, E. and Fudenberg, D. (1990). Rational behavior with payoff uncertainty. *Journal of Economic Theory*, 52(2):243–267.
- Dekel, E. and Wolinsky, A. (2003). Rationalizable outcomes of large private-value first-price discrete auctions. *Games and Economic Behavior*, 43(2):175–188.
- Drichoutis, A. C., Lusk, J. L., and Nayga, R. M. (2015). The veil of experimental currency units in second price auctions. *Journal of the Economic Science Association*, 1(2):182–196.
- Dyer, D., Kagel, J. H., and Levin, D. (1989). Resolving uncertainty about the number of bidders in independent private-value auctions: an experimental analysis. *The RAND Journal of Economics*, pages 268–279.
- Filiz-Ozbay, E. and Ozbay, E. Y. (2007). Auctions with anticipated regret: Theory and experiment. *American Economic Review*, 97(4):1407–1418.
- Georganas, S., Levin, D., and McGee, P. (2017). Optimistic irrationality and overbidding in private value auctions. *Experimental Economics*, 20(4):772–792.
- Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104(1):247–272.
- Gonçalves, R. and Ray, I. (2017). A note on the wallet game with discrete bid levels. *Economics Letters*, 159:177–179.
- Harsanyi, J. C. (1967). Games with incomplete information played by “bayesian” players, i–iii part i. the basic model. *Management science*, 14(3):159–182.
- Kagel, J. (1995). Auctions: A survey of experimental research,” handbook of experimental economics, edited by j. kagel and a. roth.
- Kagel, J. H. and Levin, D. (1985). Individual bidder behavior in first-price private value auctions. *Economics Letters*, 19(2):125–128.
- Kagel, J. H. and Levin, D. (1993). Independent private value auctions: Bidder behaviour in first-, second-and third-price auctions with varying numbers of bidders. *The Economic Journal*, 103(419):868–879.
- Kagel, J. H. and Levin, D. (2011). Auctions: A survey of experimental research, 1995-2010. *Handbook of experimental economics*, 2:563–637.
- Kirchkamp, O. and Reiß, J. P. (2011). Out-of-equilibrium bids in first-price auctions: Wrong expectations or wrong bids. *The Economic Journal*, 121(557):1361–1397.

- Li, Z. (2017). Nash equilibria in all-pay auctions with discrete strategy space. Technical report, Economics Discussion Papers.
- Maskin, E. and Riley, J. (2000). Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67(3):439–454.
- Maskin, E. S. and Riley, J. G. (1985). Auction theory with private values. *The American Economic Review*, 75(2):150–155.
- McKelvey, R. D., McLennan, A. M., and Turocy, T. L. (2014). Gambit: Software tools for game theory, version 14.1. 0.
- Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49.
- Noussair, C. and Silver, J. (2006). Behavior in all-pay auctions with incomplete information. *Games and Economic Behavior*, 55(1):189–206.
- Rasooly, I. (2022). Going... going... wrong: a test of the level-k (and cognitive hierarchy) models of bidding behaviour. *Journal of Political Economy: Microeconomics*.
- Reny, P. J. and Zamir, S. (2004). On the existence of pure strategy monotone equilibria in asymmetric first-price auctions. *Econometrica*, 72(4):1105–1125.
- Riley, J. G. (1989). Expected revenue from open and sealed bid auctions. *Journal of Economic Perspectives*, 3(3):41–50.
- Rothkopf, M. H. and Harstad, R. M. (1994). On the role of discrete bid levels in oral auctions. *European Journal of Operational Research*, 74(3):572–581.
- Rubinstein, A. (1991). Comments on the interpretation of game theory. *Econometrica: Journal of the Econometric Society*, pages 909–924.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25—55.
- von Neumann, J. and Morgenstern, O. (1947). Theory of games and economic behavior, 2nd rev.
- von Stengel, B. (1996). Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2):220–246.
- Wang, R. (1991). Common-value auctions with discrete private information. *Journal of Economic Theory*, 54(2):429–447.

9 Tables

The following tables summarise our computational results on the existence of pure-strategy Bayes-Nash equilibria. The first column ('Valuations') specifies the number of possible valuations and bids in the game. The subsequent columns specify the auction structure and number of bidders. An entry of 'yes' means that a pure-strategy Bayes-Nash equilibrium exists; an entry of 'no' means that it does not. A blank entry means that we did not search for the equilibrium (either because we had around found symmetric equilibria analytically or because the payoff matrices were too large for this to be computationally feasible).

Table 1: The existence of pure-strategy equilibria in the model with ties

Valuations	FPSB, $n = 2$	FPSB, $n = 3$	All-pay, $n = 2$	All-pay, $n = 3$
6	-	Yes	Yes	No
7	-	No	No	No
8	-	Yes	No	-
9	-	-	Yes	-
10	No	-	-	-

Table 2: The existence of pure-strategy equilibria in the model without ties

Valuations	FPSB, $n = 2$	FPSB, $n = 3$	All-pay, $n = 2$	All-pay, $n = 3$
6	-	Yes	Yes	Yes
7	-	Yes	No	No
8	-	-	No	-
9	-	-	Yes	-
10	-	-	-	-

Appendix A. Proofs

Proof of Proposition 1. In the model with ties, a standard argument establishes that the bidding function $\beta(v) = v$ weakly dominates all other bidding functions. Thus, each player sets $\beta(v) = v$ in the unique equilibrium in undominated strategies.

In the model without ties, the strategies $\beta(v) = v$ and $\beta(v) = v + \delta$ necessarily yield the same payoff (the only difference is that bidding $v + \delta$ allows a player to win when the highest bid of their opponents is v ; but in that case the payoff from winning is zero). Furthermore, a standard argument establishes that $\beta(v) = v$ and $\beta(v) = v + \delta$ weakly dominate all other bidding functions. Thus, the set of equilibria in undominated strategies is precisely the set of strategy profiles in which each player sets either $\beta(v) = v$ or $\beta(v) = v + \delta$. \square

Proof of Lemma 1. Consider the game in which bids and values are restricted to $X = \{0, \delta, 2\delta, \dots, x\delta\}$ for some $\delta > 0$. Thus, each player i has a value of δv_i for some $v_i \in \{0, 1, \dots, x\} \equiv X'$ and must choose to bid δb_i for some $b_i \in X'$. Consider initially the model without ties. In equilibrium, each player chooses their bid δb_i so that, for all $b'_i \in X'$,

$$\begin{aligned} (\delta v_i - \delta b_i)P(\text{win}|\delta b_i) &\geq (\delta v_i - \delta b'_i)P(\text{win}|\delta b'_i) \\ \iff (v_i - b_i)P(\text{win}|\delta b_i) &\geq (v_i - b'_i)P(\text{win}|\delta b'_i) \\ \iff (v_i - b_i)P(\delta b_i > \delta b_j \forall j \neq i) &\geq (v_i - b'_i)P(\delta b'_i > \delta b_j \forall j \neq i) \\ \iff (v_i - b_i)P(b_i > b_j \forall j \neq i) &\geq (v_i - b'_i)P(b'_i > b_j \forall j \neq i). \end{aligned}$$

Clearly, this inequality does not depend on δ . Thus, if there is an equilibrium in the game with x values and discretisation $\delta > 0$ in which player i bids $b\delta$ when her valuation is $v\delta$, then there is an analogous equilibrium in the game with x values and arbitrary discretisation $\delta' > 0$ in which player i bids $b\delta'$ when her valuation is $v\delta'$. A similar argument establishes that the set of equilibria is invariant to δ in the model with ties. \square

Proof of Lemma 2. Consider first the model with ties. For all v , bidding $b > v$ is weakly dominated by $b = 0$: bidding $b = 0$ ensures a non-negative payoff, whereas bidding $b > v$ leads either to a zero payoff (if the player loses) or a negative payoff (if the player wins). Hence, $\beta(v) \leq v$ in any equilibrium in undominated strategies; and therefore $\beta(0) = 0$. Moreover, for any $v \geq 1$, bidding $b = 0$ also weakly dominates bidding $b = v$ (since bidding $b = 0$ can lead to a strictly positive payoff if all other players bid 0). Thus, $\beta(v) \leq v - 1$ for all $v \geq 1$ in any equilibrium in undominated strategies.

Similar arguments apply in the model without ties. As before, bidding $b > v$ is weakly dominated by $b = 0$: the latter ensures a payoff of zero, whereas the former yields either

a negative or zero payoff. Hence, $\beta(v) \leq v$ for all v in any equilibrium in undominated strategies; which implies that $\beta(0) = 0$ and $\beta(1) \leq 1$. Suppose now that $v \geq 2$. If a player bids v , their payoff is necessarily zero. However, if they bid 1, their payoff is either zero (if their opponents all bid 1 or more) or positive (if their opponents all bid 0). Since $\beta(0) = 0$, there is a strictly positive probability that all opponents bid zero. Hence, bidding $b = 1$ now strictly dominates $b = v$, which implies that $\beta(v) \leq v - 1$ for $v \geq 2$. \square

Proof of Lemma 3. See Goeree et al. (2002) Proposition A1. \square

Proof of Lemma 4. See the main text. \square

Proof of Lemma 5. Since players avoid dominated bids, $\beta(0) = 0$ and either $\beta(1) = 0$ or $\beta(1) = 1$. Moreover, it is obvious that $\beta(2) = 1$ in any SE. We will now argue that the SE bidding function up to and including an arbitrary valuation $v \geq 2$ uniquely determines the SE bidding function up to and including valuation $v + 1$. By induction, the SE is therefore uniquely implied by our choice of $\beta(1)$.

To do this, let b denote the bid submitted when the value is $v \geq 2$. By monotonicity and no-jumps, there are only 2 possibilities when the value is $v + 1$:

1. $\beta(v + 1) = b$
2. $\beta(v + 1) = b + 1$

Let us examine these possibilities in turn.

First possibility: $\beta(v + 1) = b$. Then bidding b with a value of $v + 1$ must be at least as good as bidding $b + 1$, i.e.

$$\underbrace{(v + 1 - b)P(\leq b - 1)^{n-1}}_A \geq \underbrace{(v - b)P(\leq b)^{n-1}}_B$$

where $P(\leq k)$ denotes the equilibrium probability that a player bids k or lower.

Second possibility: $\beta(v + 1) = b + 1$. Then bidding $b + 1$ with a value of $v + 1$ must be at least as good as bidding b , i.e.

$$\underbrace{(v + 1 - b)P'(\leq b - 1)^{n-1}}_{A'} \leq \underbrace{(v - b)P'(\leq b)^{n-1}}_{B'}$$

Now notice that, in both cases, the equilibrium probability of submitting a bid that is $b - 1$ or less is the same. That is, $P(\leq b - 1) = P'(\leq b - 1)$ which in turn implies that $A = A'$.

On the other hand, the equilibrium probability of submitting a bid that is b or less does vary across the scenarios. In the first scenario, a player with a value of $v + 1$ always bids b . In the second, they never bid b and (by monotonicity) neither do players with higher values. Therefore, $P(\leq b) > P'(\leq b)$. Assuming that $v - b > 0$ — which is without loss of generality since $v \geq 2$ — this means that $B > B'$.

Let us now recall our two conditions (and substitute $A = A'$):

1. $A \geq B$
2. $A \leq B'$

Since $B > B'$, the truth of either condition implies the falsity of the other. Hence, at most one can hold; which means that $\beta(v + 1)$ is uniquely implied by $\beta(v)$. By induction, this then means that the SE (if it exists) is uniquely implied by the choice of $\beta(1)$. \square

Proof of Proposition 2. To demonstrate existence, we will show that if one player (say player 2) bids according to one of the two equilibrium strategies, then the other (player 1) can do no better than follow the same strategy.

Initially, suppose that player 2 sets $\beta_2(v_2) = \lfloor v_2/2 \rfloor$. Then their bids are drawn from either $\{0, 1, 2, \dots, x/2\}$ (if x is even) or $\{0, 1, 2, \dots, (x - 1)/2\}$ (if x is odd). First, we check that player 1 prefers to follow the proposed equilibrium strategy than to win for sure. It is easy to show that the case in which winning for sure is most appealing is when player 1's value is x (the maximum) and x is odd. In that case, player 1's payoff from winning for sure (by bidding an extra 1) is

$$x - \left(\frac{x - 1}{2} + 1 \right) = \frac{x - 1}{2}$$

Meanwhile, player 1 wins in the candidate SE if and only if her opponent's value is 2 (or more) lower than hers. Thus, player 1's equilibrium payoff is

$$\left(\frac{x - 1}{x + 1} \right) \left(x - \frac{x - 1}{2} \right)$$

This can be simplified to

$$\frac{x - 1}{2}$$

which is the exactly the expected payoff that player 1 would get by deviating. Therefore, she is never strictly better off by winning for sure instead of bidding in accordance with the SE.

Given that we have an 'interior' solution, what is player 1's optimal bid? Since player 2's

values follow a uniform distribution, the expected payoff of player 1 from making any bid $b_1 \in \mathbb{N}$ that does not win for sure is

$$\left(\frac{2b_1}{x+1}\right)(v_1 - b_1)$$

Let us write bids as deviations from the proposed equilibrium, i.e. let $b_1 = \lfloor v_1/2 \rfloor + i$ for some $i \in \mathbb{N}$. Suppose that v_1 is even so that $\lfloor v_1/2 \rfloor = v_1/2$. Writing $b_1 = v_1/2 + i$, player 1's expected payoff is

$$\left(\frac{v_1 + 2i}{x+1}\right)\left(v_1 - \frac{v_1}{2} - i\right)$$

This can be simplified to

$$\frac{v_1^2}{2(x+1)} - \frac{2i^2}{x+1}$$

which obviously attains a strict global maximum when $i = 0$, confirming the optimality of the proposed strategy.

Suppose instead that v_1 is odd. Then $\lfloor v_1/2 \rfloor = (v_1 - 1)/2$. Writing $b_1 = (v_1 - 1)/2 + i$, player 1's expected payoff is

$$\left(\frac{v_1 - 1 + 2i}{x+1}\right)\left(v_1 - \frac{v_1 - 1}{2} - i\right)$$

This can be simplified to

$$\frac{(v_1 - 1)(v_1 + 1)}{2(x+1)} + 2i\left(\frac{1 - i}{x+1}\right)$$

If $i = 0$ or $i = 1$, the second term is zero. However, if $i \geq 2$ or $i \leq -1$, the second term is negative. Thus, bidding the proposed strategy is at least as good as bidding one more and strictly better than bidding anything else.

Since $\beta_1(v_1) = \lfloor v_1/2 \rfloor$ is a best response for both odd and even v_1 , it is optimal for all v_1 and therefore constitutes an SE. A similar argument establishes the existence of an SE in which each player sets $\beta(v) = \lfloor v/2 \rfloor$. Finally, Lemma 5 implies that these are the only SE. \square

Proof of Corollary 6. In the continuous case, the unique symmetric equilibrium is given by $\beta(v) = v/2$. In the discrete case with discretisation δ , the SE are $\beta(\delta v) = \lfloor v/2 \rfloor \delta$ and $\beta(\delta v) = \lceil v/2 \rceil \delta$ for all values $\delta v \in X$ (i.e. for $v = 1, 2, \dots, k$). Thus, the difference between a discrete bid and continuous bid is at most $\delta/2 \rightarrow 0$ as $\delta \rightarrow 0$ (holding the valuation δv fixed). This establishes the (uniform) convergence of the discrete bidding function.

We now sketch a proof for the convergence of expected revenues. As $\delta \rightarrow 0$, each player's valuation converges in distribution to the valuation in the continuous case (i.e. a continuously and uniformly distributed random variable on the interval $[0, x]$). By the continuous mapping theorem, each player's bid converges in distribution to the bid in the continuous case (i.e. a continuously and uniformly distributed random variable on $[0, x/2]$). Applying the continuous mapping theorem again, we see that the maximum bid also converges in distribution to its continuous counterpart. Finally, by the bounded convergence theorem, the expected value of the maximum bid also converges to its continuous analogue. \square

Proof of Proposition 3. Let us suppose (for contradiction) that there were an SE and begin by studying the start of the bidding function that defines it. As before, $\beta(0) = 0$ and $\beta(1) = 0$ or $\beta(1) = 1$ (since players avoid dominated strategies). Either way, $\beta(2) = 1$ (since bidding 0 or 2 yield a payoff of 0 whereas bidding 1 yields a positive payoff).

Let us now consider the first value at which the bidding function does not increase. To this end, define k as the smallest value such that $\beta(k) = \beta(k - 1)$ and $k \geq 2$. First, we argue that k must exist. For contradiction, suppose that it does not, i.e. $\beta(v) = v - 1$ for all $v \geq 2$. In equilibrium, a player with $v = x$ would win whenever all of their opponents have values of $x - 1$ or lower. Since there are x such values, and their payoff in the event of winning is 1, their equilibrium expected payoff would be

$$\left(\frac{x}{x+1}\right)^{n-1}$$

If such a player were to bid one less, they would obtain an expected payoff of

$$2\left(\frac{x-1}{x+1}\right)^{n-1}$$

which (rearranging) is strictly larger if and only if

$$x > \frac{2^{\frac{1}{n-1}}}{2^{\frac{1}{n-1}} - 1}.$$

But this is precisely what we have assumed. Thus, if there is an SE, and a player with a value of x does *not* want to deviate, then there must be some $k \geq 2$ such that $\beta(k) = \beta(k - 1)$.

Next, we locate k . In equilibrium, a player with a value $v = k$ bids $v - 2$, thereby winning whenever their opponents all have values of $v - 2$ or lower. Since there are $v - 1$ such values,

the equilibrium expected payoff of a player with a value $v = k$ is

$$2 \left(\frac{v-1}{x+1} \right)^{n-1}$$

If the player were to bid 1 more (i.e. were to bid $v-1$), they would win (at least) whenever their opponents all have values of v or lower. Thus, their expected payoff would be at least

$$\left(\frac{v+1}{x+1} \right)^{n-1}$$

Since an equilibrium cannot have strictly profitable deviations,

$$2 \left(\frac{v-1}{x+1} \right)^{n-1} \geq \left(\frac{v+1}{x+1} \right)^{n-1} \tag{2}$$

We claim that $v = k$ is the smallest integer that satisfies inequality (2). To see this, notice that for all values smaller than k , submitting the equilibrium bid $\beta(v) = v-1$ must be at least as good as bidding $v-2$. This means that, for any $v < k$,

$$2 \left(\frac{v-1}{x+1} \right)^{n-1} \leq \left(\frac{v}{x+1} \right)^{n-1} < \left(\frac{v+1}{x+1} \right)^{n-1}$$

so such values cannot satisfy (2).

We can re-write (2) as

$$v \geq \frac{1 + 2^{\frac{1}{n-1}}}{2^{\frac{1}{n-1}} - 1}.$$

Since k is the smallest integer that satisfies this inequality,

$$k = \frac{1 + 2^{\frac{1}{n-1}}}{2^{\frac{1}{n-1}} - 1} + \epsilon$$

for some $\epsilon \in [0, 1)$.

Having located k , we now consider a bidder with a value $v = k-1$. By assumption, they prefer to bid $\beta(k-1) = k-2$ than to bid 1 less, i.e.

$$\left(\frac{k-1}{x+1} \right)^{n-1} \geq 2 \left(\frac{k-2}{x+1} \right)^{n-1}$$

Equivalently,

$$k \leq \frac{2(2^{\frac{1}{n-1}}) - 1}{2^{\frac{1}{n-1}} - 1}$$

Inserting our value for k , we find that

$$\epsilon \leq \frac{2^{\frac{1}{n-1}} - 2}{2^{\frac{1}{n-1}} - 1} < 0$$

which contradicts $\epsilon \geq 0$, implying that there is no SE. □

Proof of Lemmas 7 – 11. The arguments for Lemmas 7, 8, 10 and 11 are straightforward variants on those that establish analogous lemmas in the first-price auction. Lemma 9 (monotonicity) requires a slightly different proof, something we now provide. First, since players avoid dominated bids, $\beta(0) = 0$. Thus, bids cannot decrease when the value increases from $v = 0$ to $v = 1$. Next, consider now the bidding function starting from $v > 0$ and suppose (for contradiction) that

1. The optimal bid is B when the value is $v > 0$.
2. The optimal bid is $b < B$ when the value is $V > v$.

From (1), bidding B must be better than bidding b with a value of v :

$$P(\text{win}|B)v - B \geq P(\text{win}|b)v - b$$

which (since $v > 0$) implies that

$$P(\text{win}|B) - P(\text{win}|b) \geq \frac{B - b}{v}.$$

From (2), bidding b must be better than bidding B with a value of V :

$$P(\text{win}|b)V - b \geq P(\text{win}|B)V - B$$

which (since $V > 0$ and $V > v$) implies that

$$P(\text{win}|B) - P(\text{win}|b) \leq \frac{B - b}{V} < \frac{B - b}{v}$$

which is a contradiction. □

Proof of Proposition 4. To begin with, consider the model without ties. As a preliminary, let us note that $\beta(x) \leq x - 2$. The argument is straightforward. Since players avoid

dominated bids, $\beta(0) = \beta(1) = 0$. Furthermore, it is easy to check that $\beta(2) = 0$ for all $n \geq 2$. Otherwise, we could have $\beta(2) = 1$, but then the expected payoff from equilibrium bidding when $v = 2$ would be

$$2 \left(\frac{2}{x+1} \right)^{n-1} - 1 \leq 2 \left(\frac{2}{5} \right)^{n-1} - 1 \leq 2 \left(\frac{2}{5} \right) - 1 < 0$$

where the first inequality uses $x + 1 \geq 5$ and the second uses $n \geq 2$. Since the expected payoff is negative, the player could do better by bidding 0, confirming that $\beta(2) \neq 1$ and therefore $\beta(2) = 0$. By no-jumps, this means that $\beta(v) \leq v - 2$ for all $v \geq 2$ and in particular for $v = x$.

This result in hand, we now consider two separate cases: $n = 2$ and $n \geq 3$.

First case: $n = 2$. We will consider the three highest bids $\beta(x)$, $\beta(x - 1)$ and $\beta(x - 2)$, i.e. the ending of the bidding function. We will argue that none of the possible endings can form part of an SE.

Step 1. We cannot have two or more identical bids at the end, i.e. $\beta(x) = \beta(x - 1)$.

Suppose (for contradiction) that there were two or more such bids at the end and consider a player with $v = x$. If the player increased their bid by 1, the cost would be 1. Meanwhile, the benefit would be

$$v\Delta P \geq x \left(\frac{2}{x+1} \right) > 1$$

since $x > 1$. Since the deviation would be strictly profitable, we cannot have $\beta(x) = \beta(x - 1)$ in an SE. (This argument implicitly assumes that $\beta(x) \leq x - 2$; if not, the deviation would be unprofitable.)

Step 2. We cannot have two consecutive increases, i.e. $\beta(x) = \beta(x - 1) + 1 = \beta(x - 2) + 2$.

Suppose (for contradiction) that there were two consecutive increases and consider a bidder with $v = x$. If they reduced their bid by 1, the benefit would be 1. Meanwhile, the cost would be

$$\Delta P v = x \left(\frac{1}{x+1} \right) < 1$$

since $x > -1$. Since the deviation would be strictly profitable, we cannot have $\beta(x) = \beta(x - 1) + 1 = \beta(x - 2) + 2$ in an SE.

Step 3. We cannot have $\beta(x) = \beta(x - 1) + 1 = \beta(x - 2) + 1$.

Suppose (for contradiction) that this did form part of an SE bidding function and consider

a bidder with $v = x - 1$. If they increased their bid by 1, the cost would be 1. Meanwhile, the benefit would be

$$v\Delta P \geq x \left(\frac{2}{x+1} \right) > 1$$

since $x > 3$. Since the deviation would be strictly profitable, we also cannot have $\beta(x) = \beta(x-1) + 1 = \beta(x-2) + 1$ in an SE.

Now, by Lemmas 2 and 3 (no gaps and monotonicity), there are only 4 ways in which the bidding function can end:

1. $\beta(x-2), \beta(x-2), \beta(x-2)$
2. $\beta(x-2), \beta(x-2), \beta(x-2) + 1$
3. $\beta(x-2), \beta(x-2) + 1, \beta(x-2) + 1$
4. $\beta(x-2), \beta(x-2) + 1, \beta(x-2) + 2$

We can rule out the first and third possibilities since they involve two or more identical bids at the end (Step 1). We can rule out the fourth since it involves two consecutive increases (Step 2). By Step 3, we can rule out the second and final possibility. Since we can rule out all of the possibilities, there is no SE when $n = 2$.

Second case: $n \geq 3$. Here the argument is simpler: we simply show that a player with $v = x$ would prefer to win for sure than bid in accordance with their equilibrium strategy.

When $v = x$, following equilibrium yields a payoff of at most

$$x \left(\frac{x}{x+1} \right)^{n-1} - \beta(x).$$

If the player were to bid an extra 1, they would get

$$x - \beta(x) - 1.$$

They are strictly better off winning for sure if

$$x - \beta(x) - 1 > x \left(\frac{x}{x+1} \right)^{n-1} - \beta(x).$$

This can be simplified to

$$n - 1 > \frac{\ln(x) - \ln(x-1)}{\ln(x+1) - \ln(x)} \equiv g(x)$$

Assume that $x \geq 2$ (i.e. there are at least 3 possible bids/values). Then it can be shown that $g(x)$ attains its maximum at $x = 2$ where $g(x) \approx 1.7$. Thus, if $n \geq 3$ (so $n - 1 \geq 2$), it must be that $n - 1 > g(x)$. In other words, the bidder with the highest value would strictly prefer to win for sure as opposed to bidding in accordance with an SE.

Next, consider the model with ties. It will be helpful to start with a simple lemma:

Lemma 12 (No jumps at the top) *If $n = 2$ or $n = 3$, then $\beta(x - 2) = b$ implies $\beta(x - 1) \leq b + 1$; and similarly, $\beta(x - 1) = b$ implies $\beta(x) \leq b + 1$.*

Proof. Consider initially the case of $n = 2$. Suppose (for contradiction) that there is a jump at some value v , i.e. $\beta(v) = b$ but $\beta(v + 1) = b + m$ where $m \geq 2$. If a player with a value of $v + 1$ deviated to bidding $b + m$, they would gain $b + m - (b + 1) = m - 1 \geq 1$ for sure. However, they would lose (in expectation)

$$(v + 1)\Delta P = (v + 1)\frac{P(b + m)}{2} = (v + 1)\frac{\#(b + m)}{2(x + 1)}.$$

In equilibrium, the loss from deviating cannot exceed the gain of (at least) 1, yielding

$$\#(b + m) \geq \frac{2(x + 1)}{v + 1} > 2$$

since $v + 1 \leq x$. Thus, if there is a jump at a value v , it must be that $\#(b + m) > 2$. But then there cannot be a jump when the value is either $x - 1$ or $x - 2$: for in either case, there are fewer than 2 values strictly greater than v , so necessarily $\#(b + m) \leq 2$.

Consider now the case of $n = 3$. By a similar argument, we can show that $\#(b + m) \geq 2$ if there is a jump at a value v . To see this, suppose (for contradiction) that there is a jump at v but that $\#(b + m) = 1$. As before, the benefit from deviating to a bid of $b + 1$ when your value is $v + 1$ would be $m - 1 \geq 1$. Meanwhile, the expected cost would be

$$\begin{aligned} (v + 1)\Delta P &= (v + 1) \left(\frac{P(b + m)^2}{3} + P(b + m)P(\leq b) \right) \\ &= (v + 1)P(b + m) \left(\frac{P(b + m)}{3} + P(\leq b) \right) \\ &= \frac{v + 1}{x + 1} \left(\frac{1}{3(x + 1)} + P(\leq b) \right) \\ &\leq \frac{v + 1}{x + 1} \left(\frac{1}{3(x + 1)} + \frac{x}{x + 1} \right) < 1 \end{aligned}$$

since $v + 1 < 1 + x$. So the cost of deviating would be smaller than the benefit, which implies

that this cannot be an SE. Equivalently, if there is a SE with a jump at a value of v , it must be that $\#(b + m) \geq 2$.

We now show how this inequality rules out jumps at the top. When $v = x - 1$, the argument is trivial: there is only one valuation strictly greater than v , so necessarily $\#(b + m) \leq 1$. Since this contradicts $\#(b + m) \geq 2$, we cannot have a jump at $x - 1$.

Consider now the possibility of a jump when the value is $x - 2$. In any SE with such a jump, $\#(b + m) \geq 2$ and so $\beta(x - 1) = \beta(x) = b + m$. But in that case, a player with a value of x would prefer to win for sure (by bidding one more) than to bid in accordance with equilibrium. To see this, note that such a player's equilibrium payoff would be

$$x \left[\left(\frac{x-1}{x+1} \right)^2 + \frac{1}{3} \left(\frac{2}{x+1} \right)^2 + \left(\frac{2}{x+1} \right) \left(\frac{x-1}{x+1} \right) \right] - \beta(x)$$

whereas they would obtain a (certain) payoff of $x - \beta(x) - 1$ by bidding one more. It can be shown that this latter quantity is larger provided that $x \geq 2$. Thus, the deviation would be strictly profitable, implying that there is also no SE with a jump at $x - 2$. \square

This lemma in hand, we can now demonstrate the non-existence of an SE. Consider initially the case of $n = 2$. As before, we will study the last three possible bids and argue that none of the possible sequences can form part of an SE.

Case 1 ($n = 2$): $\beta(x) = \beta(x - 1) + 1 = \beta(x - 2) + 2$.

Consider deviating to bidding $\beta(x - 1)$ when $v = x$. The gain from doing so would equal 1. Meanwhile, the (expected) cost would be

$$v\Delta P = x \left(\frac{0.5}{x+1} + \frac{0.5}{x+1} \right) = \frac{x}{x+1} < 1.$$

Since the cost is lower than the gain, this is a strictly profitable deviation — which means that this sequence of bids cannot form part of an SE.

Case 2 ($n = 2$): $\beta(x) = \beta(x - 1) = \beta(x - 2)$.

Consider bidding one more when $v = x$. The cost of doing so would equal 1. Meanwhile, the (expected) benefit would be

$$v\Delta P \geq x \left(\frac{1.5}{x+1} \right) > 1$$

provided that $x \geq 3$. Since the benefit exceeds the cost, this is a strictly profitable deviation,

ruling out this sequence of bids in an SE.

Case 3 ($n = 2$): $\beta(x) = \beta(x - 1) + 1 = \beta(x - 2) + 1$.

Consider bidding one more when $v = x - 1$. The cost of doing so equals 1. The (expected) benefit is

$$v\Delta P \geq (x - 1) \left(\frac{1}{x + 1} + \frac{0.5}{x + 1} \right) > 1$$

provided that $x \geq 6$. Again, this is a strictly profitable deviation, which means that this is not an SE.

Case 4 ($n = 2$): $\beta(x) = \beta(x - 1) = \beta(x - 2) + 1$.

Consider bidding 1 more when $v = x - 2$. The cost is 1, whereas the (expected) benefit is

$$v\Delta P \geq (x - 2) \left(\frac{0.5}{x + 1} + \frac{1}{x + 1} \right) > 1$$

provided that $x \geq 9$. So this is (once again) not an SE.

Given monotonicity and no-jumps, this exhausts all the cases, showing that there is no SE with 2 bidders.

Consider now the case of $n = 3$. The argument has a similar structure to before:

Case 1 ($n = 3$): $\beta(x) = \beta(x - 1)$.

In that case, a bidder would prefer to bid 1 more when $v = x$; in fact, this follows directly from the proof of the previous lemma (letting $m = 1$).

Case 2 ($n = 3$): $\beta(x) = \beta(x - 1) + 1$.

Consider a bidder with $v = x - 1$ who deviates by bidding one more. The cost of doing so is 1. Meanwhile, the expected benefit is $v\Delta P$. Before the deviation,

$$P(\text{win}) \leq \left(\frac{x - 1}{x + 1} \right)^2 + \frac{1}{3} \left(\frac{1}{x + 1} \right)^2 + \left(\frac{x - 1}{x + 1} \right) \left(\frac{1}{x + 1} \right)$$

After the deviation,

$$P(\text{win}) = \left(\frac{x}{x + 1} \right)^2 + \frac{1}{3} \left(\frac{1}{x + 1} \right)^2 + \left(\frac{x}{x + 1} \right) \left(\frac{1}{x + 1} \right)$$

A quick calculation thus reveals that

$$v\Delta P \geq \frac{2vx}{(x+1)^2} = \frac{2x(x-1)}{(x+1)^2} > 1$$

provided that $x \geq 5$. Thus, the benefits from bidding 1 more exceeds the costs, confirming that this cannot comprise part of an SE. Furthermore, by monotonicity and no-jumps, these are the only possible cases — implying that there is no SE at all.

Finally, consider the case of $n \geq 4$. To begin, we find an upper bound on the probability with which a bidder with a value of x (the maximum) can win in equilibrium. In the best case scenario, $\beta(x) > \beta(x-1)$, i.e. opponents with strictly lower values submit strictly lower bids. In that case, a player with a value of x would win with probability

$$\left(\frac{x}{x+1}\right)^{n-1} + \left(\frac{x}{x+1}\right)^{n-2} \left(\frac{1}{x+1}\right) \binom{n-1}{1} \left(\frac{1}{2}\right) + \dots + \left(\frac{1}{x+1}\right)^{n-1} \frac{1}{n}$$

since they win if all opponents have lower values (the first term), if all but one has a lower value and they win the coin flip that decides the tie (the second term, note that this event can happen in $n-1$ ways since they have $n-1$ opponents to tie with), ..., or finally if all their opponents also have the maximum value but they are still chosen to win (the last term). If we define $p \equiv 1/(x+1)$, we can write the sum $s(n, p)$ more compactly as

$$s(n, p) = \sum_{i=1}^n (1-p)^{n-i} p^{i-1} \binom{n-1}{i-1} \frac{1}{i}. \quad (3)$$

Let us now evaluate (3) in closed form. As a preliminary, we note that

$$\binom{n-1}{i-1} \frac{1}{i} = \frac{(n-1)!}{(n-i)!(i-1)!} \frac{1}{i} = \frac{1}{n} \frac{n!}{(n-i)! i!} = \frac{1}{n} \binom{n}{i}.$$

We can thus write (3) as

$$\begin{aligned}
s(n, p) &= \sum_{i=1}^n (1-p)^{n-i} p^{i-1} \frac{1}{n} \binom{n}{i} \\
&= \frac{1}{np} \sum_{i=1}^n (1-p)^{n-i} p^i \binom{n}{i} \\
&= \frac{1}{np} \left(\sum_{i=0}^n (1-p)^{n-i} p^i \binom{n}{i} - (1-p)^n \right) \\
&= \frac{1}{np} \left([(1-p) + p]^n - (1-p)^n \right) \\
&= \frac{1 - (1-p)^n}{np}
\end{aligned} \tag{4}$$

where the penultimate equality follows from the binomial theorem. We claim that $s(n, p)$ is decreasing in n . To see this, consider the difference between each term and the next, i.e.

$$s(n, p) - s(n+1, p) = \frac{1 - (1-p)^n}{np} - \frac{1 - (1-p)^{n+1}}{(n+1)p} = \frac{1 - (1-p)^n(1+pn)}{n(n+1)p}. \tag{5}$$

Now notice that

$$(1-p)^{-n} = \sum_{k=0}^{\infty} \binom{n+k-1}{k} p^k = 1 + np + \dots > 1 + np$$

which in turn implies that

$$(1-p)^n(1+pn) < 1$$

and so the difference in (5) is positive. This means that $s(n, p) > s(n+1, p)$, i.e. our upper bound on the probability of winning falls as the number of bidders increases.

We are now in a position to make the main argument. When $n = 4$, we can use (4) to bound the equilibrium expected payoff that a player with a value of x can receive. Recalling that $p \equiv 1/(x+1)$, we see that their payoff cannot exceed

$$x \left(\frac{1 - \left(\frac{x}{x+1}\right)^4}{4 \left(\frac{1}{x+1}\right)} \right) - \beta(x). \tag{6}$$

Meanwhile, if they were to bid 1 more, they would win the auction for sure, getting

$$x - \beta(x) - 1 \tag{7}$$

It is straightforward but tedious to check that the quantity in (7) is strictly larger than the quantity in (6) if $x > 3.465$. Thus, if $x \geq 4$ (i.e. there are at least 5 values), the deviation is strictly profitable; with the implication that there is no SE.

Finally, recall that $s(n, p)$ is decreasing in the number of bidders n . Thus, if the number of bidders rises, we can impose an even tighter upper bound on the equilibrium payoff of a bidder with a value of x . In contrast, however, the payoff from the deviation does not change: it is fixed at $x - \beta(x) - 1$. Hence, if the deviation is strictly profitable when $n = 4$, it must also be strictly profitable for arbitrary $n \geq 4$. This concludes the proof. \square

Proof of Proposition 5. First, we show that, if n is sufficiently large, then all players with values $v \leq x$ must bid 0 with probability 1 in any SE. To this end, consider any SE and observe that, for such players, equilibrium expected pay-offs π_{eq} can be bounded as follows:

$$\pi_{eq}(v, b) = vP(\text{win}|b) - b \leq v \left(\frac{x}{x+1} \right)^{n-1} - b \rightarrow -b \text{ (as } n \rightarrow \infty) \quad (8)$$

where the inequality holds since such types never win against opponents with valuations of x . Since expected payoffs become negative for any $b \geq 1$ if n is sufficiently large, this implies that such types must all bid 0.

Given the logic of lemmas 7 – 10, we see that players with a valuation of x must then randomise over an set of consecutive integers whose smallest element is either 0 or 1, and whose largest element we denote k (see Rasooly (2022) for details). In fact, the smallest element must be 0: for it were 1, then

$$\pi_{eq}(v = x, b = 1) = vP(\text{win}|b) - b = v \left(\frac{x}{x+1} \right)^{n-1} - b \rightarrow -b \text{ (as } n \rightarrow \infty) \quad (9)$$

and so the $v = x$ type would lose money in expectation in the large n limit (and would prefer to bid 0). Given that the $v = x$ type randomises over the interval $\{0, 1, \dots, k\}$, and is therefore indifferent between all bids in this interval, we see that

$$\pi_{eq}(v = x) = \sum_{i=0}^k \pi_{eq}(v = x|b = i)P(b = i) = \sum_{i=0}^k \pi_{eq}(v = x|b = 0)P(b = i) = 0 \quad (10)$$

But then it must be that $k = x - 1$: if $k \leq x - 2$ (the only other non-dominated possibility), the type could earn a strictly positive expected payoff by deviating to a bid of $k + 1$. \square

Appendix B. FPSB with fair tie-breaking

We begin by considering the case of $n = 2$ and proving a qualified version of our no-jumps lemma:

Lemma 13 *Let $n = 2$. Then there is no SE with a value v at which $\beta(v - 1) = b - k$ where $k \geq 2$ but $\beta(v) = b$ provided that $\beta(v) \geq (2 - \sqrt{3})v \approx 0.27v$.*

Proof. Suppose that there is a jump in the SE bidding function, i.e. there exists a $k \geq 2$, $v \in X$ and $b \in X$ such that $\beta(v - 1) = b - k$ but $\beta(v) = b$. Let m denote the number of values that induce a bid of b . In equilibrium, bidding b with a value of v must be at least as good as bidding $b - 1$:

$$\pi(v, b) = (v - b) \left(\frac{v + 0.5m}{S} \right) \geq \pi(v, b - 1) = (v - b + 1) \frac{v}{S}.$$

Hence,

$$m \geq \frac{2v}{v - b}.$$

If we write $b = \alpha v$, our inequality becomes

$$m \geq \frac{2}{1 - \alpha}. \tag{11}$$

Intuitively, the number of b bids (m) must be large so that the loss of probability of winning from decreasing one's bid is large — that way, the deviation is unprofitable.

Now, let v' denote the maximum value that submits a bid of b (so that $v' \equiv v + m - 1$). In equilibrium, bidding b with a value of v' must be at least as good as bidding $b + 1$:

$$\pi(v', b) = (v' - b) \left(\frac{v' + 1 - 0.5m}{S} \right) \geq \pi(v', b + 1) \geq (v' - b - 1) \left(\frac{v' + 1}{S} \right)$$

Rearranging, we obtain

$$m \leq \frac{2(v' + 1)}{v' - b}. \tag{12}$$

Intuitively, the number of b bids (m) must be small so that the gain in probability from bidding 1 more is small — this makes the deviation unprofitable.

To argue that we cannot have jumps an SE, we will claim that (11) implies the falsity of

(12). That way, it is impossible for both inequalities to hold. To do this, we will argue that

$$\frac{2}{1-\alpha} > \frac{2(v'+1)}{v'-b} \iff v' > v + \frac{1}{\alpha} - 1$$

(writing $b = \alpha v$), which holds since

$$v' \geq v + \frac{2}{1-\alpha}$$

and

$$v + \frac{2}{1-\alpha} > v + \frac{1}{\alpha} - 1$$

provided that $\alpha > 2 - \sqrt{3}$. □

Proposition 6 *In the first-price auction with two bidders and an odd maximum value, there is an SE in which $\beta(v) = \lfloor v/2 \rfloor$.*

Proof. The argument is essentially the same as before. The only difference is that the probability of winning is now the probability that one's bid is strictly higher than the opponent's bid plus the probability that it is the same multiplied by 1/2. □

Proposition 7 *With two bidders and $x \geq 10$, the only possible SE is $\beta(v) = \lfloor v/2 \rfloor$.*

Proof. Since players avoid dominated strategies, $\beta(0) = 0$, $\beta(1) = 0$ and $\beta(2) = 1$. Plainly, then, $\beta(v) = \lfloor v/2 \rfloor$ holds up to (and including) $v = 2$. We will now argue that, if it holds up to any $v \geq 2$, it must hold up to $v + 1$. By induction, it then must hold for all v .

First case: the value is odd. Let us denote the value by $2v + 1$. Since $\beta(v) = \lfloor v/2 \rfloor$ holds up to the value $2v$, we know that $\beta(2v) = v$ (since $2v$ is even) and $\beta(2v - 1) = v - 1$ (since $2v - 1$ is odd). We wish to demonstrate that $\beta(2v + 1) = v$ (i.e. the bidding function does not increase).

To do so, let us suppose (for contradiction) that $\beta(2v + 1) \neq v$. Since $\beta(2v + 1) \geq (2v + 1)/2 \geq (2 - \sqrt{3})v$, we may apply the no-jumps lemma, which implies that $\beta(2v + 1) \leq v + 1$. By monotonicity, $\beta(v + 1) \geq v$; and so the assumption that $\beta(2v + 1) \neq v$ implies that $\beta(2v + 1) = v + 1$. Now, in equilibrium, a player with a value of $2v + 1$ must prefer to bid $v + 1$ than to bid only v , i.e.

$$\pi(2v + 1, v + 1) = v \left(\frac{2v + 1 + 0.5\#(v + 1)}{S} \right) \geq \pi(2v + 1, v) = (v + 1) \left(\frac{2v + 0.5}{S} \right)$$

where $\#(v + 1)$ denotes the number of values that induce a bid of $v + 1$. This simplifies to

$$\#(v + 1) \geq 3 + \frac{1}{v}$$

which implies (since $v > 0$) that $\#(v + 1) > 3$; and therefore (since $\#(v + 1)$ is an integer) that $\#(v + 1) \geq 4$. Hence, we conclude that there must be at least 4 valuations that induce a bid of $v + 1$.

We now argue that this leads to a contradiction. There are two (sub)cases to consider. In the first case, suppose that $2v + 4 \notin X$ (i.e. the value $2v + 1$ is sufficiently close to the end of the bidding function that $2v + 4$ exceeds the maximum value). In that case, it is clearly impossible that $\#(v + 1) \geq 4$, so we have our contradiction immediately. In the second case, suppose that $2v + 4 \in X$ and consider the situation when $v = 2v + 4$. We claim that such a bidder would be made strictly better off by bidding 1 more:

$$\begin{aligned} \pi(2v + 4, v + 2) &\geq (v + 2) \left(\frac{2v + 1 + \#(v + 1)}{S} \right) > \\ \pi(2v + 4, v + 1) &= (v + 3) \left(\frac{2v + 1 + 0.5\#(v + 1)}{S} \right) \end{aligned}$$

which simplifies to

$$\#(v + 1)(v + 1) > 4v + 2$$

which holds since $\#(v + 1) \geq 4$. Since this deviation is profitable, there is no SE in which $\beta(2v + 1) = v + 1$, implying that $\beta(2v + 1) = v$. This concludes the proof for the odd case.

Second case: the value is even. Let us denote the value by $2v$ and assume (initially) that $2v \geq 12$, i.e. the value is 6 or higher. (We will consider the case where $2v < 12$ later.) By the inductive hypothesis, $\beta(v) = \lfloor v/2 \rfloor$ holds up to $2v - 1$. In particular, this means that $\beta(2v - 1) = v - 1$, $\beta(2v - 2) = v - 1$, $\beta(2v - 3) = v - 2$ and $\beta(2v - 4) = v - 2$.

We wish to show that $\beta(2v) = v$ (i.e. the bidding function increases). To do this, suppose (for contradiction) that $\beta(2v) \neq v$. By monotonicity, $\beta(2v) \geq v - 1 > (2 - \sqrt{3})2v$. Thus, we may again apply the no-jumps lemma; so our assumption that $\beta(2v) \neq v$ implies that $\beta(2v) = v - 1$.

Now, one requirement in equilibrium is that a player with a value of $2v - 3$ cannot strictly benefit by bidding one more than they are supposed to. That is,

$$\pi(2v - 3, v - 2) = (v - 1) \left(\frac{2v - 4 + 1}{S} \right) \geq \pi(2v - 3, v - 1) = (v - 2) \left(\frac{2v - 2 + 0.5\#(v - 1)}{S} \right)$$

Hence,

$$\begin{aligned} (v-1)(2v-3) &\geq (v-2)(2v-2+0.5\#(v-1)) \\ \iff \#(v-1) &\leq 2\left(\frac{v-1}{v-2}\right) \end{aligned} \tag{13}$$

Intuitively, the number of values that induce a bid of $v-1$ must be small: otherwise, the gain from increasing one's bid to $v-1$ will be sufficiently large that the player will prefer to deviate.

Since $v \geq 6$, inequality (13) implies that

$$\#(v-1) \leq 2\left(\frac{5}{4}\right) = 2.5.$$

Plainly, however, this is inconsistent with our assumption that $\beta(2v) = v-1$: for in that case, we must have $\#(v-1) \geq 3$.

The previous argument establishes that, if $\beta(2v) = \lfloor v/2 \rfloor$ holds up to all even $v \geq 6$, then it must hold for all even v . To complete the argument, it remains to check the cases of $v = 4$ and $v = 2$. (When $v = 0$, $\beta(2v) = \lfloor v/2 \rfloor$ holds trivially.)

To begin, consider the case of $v = 2$. By contradiction, assume that $\beta(2) \neq 1$. Clearly, this implies that $\beta(2) = 0$. If this is an equilibrium, a player with a value of 2 must then prefer to bid 0 than to bid 1:

$$\pi(2, 0) = 2\left(\frac{0.5\#(0)}{S}\right) \geq \pi(2, 1) = \frac{\#(0) + 0.5\#(1)}{S} \iff \#(1) \leq 0$$

Thus, if this is an equilibrium, then there can be no bids equal to 1; which means that the bids are all 0 or greater than 2. By the no-jumps lemma, players then must bid 0 for all values satisfying

$$2 \geq (2 - \sqrt{3})v \iff v \leq \frac{2}{2 - \sqrt{3}} \approx 7.46$$

Hence, $\beta(v) = 0$ for all $v \in \{0, 1, \dots, 7\}$. It is then easy to check, however, that a player with a $v = 7$ would then strictly benefit if they deviated to a bid of 1, contradicting the supposed optimality of equilibrium play.

Finally, consider the case of $v = 4$. By the usual argument, assuming (for contradiction) that $\beta(4) \neq 2$ implies that $\beta(4) = 1$. Furthermore, a player with a value $v = 4$ must prefer

their equilibrium bid to bidding one more:

$$\pi(4, 1) = 3 \left(\frac{2 + 0.5\#(1)}{S} \right) \geq \pi(4, 2) = 2 \left(\frac{2 + \#(1) + 0.5\#(2)}{S} \right)$$

or equivalently

$$4 \geq \#(1) + 2\#(2) \tag{14}$$

Since $\beta(2) = \beta(3) = \beta(4) = 1$, $\#(1) \geq 3$. From (14), we thus conclude that $\#(2) = 0$, i.e. there are no bids of 2. Applying the no-jumps lemma, we infer that $\beta(v) = 1$ for all v that satisfy

$$3 \geq (2 - \sqrt{3})v \iff v \leq \frac{3}{2 - \sqrt{3}} \approx 11.2$$

Thus, $\beta(v) = 1$ for all $v \in \{4, 5, \dots, 10\}$. (Since $x \geq 10$, these values are all well defined.) It is easy to check, however, that a player with a value $v = 10$ would strictly benefit by increasing their bid from $b = 1$ to $b = 2$. This concludes the proof for the even case. \square

Corollary 14 *If $x \geq 10$, there is no SE in the first-price auction with 2 bidders and an even maximum value.*

Proof. From the previous, any SE must have the form $\beta(v) = \lfloor v/2 \rfloor$. However, consider the situation when $v = x$ (the maximum). Equilibrium play calls for $b = 0.5x$, so the equilibrium payoff is

$$\frac{x}{2} \left(\frac{x}{x+1} + \frac{1}{2(x+1)} \right)$$

However, if the player deviated to $b = 0.5x - 1$, they would get

$$\left(\frac{x}{2} + 1 \right) \left(\frac{x-2}{x+1} + \frac{1}{x+1} \right)$$

It can be shown that the latter quantity is strictly larger assuming that $x \geq 5$. Thus, there is no SE in which $\beta(v) = \lfloor v/2 \rfloor$ – and therefore no SE at all. \square

Going... going... wrong: a test of the level- k (and cognitive hierarchy) models of bidding behaviour*

Abstract

In this paper, we design and implement an experiment aimed at testing the level- k model of auctions. We begin by identifying (simple) environments that optimally disentangle the predictions of the level- k model from the natural benchmark of Bayes-Nash equilibrium. We then implement these environments within a virtual laboratory in order to see which theory can best explain observed bidding behaviour. Overall, our findings suggest that, despite its notable success in predicting behaviour in other strategic settings, the level- k model (and its close cousin, cognitive hierarchy) cannot explain behaviour in auctions.

KEYWORDS: auction, behavioral game theory, experimental design, level- k models

JEL CODES: C72, C90, D44, D83, D90

**Links:* [algorithm](#); [pre-registration](#); [experiment](#).

1 Introduction

Although the study of auctions remains dominated by equilibrium based approaches, in recent years more ‘behavioural’ alternatives have emerged. Most prominent among these is Crawford and Iriberry’s (2007) level- k model, which posits that individuals iteratively best respond to a naive ‘first thought’ about how one might bid in an auction. If accurate, the level- k model would substantially revise our understanding of how auctions ought to be designed (Crawford et al., 2009; De Clippel et al., 2019); and would similarly revise how we think about individual reasoning in incomplete information settings. Understanding whether the level- k model is accurate is thus a critical issue.

Surprisingly, however, the existing literature provides little insight into the predictive performance of level- k and whether it can outperform the natural benchmark of Bayes-Nash equilibrium. There is some evidence that the level- k model might be able to explain the ‘winner’s curse’ in common value auctions (Crawford and Iriberry, 2007; Costa-Gomes and Shimoji, 2015); but also evidence that it might do so for the wrong reasons (Ivanov et al., 2010). In addition, there does not appear to be any serious comparison of level- k and equilibrium in the rather simpler independent private setting (IPV) that forms the starting point for most auctions research. For example, while Crawford and Iriberry (2007) do compare level- k and equilibrium using an IPV dataset, the auction structure that generated this dataset leads level- k and equilibrium to make near identical predictions — making a formal comparison of the models almost impossible.^{1,2}

It is also difficult to learn very much about the relative performance of level- k and equilibrium using existing experimental data sets since these have often been generated by first-price auctions with uniformly distributed valuations (Kagel and Levin, 1995). Unfortunately, the predictions of level- k and equilibrium entirely coincide in such settings (for any $k \geq 1$). Moreover, while it is easy to disentangle level- k and equilibrium using ‘exotic’ value distributions, it is doubtful that such distributions would be understood by all subjects in any actual experiment. The challenge is therefore to find an experimental design that is both sufficiently simple for subjects to understand and yet dramatically disentangles the models.

¹Indeed, Crawford and Iriberry (2007) restrict themselves to comparing level- k with a rival behavioural theory, quantal response equilibrium.

²See also Kirchkamp and Reiß (2011) for a brief experimental comparison of level- k and equilibrium, albeit in a paper not devoted to level- k reasoning. Since the goal of their paper was not to assess the performance of level- k , it is perhaps not surprising that their experiments do not cleanly disentangle level- k and equilibrium: for example, in their first experiment, the predictions of the models co-incide entirely. See also Gillen (2009), An (2017), Hortaçsu et al. (2019) and Galavotti et al. (2018) for work on the level- k model in the context of observational data.

We begin by proposing an environment that does exactly this: a discrete³ all-pay auction with uniformly distributed values. We first analyse the equilibrium predictions in this environment, proving that it possesses exactly one symmetric equilibrium. We thus extend a well known result from continuous auction theory to the discrete auction model – and do so using completely novel (and elementary) arguments.⁴ In the course of our proof, we construct an algorithm⁵ that is guaranteed to identify this equilibrium, finding that it closely approximates the continuous equilibrium of traditional auction theory given the parameters chosen in our experiment. In this equilibrium, bids increase (quadratically) in valuations, rising to around half of valuations at the maximum possible valuation that players can draw.

We then conduct a similar analysis of the level- k model but uncover a surprising contrast. Solving the level-1 player’s optimisation problem reveals a corner solution: such types bid zero for all valuations. As a result, level-2 and level-3 types also bid close to zero; the former ‘overcutting’ the level-1s by (almost always) bidding 1 and the latter ‘overcutting’ the level-2s by (almost always) bidding 2. Clearly, this is very different from equilibrium; and so we have found an environment that dramatically separates the models. In addition, this environment leads the models to possess very different comparative statics. For example, in the level- k model making the bid discretisation more coarse substantially inflates the predicted bids; whereas in the equilibrium model, it has very little effect.

Next, we conduct a similar exercise in the first-price auction, again assuming uniformly distributed values to make the environment experimentally implementable (and to force the two level-0 specifications proposed by Crawford and Iriberri (2007) to coincide – we elaborate on this point later on). To break the equivalence of the level- k and equilibrium models, we introduce the possibility that subject bids are ‘cancelled’ (in which case they lose the auction automatically). We then solve for the cancellation probability that maximises the ‘distance’ between the two theories. This turns out to separate the models in a similar way: level- k bids are once again close to zero, whereas equilibrium bids are roughly a quarter of valuations. Thus, we once again manage to disentangle the models in an environment simple enough to be implemented in an experiment.

We then conduct an experiment designed to put these insights into practice. The basic

³We work with discrete models throughout both for the sake of realism (any actual auction must be discrete) but also because the level- k model would be ill-defined in a continuous action setting (we elaborate more upon these points later on).

⁴Our argument holds for arbitrary value distributions and can be straightforwardly modified to allow for factors like risk aversion. It also extends immediately to the first-price auction; indeed, we suspect that it could be extended to fairly arbitrary auction structures (although we have not formally verified this).

⁵To aid future experimenters, we have implemented the algorithm and made it available online. See: <https://auctionsolver.herokuapp.com/>

idea of the experiment is straightforward. We create the aforementioned environments in a virtual laboratory, allowing us to see which of the models under consideration better predict the observed bids. However, the experiment also possesses various other features aimed at testing the predictive success of the level- k model. For example, we elicit individuals' levels using a variant on Arad and Rubinstein (2012)'s 11-20 game; study whether behaviour varies with the bid discretisation in the manner predicted by level- k theory; and finally ask individuals to explain why they bid in the way that they did to see if the level- k model accurately characterises their conscious reasoning processes.

The experiment yields four central findings. First, if individual levels are restricted to the 1–3 range (to capture an intuitive limit on the number of thinking steps that individuals can be expected to perform), then the level- k predictions are far lower than the observed bids. In addition, the level- k predictions are substantially less accurate than those of equilibrium. These conclusions hold across all the auction structures and treatments that we consider and under a variety of approaches to model testing and comparison. Moreover, this conclusion continues to hold if one allows individuals to vary across auction rounds but retains the requirement that individual levels are 'plausible' (i.e. in the 1–3 range).

Second, if one fits the model without any such plausibility restrictions, one finds levels typically estimated in the 30–35 range. We view this almost as *reductio ad absurdum* of the model since it conflicts with almost all prior experimental work on this issue (Crawford et al., 2013), clashes with basic intuition about the number of thinking steps that individuals can plausibly be expected to conduct, and sits uneasily beside findings from linguistics that individuals struggle to even comprehend higher order statements whose order exceeds three (Arad and Rubinstein, 2012). Moreover, methods for penalising model flexibility (e.g. the Bayesian information criterion) generally do not support the inclusion of these very high levels.

Third, individuals' levels as estimated from the auctions bear almost no resemblance to their levels as inferred from a game known to trigger level- k reasoning (the 11-20 game). This is hard to square within the context of the level- k model: for if levels reflect either cognitive sophistication or beliefs about the cognitive sophistication of one's opponents, then one would expect to find a positive correlation between estimated levels across games. Conversely, the absence of positive correlation across any two games can be taken as *prima facie* evidence that level- k reasoning is not operative in at least one of the games under consideration.

Finally, we find that subjects are unlikely to cite iterated reasoning when asked to explain why they bid in the way that they did. In other words, subjects are unlikely to offer

conjectures about the bidding strategies of their opponents – and even less likely to ground these conjectures in what they take to be their opponents’ conjectures about their own bidding strategy. Admittedly, we do see evidence of iterated reasoning in the explanations given by at least one subject – and this subject was (perhaps not coincidentally) likely to bid in line with the predictions of the level- k model. However, such reports are very rare, suggesting that the level- k model captures at best a small fraction of observed bidding behaviour.

We then conduct a series of extensions and robustness checks. For example, we consider the impact of incorporating risk aversion into the models, allowing subjects to best respond to a distribution over levels (as in Camerer et al. (2004)), changing the L_0 specification, as well as conducting various tweaks to our empirical analyses (e.g. including subjects who submit dominated bids). We do not find that any of these exercises alter our substantive conclusions.

The remainder of this article is structured as follows. Section 2 outlines an equilibrium and level- k analysis of the discrete all-pay and first-price auctions. Section 3 presents our experimental design. Section 4 contains our empirical results and Section 5 our robustness checks. Finally, Section 6 concludes with a discussion of why the level- k model does so poorly in the auction setting despite its well documented successes in other domains of strategic behaviour.

2 Separating level- k from equilibrium

2.1 The all-pay auction

To begin, let us describe our first environment which sharply disentangles the level- k and equilibrium models: a (discrete) all-pay auction with uniformly and independently distributed values. There are $n \geq 2$ bidders, each assumed to be risk neutral (we relax this assumption in Section 5). As in our experiment, each player’s valuation v_i is drawn from the set $\mathbb{X} = \{0, 1, \dots, x\}$ for some $x \in \mathbb{N}^+$; and the bids that they may submit are restricted to the very same set \mathbb{X} . Let $\beta: \mathbb{X} \rightarrow \mathbb{X}$ denote a generic pure strategy (i.e. ‘bidding function’). We study the case of discrete values and bids since (i) this is inevitably the set-up in the experiment that follows, as indeed it must be in any auction experiment; (ii) if we had instead assumed that the bids were continuous, as is more standard in the literature, then the predictions of the level- k model would be ill-defined.⁶

⁶In particular, the level-2 player would not possess a best response: for any number $\epsilon > 0$ that they might bid, there exists a feasible bid $\epsilon' < \epsilon$ that would yield them a higher expected pay-off.

Since this is an all-pay auction, individuals need to pay their bid no matter whether they win the object. However, they win the object if and only if their bid is strictly higher than all bids submitted by their opponents. Note that we assume that individuals cannot win the object in the event of a tie. While this is a slightly unusual assumption, it matches the tie-breaking rule that we adopt in the experiment. This tie-breaking rule, in turn, was chosen because it tremendously simplifies the equilibrium (and to a lesser extent the level- k) analysis.⁷

To begin with, we consider what Bayes-Nash Equilibrium (henceforth ‘equilibrium’) predicts that individuals will do in this situation. In particular, we will focus on *symmetric* equilibria. We do this for three reasons. First of all, it is unclear how individuals could be expected to coordinate on asymmetric equilibria without any communication in a one-shot game. Second, symmetric equilibria are analytically more tractable. Third, based on computational experiments reported in Rasooly and Gavidia-Calderon (2020), one may suspect that discrete auction games like those described above do not possess asymmetric equilibria which are substantially different from their symmetric equilibrium.⁸ Therefore, one might hope that the restriction to symmetric equilibria is ‘almost’ without loss of generality.

Unfortunately, since our auction is discrete we cannot directly apply results from ‘standard’ auction theory. We can, however, establish the following very useful result (which holds entirely independently of the distribution from which each player’s valuation is drawn).

Proposition 1 *The discrete all-pay auction has exactly one symmetric equilibrium.*

Proposition 1 contains two claims: that a symmetric equilibrium exists and that it is unique. Establishing the existence part is straightforward: since the game is finite, one can essentially just apply Harsanyi (1967). However, establishing uniqueness is considerably more involved: see Appendix A for the details. Importantly, the uniqueness proof is based on an *algorithm* that explicitly constructs the (candidate) equilibrium (see Appendix B for examples of the algorithm in action). Thus, our result does not merely tell us that there is exactly one

⁷While the tie-breaking rule assumed here mirrors that used in our experiment, one can obtain quantitatively similar equilibrium and level- k predictions using alternative tie-breaking rules (assuming that the number of bids and valuations is sufficiently large to make ties somewhat improbable).

⁸Rasooly and Gavidia-Calderon (2020) are able to completely enumerate all equilibria of a series of discretised all-pay auctions through a brute search of all possible strategy profiles. The exercise shows that, in the games that they consider, any asymmetric equilibria are near identical to the (unique) symmetric equilibrium. While this provides evidence against the existence of ‘very asymmetric’ equilibria, the auctions studied computationally by Rasooly and Gavidia-Calderon (2020) are smaller than those considered in our experiment, hindering a direct application of their results.

symmetric equilibrium, but also shows us exactly how to compute it.⁹

Having developed this algorithm, we now use it to compute the symmetric equilibrium in the case of the distribution (uniform) used in our experiment. Figure 1 displays the result for the case of $x = 15$ and $n = 3$: a dot represents a bid that is submitted with positive probability. As can be seen, the equilibrium is in (properly) mixed strategies; but it is closely approximated, quantitatively speaking, by the continuous equilibrium¹⁰ of ‘textbook’ theory:

$$\beta(v) = \left(\frac{n-1}{n} \right) \frac{v^n}{x^{n-1}} \quad (1)$$

As in the continuous model, equilibrium bids increase slowly when valuations are low but quickly when valuations are high. When values equal the maximum x , bids are roughly equal to the equilibrium bid in the first-price auction, i.e. $x(n-1)/n$.

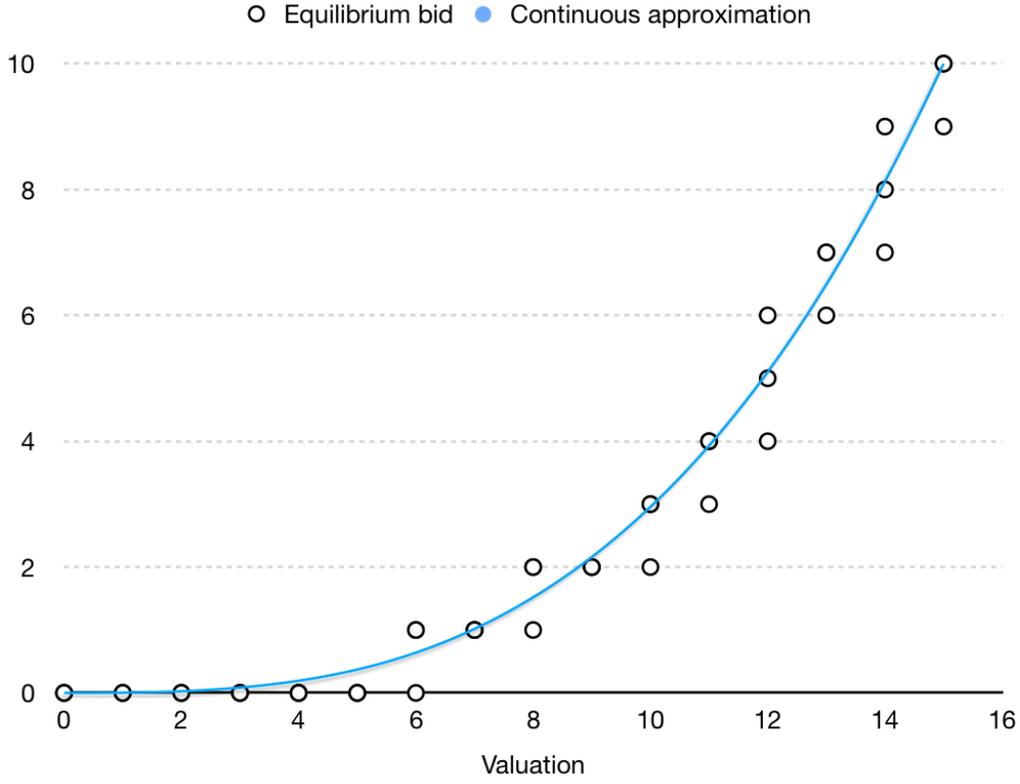
Having discussed the predictions of equilibrium (our main benchmark), we now turn to the predictions of the level- k model. As formulated by Crawford and Iriberri (2007), the model begins with a ‘naive’ type of player (the level-0) who either bids their valuation (the ‘truthful’ specification) or otherwise bids any element of the strategy space \mathbb{X} with equal probability (the ‘random’ specification). This is meant to capture a player’s ‘first thought’ about how their opponents will play the game. Higher types are then defined inductively: for $k \geq 1$, a level- k player submits the bid that maximises their expected payoff given that their opponents are all level $k-1$ (see Crawford and Iriberri (2007) for interpretation and elaboration). In the following, we will write $\beta^k: \mathbb{X} \rightarrow \mathbb{X}$ to denote a level- k player’s bidding strategy.

We now derive the behaviour of the types. Assume, as in our experiment, that values are uniformly distributed on \mathbb{X} (retreating from the generality of Proposition 1). Given the ‘truthful’ specification ($b = v$), level-0 bids are then uniform on \mathbb{X} . Note, however, that this is also the ‘random’ specification. We thus see that, quite aside from making the experimental rules easier for subjects to grasp, imposing uniform values ensures that the two level-0 specifications lead to identical predictions, a great advantage when it comes to testing the model.

⁹This is invaluable from an experimental perspective since ‘brute force’ approaches to equilibrium computation take several hours even for auctions with a handful of possible values and essentially intractable for auctions of the size considered in our experiment (Rasooly and Gavidia-Calderon, 2020). In contrast, our algorithm is able to compute the symmetric equilibria of our experimental games in less than a second.

¹⁰See, for instance, Klemperer (1999) for a derivation.

Figure 1: Equilibrium in the all-pay auction



Notes: This figure plots the symmetric equilibrium of the all-pay auction with uniform values, $n = 3$, and $x = 15$. A dot represents a bid that is submitted with positive probability.

Given that level-0 bids are uniform, the level-1 player solves the problem

$$\max_{b \in \mathbb{X}} \pi(v, b) \equiv v\mathbb{P}(\text{win}|b) - b = v \left(\frac{b}{x+1} \right)^{n-1} - b. \quad (2)$$

This leads to a corner solution: bidding 0 for all values $v \in \mathbb{X}$. To see why this is true, note that the benefit to such a player of increasing their bid by 1 is the gain in probability $\Delta\mathbb{P}$ multiplied by their valuation v . If $n = 2$, then $\Delta\mathbb{P} = 1/(x+1)$; so the benefit of bidding one more is $v/(x+1)$. Meanwhile, the cost of bidding one more is simply 1. Since $v/(x+1) < 1$ for all $v \in \mathbb{X}$, the benefit of increasing one's bid is always lower than the cost – which means that the optimal bid is 0. Moreover, if it is optimal to bid 0 when $n = 2$, it must be optimal for any $n \geq 2$: adding more bidders decreases the expected payoff from submitting any positive bid, while leaving the pay-off from bidding 0 fixed.

It is then straightforward to derive the strategies of higher levels. For now, assume that a level- k player breaks ties by choosing the lowest optimal bid (while this assumption is natural

since it allows us to avoid dominated bids, we relax it in Section 5). Given this assumption, a level-2 player sets $\beta(v) = 1$ for every value $v \geq 2$ ('over-cutting' their level-1 opponents); and sets $\beta(v) = 0$ otherwise. Similarly, a level-3 player sets $\beta(v) = 2$ for all $v \geq 3$; and sets $\beta(v) = 0$ otherwise. Thus, for all 'reasonably low' levels, a level- k type player bids

$$\beta_{LK}(v) = \begin{cases} k - 1 & \text{if } v \geq k \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

The strategy β_{LK} describes level- k behaviour at 'reasonably low' levels. At higher levels, however, this simple pattern breaks down. The strategy β_{LK} calls for a level- k player with a value of $k - 1$ to bid 0, thereby earning a (certain) payoff of 0. There is another option, however: bid 1 in the hope of winning when one's opponents all bid 0. As k rises, this second option becomes more attractive; and eventually it pays to deviate from β_{LK} . It is easy to check, however, that level- k players must bid β_{LK} provided that

$$k \leq (x + 1)^{\frac{n-1}{n}} + 1; \quad (4)$$

an inequality that gives concrete meaning to the phrase 'reasonably low' in the previous paragraph.

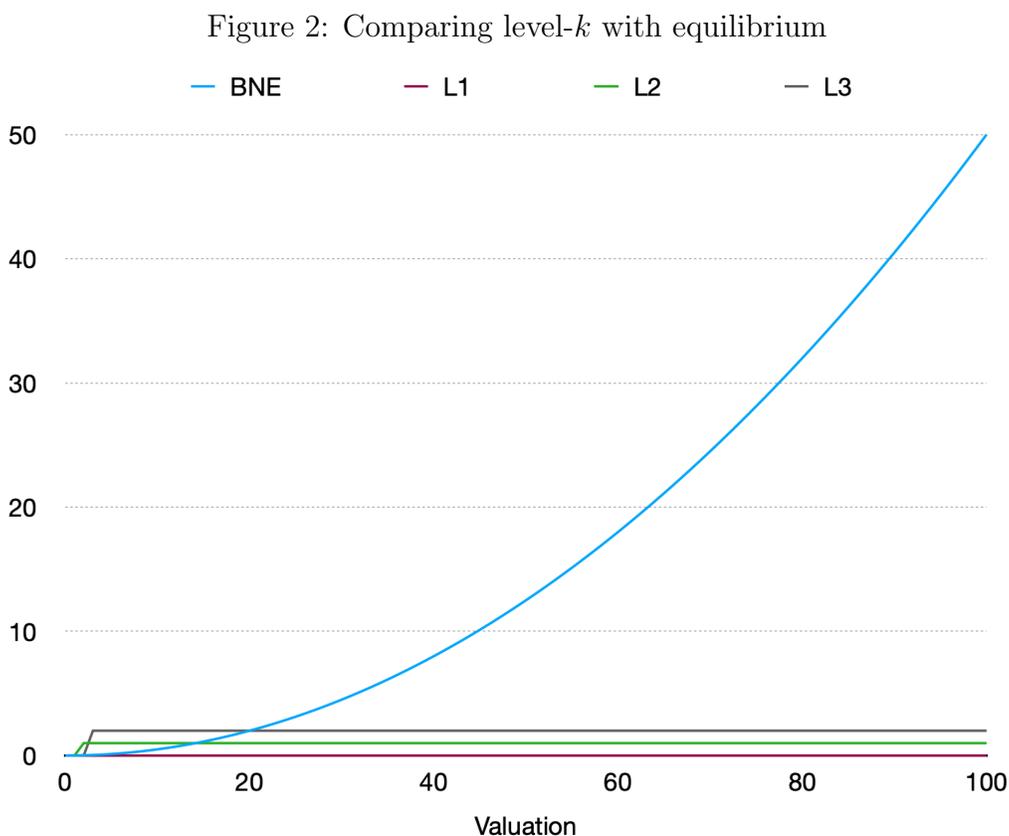
Finally, we argue that level- k can never coincide with equilibrium; with the implication that level- k bidding functions must cycle as $k \rightarrow \infty$. To prove this, note that, by Proposition 4 of Rasooly and Gavidia-Calderon (2020), the discrete all-pay auction with uniform values does not possess a symmetric equilibrium in pure strategies. Of course, it does have a symmetric equilibrium (see our Proposition 1), so this must be in mixed strategies. Now, given our assumption about level- k tie breaking, it is clear that a level- k player could never choose a mixed strategy. As a result, for any $k \in \mathbb{N}^+$, level- k can never coincide with the symmetric equilibrium. Moreover, this means that the level- k predictions must cycle. To see this, suppose that they did not cycle. Since the set of bidding functions is finite, there would need to exist some $k \in \mathbb{N}$ such that β^k coincides with $\beta^{k'}$ for all $k' > k$. (Otherwise, a strategy would need to re-appear at some stage – which would necessarily trigger a cycle.) But then β^k would be a symmetric, pure strategy equilibrium, contradicting Rasooly and Gavidia-Calderon (2020)'s Proposition 4.

The following proposition summarises the findings of the previous paragraphs.

Proposition 2 *For every level k that satisfies (4) and every value $v \in \mathbb{X}$, a level- k bidder sets $\beta^k(v) = \beta_{LK}(v)$. Moreover, there is no $k \in \mathbb{N}$ at which β^k coincides with equilibrium;*

with the implication that β^k cycles as $k \rightarrow \infty$.

The figure below plots both the equilibrium predictions (or technically the continuous approximation) and level- k predictions (for levels 1–3). As can be seen, the auction structure has done an excellent job disentangling the predictions of the models. First, given any reasonable calibration of the levels, the level- k model predicts bids that are substantially lower than those predicted by equilibrium. Second, while the equilibrium bidding function is generally increasing, the level- k bidding functions are flat at all but one value, providing a second contrast which one can test. Finally, our non-convergence result means that the theories remain separated at arbitrarily high levels. It follows then, that we have found a powerful way to disentangle the models, and to do so without abandoning the uniform values specification which is so useful when implementing any actual experiment.



Notes: This figure plots the equilibrium and level- k predictions for the all-pay auction with uniform values, $n = 2$, and $x = 100$.

2.2 The first-price auction

It is clearly desirable to find a second auction structure which sharply separates the models – that way, we can check the relative performance of the models in two different strategic

settings. To this end, we now examine the possibility of doing this within the first-price auction. As noted by Crawford and Iriberri (2007), the level- k and equilibrium models coincide (for every $k \geq 1$) in the first-price auction with uniformly distributed values. To break this equivalence, we now suppose that each player's bid only goes through with probability $p \in (0, 1)$; otherwise their bid is cancelled and they cannot win the auction. We now consider what probability p maximally disentangles the theories.

Formally, the model structure is much as before: for example, values and bids are restricted to the same set of consecutive integers \mathbb{X} . However, there are two differences. First, since this is now a first-price auction, a player pays their bid if and only if they win the object. Second, as just mentioned, each player's bid is now cancelled with probability $1 - p \in (0, 1)$; and a player wins if and only if their (non-cancelled) bid is strictly higher than all the non-cancelled bids submitted by their opponents.

As before, we begin by demonstrating that a symmetric equilibrium exists and is unique¹¹:

Proposition 3 *If $p \in (0, 1)$, then the discrete first-price auction has exactly one symmetric equilibrium.*

The arguments used to establish this result are very similar to those used to prove Proposition 1. In particular, the argument is based on the very same algorithm that constructs the equilibrium. As before, using this algorithm reveals that the equilibrium bids are well approximated by the continuous equilibrium, which is now (under uniform values)¹²

$$\beta(v) = \left(\frac{n-1}{n}\right)v - \frac{x(1-p)}{np} \left[1 - \left(\frac{1-p}{1-p+p(v/x)}\right)^{n-1}\right] \quad (5)$$

As can be seen, introducing the possibility that bids are cancelled makes the bidding less aggressive. If $p = 1$, i.e. bids are never cancelled, our expression reduces to $\beta(v) = v(n-1)/n$.

¹¹While our uniqueness result still holds under arbitrary value distributions, it now requires the restriction that $p < 1$; the role of this assumption is to ensure that all bids have positive probability of winning and thus to rule out equilibria involving dominated bids. Since this restriction holds in our experiment, this technical point need not concern us here.

¹²Equation (5) may be derived by noting the formal equivalence of our model with a model with an uncertain number of bidders. One can then apply Theorem 1 of Harstad et al. (1990). Alternately, one can note that our model is also equivalent (at least as far as undominated equilibria are concerned) with a model in which each player's value distribution has a mass of probability on zero. One can then use the 'standard formula' (see, e.g., Krishna (2009))

$$\beta(v) = v - \int_0^v \frac{G(y)}{G(v)} dy,$$

and insert $G(v) = (1 - p + pv/x)^{n-1}$ (recalling that $G(v)$ is the probability that the maximum valuation is v or lower).

In other words, we recover the standard formula for equilibrium with uniformly distributed values. Moreover, it is easy to check that for any $v \in \mathbb{X}$, $\beta(v)$ is strictly increasing in p . Finally it can be shown (using L'Hôpital's rule) that $\beta(v) \rightarrow 0$ as $p \rightarrow 0$ for all $v \in \mathbb{X}$. This is also not surprising: if one is almost certain to win the auction anyway (in the event that one's own bid is not cancelled), then one may as well bid $b \approx 0$.

We now turn to the predictions of the level- k model. Given that level-0 bids are either cancelled or uniform on \mathbb{X} , a level-1 player solves the problem

$$\max_{b \in \mathbb{X}} (v - b)\mathbb{P}(\text{win}|b) = (v - b) \left(1 - p + \frac{pb}{x + 1}\right)^{n-1} \quad (6)$$

Since the objective function in (6) is single peaked, it is obvious that the integer that maximises this function can be obtained by finding the $b \in [0, x]$ that maximises the function and rounding up or down. In other words, we have a good continuous approximation of the exact level-1 bidding function. To simplify formulae, we will therefore proceed as if level-1 types could choose any $b \in [0, x]$ (although we will always use the exact model predictions when analysing the data from our experiment). Routine optimisation subject to the inequality constraint then reveals that

$$\beta^1(v) = \max \left\{ \left(\frac{n-1}{n}\right)v - \left(\frac{1-p}{p}\right)\frac{x}{n}, 0 \right\} \quad (7)$$

As in equilibrium, level-1 bids are increasing in the probability p . Moreover, for all $p \leq 1/n$, the level-1 player sets $\beta(v) = 0$ for all $v \in \mathbb{X}$, i.e. we recover the same corner solution as in the all-pay auction. Intuitively, if p is quite low, then there is a large chance that the player's opponents' bids are all cancelled – so the player may as well just bid 0 and receive their entire valuation whenever they ‘get lucky’.

Our goal is to disentangle the models. To this end, we define the distance between the theories as

$$d(p) \equiv \int_0^x |\beta(v) - \beta^1(v)| dv. \quad (8)$$

Of course, the distance $d(p)$ is simply the total area between the (approximated) level-1 and equilibrium bidding functions. We now study the probability p^* that maximises this distance. Recall from above that the level-1 player bids $\beta(v) = 0$ for all $v \in \mathbb{X}$ provided that $p \leq 1/n$. Meanwhile, every equilibrium bid $\beta(v)$ is strictly decreasing in p (moving closer to 0). From this, we see that we could never maximise (8) by choosing some $p \in [0, 1/n)$. The reason is that increasing p (by a sufficiently small $\epsilon > 0$) would increase every equilibrium bid $\beta(v)$ while leaving the level- k bids unchanged, thus increasing the distance between the

two theories. In other words, we have the following easy result.

Proposition 4 $p^* \geq 1/n$.

Unfortunately, p^* cannot be solved analytically. Nonetheless, we can compute the optimal p numerically given various values of n (see Appendix C). We find, for example, that $p^* \approx 0.54$ if $n = 2$, $p^* \approx 0.34$ if $n = 3$ and $p^* \approx 0.26$ if $n = 4$. Overall, then, choosing $p = 1/n$ seems like a good approximation of the optimal probability (the loss in terms of distance is small). This has the further advantage of making the probabilities easier for subjects to grasp than the exact solutions would be; so we ultimately opt for $p = 1/n$ in our experiment.

Assuming that $p = 1/n$, it is then straightforward to derive the exact level- k bidding strategies. As noted, the level-1 player sets $\beta^1(v) = 0$ for all $v \in \mathbb{X}$. As a result, the level-2 player bids 1 (when their value is high) or bids 0 (when their value is low). Similarly, the level-3 player bids 2 for high values and 0 otherwise. This simple pattern characterises the behaviour of level- k bidders when k is ‘reasonably low’ relative to the number of valuations $x + 1$. More formally, we can say the following:

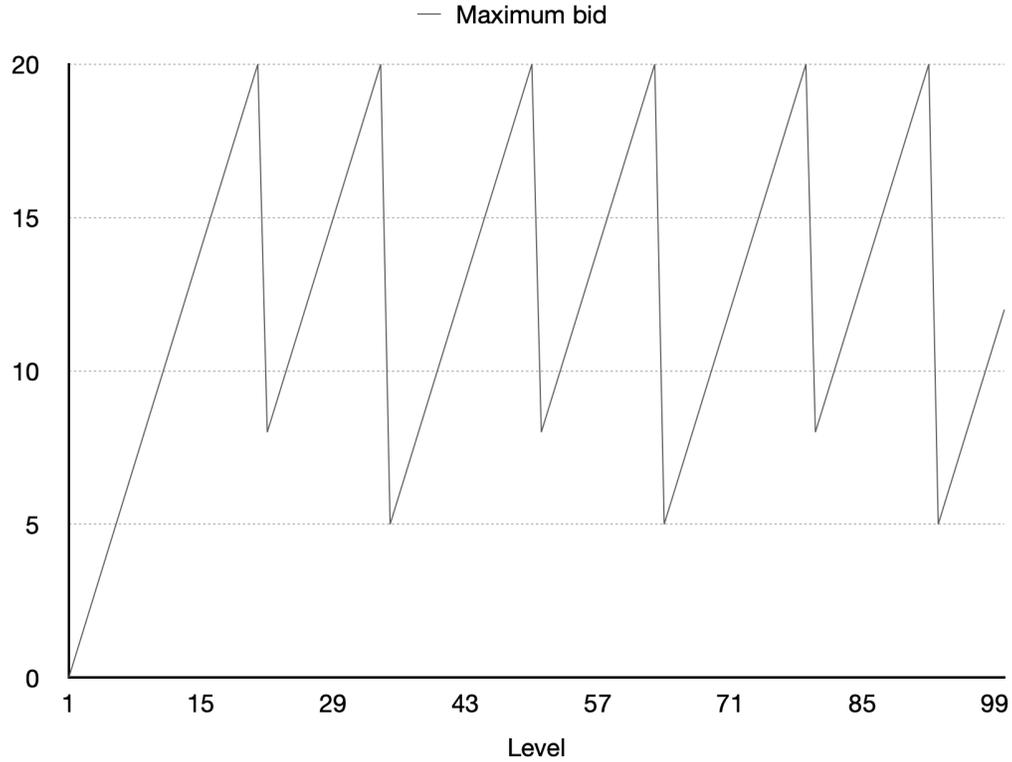
Proposition 5 *Fix a set of types $\mathbb{K} = \{1, 2, \dots, \bar{k}\}$, let $p = 1/n$, and define*

$$v^*(k) = \frac{k - 1}{1 - (1 - p)^{n-1}}$$

Then if x is sufficiently large, every level- k type in \mathbb{K} bids $\beta^k(v) = k - 1$ for $v > v^(k)$ and $\beta^k(v) = 0$ otherwise.*

Given the parameters ultimately chosen for our experiment ($n = 2$, $p = 1/2$, $x = 100$), level- k behaviour once again cycles at very high levels. To illustrate this, Figure 3 plots the maximum bid submitted by the first 100 levels (the full level- k bidding functions are available in the supplementary materials). As can be seen, the model produces cycles of length 29 (so level-7 coincides with level-36, which in turn coincides with level 65, etc.) Moreover, far from cycling locally around the equilibrium predictions, level- k cycles are quite dramatic: the maximum bid ranges from 0 to 20. This is all the more surprising once we realise that all bids above 50 are strictly dominated (by bidding zero), so the highest level- k bid traces out around 40% of the non-dominated strategy space.

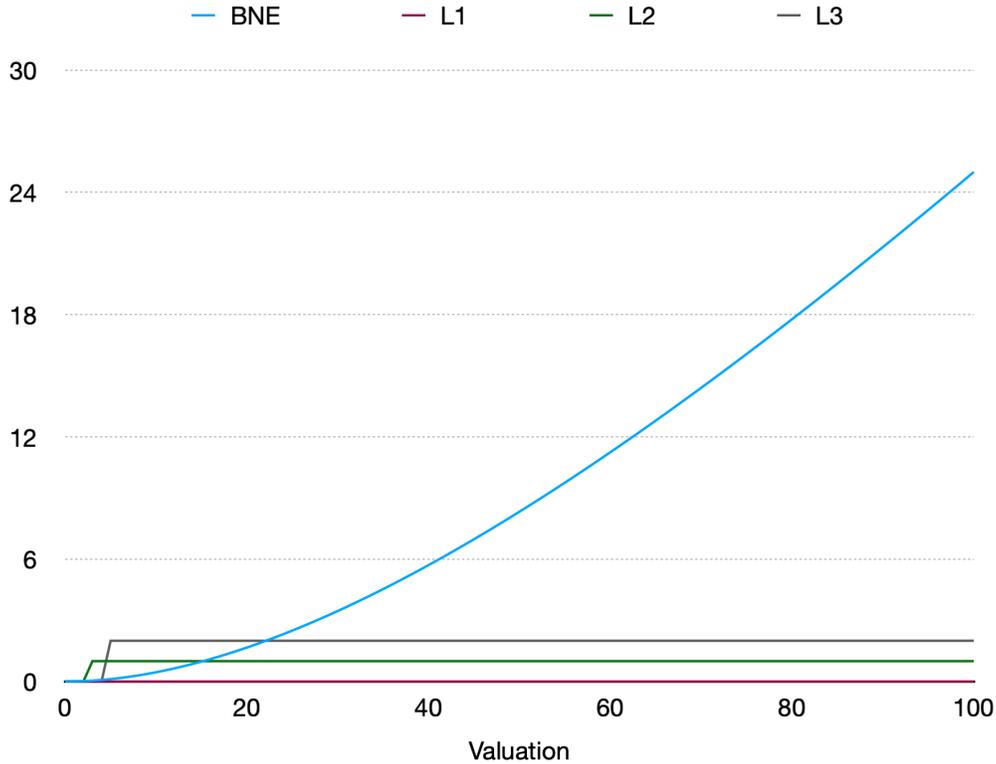
Figure 3: Level- k cycling



Notes: This figure depicts level- k cycling in the first-price auction with uniform values, $n = 2$ and $x = 100$. Specifically, it plots the maximum bid submitted by a level- k player, for all $k \in \{1, 2, \dots, 100\}$.

Finally, Figure 4 compares the models in the case of $n = 2$ and $p = 1/2$. As can be seen, we have once again managed to disentangle the theories – and in a dramatic way. First, given any plausible calibration of the levels – recall that these are supposed to be in the 1–3 region – level- k bids are significantly lower than equilibrium bids. Second, while equilibrium bids are increasing in values, level- k bids are almost always non-increasing. Finally, the non-convergence of level- k to equilibrium means that the models remain separated even if one were to ascribe implausibly high levels to the experimental participants.

Figure 4: Comparing level- k with equilibrium



Notes: This figure plots the equilibrium and level- k predictions for the first-price auction with uniform values, $n = 2$, $p = 0.5$, and $x = 100$.

3 Experimental design

To test the insights of the previous section, we conducted an experiment with 84 participants in collaboration with the Nuffield Centre for Experimental Social Sciences. All participants were students at the University of Oxford and no participant took part in more than one session. The experiment was conducted online and written in oTree (Chen et al., 2016).¹³

The essential idea of the experiment is straightforward – to implement the two aforementioned auction structures in a virtual laboratory and ask subjects how much they wanted to bid. It is then a relatively straightforward matter to check which of the two theories under consideration can better predict (and rationalise) the observed bidding behaviour.

More precisely, each subject in the experiment participated in a first-price auction (as described previously) with $n = 2$, $p = 1/2$ and an all-pay auction (as described in the previous section) with $n = 2$. Following the auctions, subjects played a modified version of Arad and

¹³The full experimental instructions can be viewed here: <https://auction-experiment.herokuapp.com/>

Rubinstein (2012)’s 11-20 game that is proposed in Alaoui and Penta (2016). The purpose of this game was to allow for a rough calibration of each participant’s level.¹⁴ Finally, we measured each subject’s risk aversion using the ‘bomb’ risk elicitation task (Crosetto and Filippin, 2013).

Before participating in either of the auctions, subjects were given extensive quizzes to check their understanding of the rules. They were not permitted to proceed past the quiz until they had answered a full set of questions correctly. We view this as an important feature of the study: subjects are unlikely to make meaningful choices if they cannot even understand the rules of the game that they are playing.

Each subject bid in two rounds of each of the auction types.¹⁵ Moreover, we did not inform subjects of the bids submitted by their opponents in previous rounds. We made these two choices – only two rounds coupled with very little feedback – since the level- k model is supposed to be a model of initial play, as opposed to a model of the outcomes of a long-run learning process (Crawford and Iriberri, 2007).¹⁶ In other words, we sought to test the level- k model within its intended domain of applicability.

Subjects were informed (and regularly reminded) that they would never bid against the same opponent twice across the rounds of an auction. That is, we used the ‘perfect stranger matching’ protocol. We opted for this protocol since the two theories under consideration are static theories – in the sense that they treat the game as a one-shot interaction – and we did not want to complicate the analysis by introducing learning or reputational considerations.

Bidding was incentivised through a procedure similar to that employed in Filiz-Ozbay and Ozbay (2007). In every round, individuals were asked how much they would want to bid given ten possible valuations that they might have. They were informed that they should bid carefully since one of these valuations would be their actual valuation and they would be committed to bidding the amount they had entered. In other words, one in ten bids were selected to ‘count’. We opted for this procedure, as opposed to incentivising every bid, since it allowed us to collect more data given a fixed number of rounds, thereby further reducing feedback or learning effects (and making the environment more hospitable to the level- k

¹⁴We use Alaoui and Penta (2016)’s modification of the original game since it prevents cycling (level- k behavior coincides for all $k \geq 9$). As Alaoui and Penta (2016) note, this is then useful if one wishes to calibrate individual levels without the potentially controversial assumption that each individual’s level is the smallest possible level that rationalises their observed behaviour.

¹⁵They were also asked to explain the general considerations determining how they chose to bid after every second round.

¹⁶Indeed, it is for this reason that Crawford and Iriberri (2007) only test their model using data from the first five rounds of the experiments they examine.

model).¹⁷

Finally, we randomly allocated the participants into two treatments: an ‘integer bid’ treatment (to which 2/3 of participants are allocated), and a treatment in which bids were constrained to be multiples of 5. The idea here was to check whether varying the bid treatment has the strong effects predicted by level- k theory or the relatively weak effects predicted by the equilibrium model. In the level- k model, increasing the discretisation by a factor of 5 increases the predicted bids by roughly a factor of 5, assuming that levels are sufficiently low for individual behavior to be characterised by Propositions 2 and 5 (see the supplementary materials for an exact description of the effects of the bid discretisation on the model’s predictions). In contrast, varying the bid discretisation does not have large quantitative effects on equilibrium bids: for example in the first-price auction the average equilibrium bid increases by just 4.7% (again, see the supplementary materials).

The experiment was pre-registered along with a full analysis plan.¹⁸ The experiment also received Ethics Approval from both the University of Oxford’s and Centre for Experimental Social Sciences’ institutional review boards.

4 Results

4.1 Pooling the data

We begin by asking which theory best explains the overall patterns observed in the bidding data. To this end, we pool all the valuations and bids, compute the mean observed bid for every valuation and plot this against the corresponding predictions of the two models. For the moment, we assume that levels cannot exceed three¹⁹ and drop the handful of subjects who submitted at least one dominated bid in one of the auctions.

Figure 5 displays the results for the case of the first price auction with integer bids (see also Figure 8 for the treatment in which bids are multiples of 5). As found in many prior studies of first-price auctions, we observe substantial overbidding relative to the (risk-neutral) equilibrium benchmark (Kagel and Levin, 1995). Indeed, the average bid is around 46% of the average value, whereas equilibrium predicts that it should be around 20%. However, while

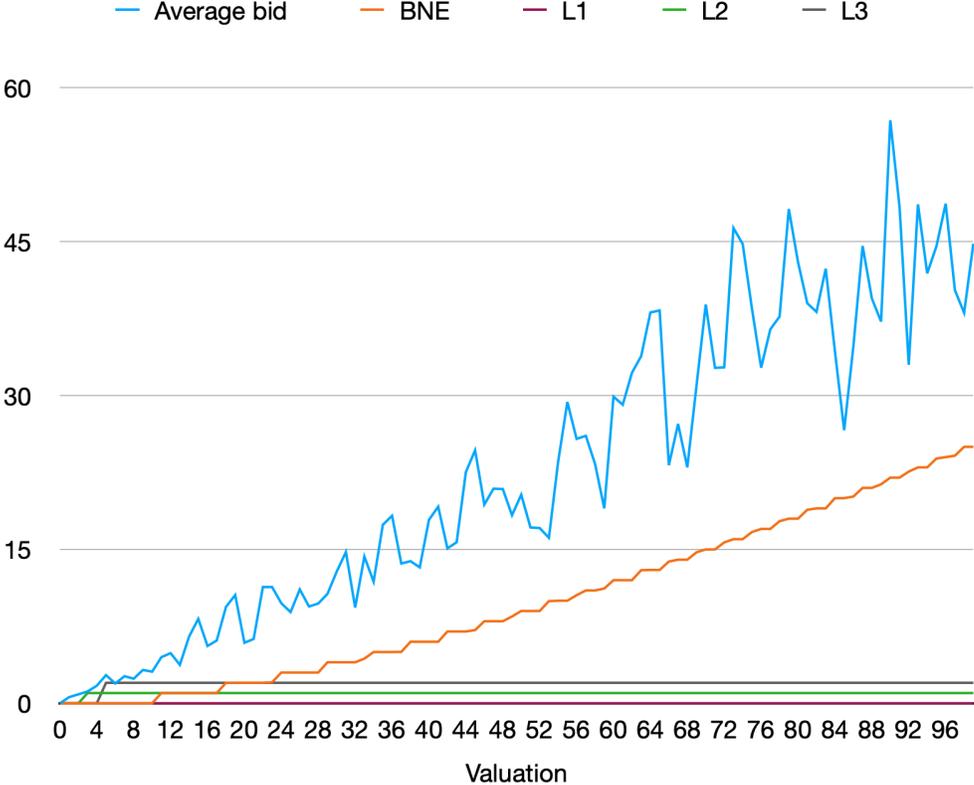
¹⁷On average, subjects earned £9.60 over the course of around 40 minutes. Total earnings ranged from £3.60 to £16.00 with a standard deviation of £2.34.

¹⁸See <https://www.socialscisearch.org/trials/8011>.

¹⁹We impose this constraint because all of the evidence of which we are aware suggests that insofar as individuals use level- k reasoning, it rarely goes beyond the third level. Indeed, this is what we find when examining the responses of our subjects in the 11-20 game: out of those who mentioned level- k reasoning when asked to explain their choices, not one made a choice consistent with more than three levels of reasoning.

the equilibrium model struggles to match the data, it is clear that the level- k model performs considerably worse, predicting bids that are at least an order of magnitude lower than those commonly observed. As a result, it is clear that equilibrium provides a substantially better (albeit highly imperfect) description of the average bidding function observed in the experimental data.

Figure 5: Predictions and data in the first-price auction

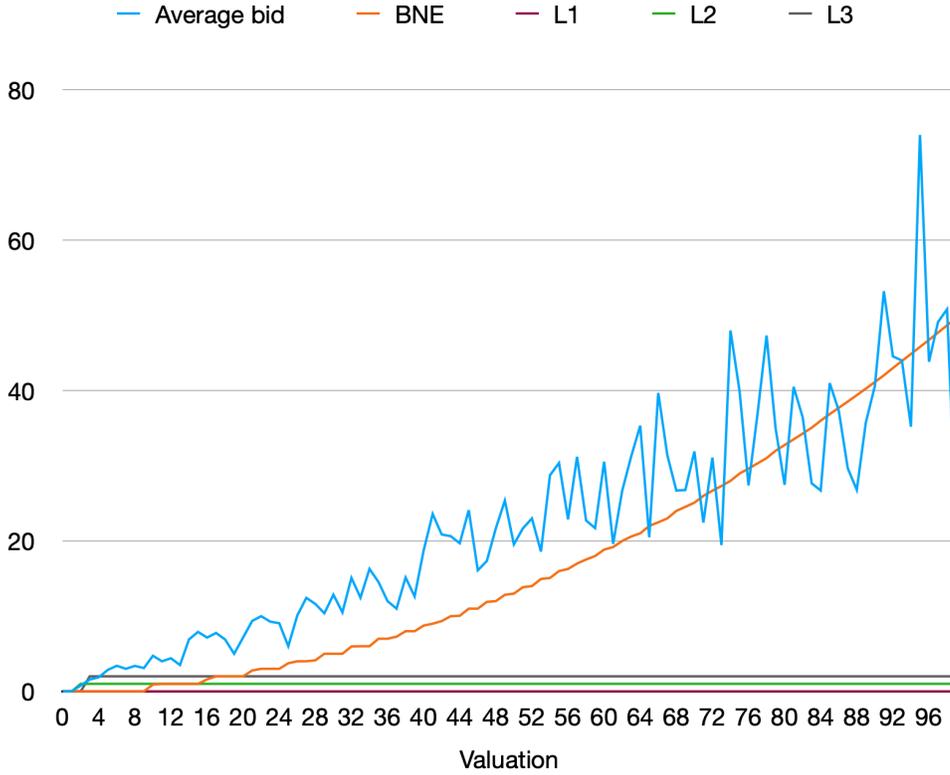


Notes: This figure plots the average observed bid for every possible valuation, along with the associated equilibrium and level- k predictions.

We then conduct a similar analysis of the all-pay auction. As shown in Figure 6, equilibrium does a somewhat better job of tracking mean bids in this environment (see also Figure 9 for the case where bids are multiples of 5). However, over-bidding relative to equilibrium is again observed on average at low to medium values.²⁰ As before, the level- k model predicts bids that are orders of magnitude below those that are observed on average, at least if one imposes the restriction that levels cannot exceed 3. Thus, our initial examination of the data suggests a very clear ranking of the two models under consideration.

²⁰The overbidding that we observe is consistent with results from the only other experimental study of IPV all-pay auctions of which we are aware (Noussair and Silver, 2006). Unfortunately, however, we were unable to obtain the data from this experiment and so were unable to use it to complement our analyses.

Figure 6: Predictions and data in the all-pay auction



Notes: This figure plots the average observed bid for every possible valuation, along with the associated equilibrium and level- k predictions.

To verify these conclusions more formally, we also compute the root-mean-square prediction errors of both of the theories. In the case of equilibrium, this is straightforward since the model is ‘parameter free’. In the case of level- k , we obtain a prediction by assigning each datapoint the level from the 1 – 3 range that minimises the model’s prediction error (this is a generous assumption that will tend to overstate the model’s predictive performance). Table 1 displays the results. As can be seen, both models make size-able prediction errors. Consistent with the graphical evidence, however, the prediction errors of level- k are substantially greater than those of equilibrium.

Table 1: Prediction errors

	T1 FP	T1 AP	T2 FP	T2 AP
Equilibrium	15.1	13.3	14.1	14.3
Level- k	21.3	19.8	21.4	24.0

Notes: T1 and T2 denote the integer bid and multiples of five treatments; FP and AP abbreviate ‘first-price’ and ‘all-pay’.

Next, we examine the effect of varying the bid discretisation. Recall that the level- k model predicts a large effect from changing the discretisation from integers to multiples of 5: bids should increase by roughly a factor of 5. In contrast, the equilibrium model predicts much smaller effects. Given these very different predictions, we now conduct a comparison of the integer bid and multiples of 5 bid treatments.

As a preliminary, we confirm that the treatments are balanced, suggesting that the randomisation worked as expected (see Table 3). Next, we inspect the average bids across the two treatments and auction structures (see Table 4). As can be seen, changing the bid discretisation makes essentially no difference to the first-price auction (in both cases, the average bid is about 46% of the average valuation). In the all-pay auction, restricting bids to multiples of five does appear to slightly increase the average bids (which rise from 45% to 50% of the average valuation), but clearly not in the dramatic way predicted by the level- k model. In neither auction, then, do the data support the surprising comparative static prediction of level- k theory.

4.2 Structural estimates

In light of the evidence above, it is clear that the level- k model cannot count for observed average behaviour, at least when calibrated with a plausible distribution of types. However, this still leaves at least two issues unresolved. First, even if level- k cannot describe the average bidder, it is still possible that it accounts for a substantial minority of the observed behaviour. Second, it remains to be seen whether the model might be rescued by allowing for a greater number of possible types. We now address both of these issues.

To do this, we follow the literature in estimating a mixture of types model (see, e.g., Stahl and Wilson (1994), Bosch-Domenech et al. (2002) and Crawford and Iriberri (2007)). To obtain stochastic versions of the relevant theories, we assume that subjects submit bids with probabilities that are proportional to a normal distribution whose mean is the predicted bid that they are ‘supposed’ to submit (according to the relevant theory). More formally, let $\mathbb{K} \subset \mathbb{N}^+$ denote the set of types, indexing equilibrium by $k = 0$ and level- k types in the obvious way. Given that a player i has value $v_i \in \mathbb{X}$ and type $k \in \mathbb{K}$, we assume that the chance they submit a bid $b_i \in \mathbb{X}$ is given by:

$$\mathbb{P}(b_i|v_i, k) \propto \exp \left[-\frac{1}{2} \left(\frac{b_i - b_p}{\sigma_k} \right)^2 \right] \quad (9)$$

where b_p is the bid predicted by their type (given their valuation), σ_k is a noise parameter

that controls the variance of their bids, and the constant of proportionality is chosen to ensure that the probabilities sum to 1.²¹ This produces a single peaked distribution whose mode is the predicted bid b_p . As $\sigma_k \rightarrow \infty$, this distribution converges to uniform noise; and as $\sigma_k \rightarrow 0$, it converges to a degenerate distribution which places all probability mass on the predicted bid b_p .

For now, assume that each individual's type is fixed (though we relax this in many of our analyses). Then the probability that an individual submits a vector of (twenty) bids \mathbf{b}_i given their vector of values \mathbf{v}_i is (assuming independence of errors):

$$\mathbb{P}(\mathbf{b}_i|\mathbf{v}_i, k) = \prod_{i=1}^{20} \mathbb{P}(b_i|v_i, k) \quad (10)$$

We assume that the individual we observe is drawn randomly from the population; and write $p_k \in [0, 1]$ to denote the population proportion of type k (obviously, $\sum_k p_k = 1$). By the law of total probability, the unconditional probability of their bids is

$$\mathbb{P}(\mathbf{b}_i|\mathbf{v}_i) = \sum_{k \in \mathbb{K}} \mathbb{P}(\mathbf{b}_i|\mathbf{v}_i, k) p_k \quad (11)$$

and so the ‘likelihood’ of all the observed decisions is (combining Equations 9 - 11)

$$L = \prod_{i=1}^n \mathbb{P}(\mathbf{b}_i|\mathbf{v}_i) \propto \prod_{i=1}^n \left(\sum_{k \in \mathbb{K}} \left[\prod_{i=1}^{20} \left(\exp \left[-\frac{1}{2} \left(\frac{b_i - b_p}{\sigma_k} \right)^2 \right] \right) \right] p_k \right) \quad (12)$$

In every estimation exercise, we compute the parameters $\mathbf{p} = (p_k)_{k \in \mathbb{K}}$ and $\boldsymbol{\sigma} = (\sigma_k)_{k \in \mathbb{K}}$ that maximise the likelihood (12). We then calculate standard errors with the jack-knife. While we consider the case where individual levels are fixed (as above), we also consider the case where individual levels are permitted to vary by round – although we do insist throughout that levels are fixed within a round.²²

As a preliminary exercise, we compare an equilibrium only model with the level- k model in which levels are restricted from 1 – 3. To avoid introducing too great an imbalance in the number of parameters between the two respective models, we restrict both to a single noise parameter σ . Table 5 displays the results. As one might have expected from Figures 5 - 6,

²¹We also experimented with a model containing logit errors. However, these had the unfortunate tendency of preferring level-1 to all higher levels even though level-1 predictions were the furthest from the observed bids (recall that level-1 types bid zero regardless of their valuation).

²²Recall that all bids within a round must be made simultaneously, so there is no possibility for learning within rounds. In contrast, individuals could conceivably learn between rounds – although the scope for such learning is limited by the near-total absence of feedback.

equilibrium does a better job than any of the level- k models at explaining the data, attaining a higher log-likelihood in all specifications. This advantage becomes even greater if we punish the level- k models for their extra flexibility; for example, using the Bayesian information criterion. This simply reinforces the conclusion that we reached earlier: equilibrium does a better job than the level- k model at accounting for the overall patterns in the data.

We now examine the extent to which the level- k model might be salvaged by adding higher levels. To this end, we now estimate a series of models, gradually adding levels until we reach the point where the model cycles. Three main points emerge from this exercise.²³ First, while the level- k model can achieve a higher likelihood than equilibrium once the full suite of levels is added, it can do this only by placing most probability on very high levels: for example, in the first price auction, most probability mass is placed on level-32. As argued previously, it is doubtful whether we can plausibly ascribe so many thinking steps to participants. Second, the level- k model fits the data with a very unusual distribution, placing some probability mass on level-1, none on intermediate levels, and the rest on the very high levels just discussed. Of course, the idea that individuals do either one thinking step or 32 thinking steps might appear to be rather strange. Finally, the inclusion of these very high levels is almost never recommended by the Bayes-Information Criterion. For all these reasons, it seems difficult to rescue the level- k model through a strategy of adding a very large number of levels.

Although a pure level- k model appears to be outperformed by equilibrium, a final question we investigate is whether a hybrid model might be able to outperform both. To this end, we now re-estimate the model for all auction structures, allowing for both equilibrium and level- k types and no longer restricting the model to a single noise parameter σ . We also return to our previous assumption that levels are in the plausible range (1-3). Table 6 displays the results. As can be seen, most probability mass is placed on equilibrium – suggesting again that the level- k model fits the data relatively poorly. However, there is still some probability mass placed on level- k types. Thus, the results from the hybrid model are consistent with the possibility that the level- k model characterises the behaviour of a small minority of subjects, even if it struggles to capture the general patterns in the bidding data.

4.3 Correlating levels

We now examine whether individual levels as estimated from the two auctions bear any resemblance to the levels as inferred from a game known to trigger level- k reasoning: a variant

²³Since the resulting tables are too large to be contained within the appendix, they have been made available [here](#).

on the 11-20 game proposed by Alaoui and Penta (2016). To this end, we calculate the level k and precision σ_k that maximises the likelihood of each individual’s bids (i.e. Equation 10). Thus, we now assign each individual a definite level, as opposed to a distribution of levels; and no longer allow for equilibrium types. After doing this for both the first-price and all-pay auction, we then compute each individual’s level as inferred from the 11-20 game, discarding those whose implied level exceeds 4.²⁴ Finally, we compute the correlation between the levels inferred from each auction and the level as inferred from the 11-20 game.

Figure 10 plots the levels estimated from the first-price and all-pay auctions against the levels as inferred from the 11-20 game. As can be seen, there is no discernible relationship between the levels; and indeed the correlations are estimated to be 0.00 and -0.09 in the case of the first-price and all-pay auctions respectively.²⁵ From the perspective of level- k theory, this is rather unexpected. For if levels reflect either cognitive sophistication or beliefs about the cognitive sophistication of one’s opponents – the two leading interpretations of individual levels – then one would expect to find a positive correlation between individual levels across games. Conversely, the absence of such a positive correlation suggests that individuals are not using level- k reasoning in at least one of the games.

4.4 Subject reports

Finally, we examine the explanations given by participants as to why they chose to bid in the way that they did. In particular, we consider whether they appealed to iterated reasoning.²⁶ Assuming that subjects’ conscious reasoning was indeed captured by the level- k model and further assuming that subjects did not for some reason wish to lie to their experimenter about their reasoning, one would expect to see iterated reasoning cited in their explanations. Conversely, one can take the absence of such reports as evidence that the level- k model did not in fact characterise their conscious reasoning.

In their explanations, the majority of subjects (around 55%) explicitly indicated an awareness

²⁴We do this partly because such levels are at odds with prior experimental evidence but also because not one of such subjects invoke level- k reasoning when asked to explain their choice.

²⁵This result reflects those of Georganas et al. (2015), who find no correlation between estimated individual levels across two families of games (albeit in a non-auction setting).

²⁶A strict definition of iterated reasoning would be reports of the form: “I bid in way X because I expected my opponent to bid in way Y; and I expected my opponent to bid in way Y because I thought they would think I would bid in way Z.” In other words, this would need to involve (1) a conjecture about the bidding strategies followed by the player’s opponents (2) an attempt to ground this conjecture using a conjecture about the conjectures of the opponents about the player’s bidding strategy. On a weaker definition, we would simply require some kind of conjecture about the opponents’ bidding strategy (along with an attempt to optimise given this conjecture). Since the level- k model allows for level-1 players, we used this weaker (and more permissive) definition when categorising the reports.

of the basic bidding trade-off: bidding more increases one’s probability of winning, but decreases one’s pay-off if one wins the auction. However, the vast majority of subjects did not appear to use iterated reasoning. Indeed, we categorise responses as invoking iterated reasoning in just two cases.²⁷

To be clear, we are not claiming that the level- k model did not characterise the reasoning given by any of the subjects. For example, one subject (who also bid in line with the model) wrote that

Since my opponent has a 50% chance of their bid being cancelled, it’s better to keep bids low in hopes theirs [sic] gets cancelled . . . I raised the bid to more than 1 because in the event my opponent is using the same strategy and bids very low, I still have a chance of winning the auction.

Thus, this subject not only considered something like the level-1 strategy (bidding ‘very low’) but then went beyond it – a clear case of level- k reasoning.²⁸ What is evident, however, is that reasoning was provided by only a small minority of subjects, a finding that is also in line with the choice data outlined in the previous section.

5 Robustness checks

We now report on the results of a variety of extensions and robustness checks.

Risk aversion. While the previous analysis assumed that individuals are risk-neutral, we now consider whether the level- k model might be salvaged by allowing individuals to be risk-averse.²⁹ To begin, consider the all-pay auction. As noted in Section 2.1, level-1 players (if assumed to be risk-neutral) are predicted to bid 0 (for all values), thereby earning a certain payoff of zero. If they submitted a positive bid, then they would be exposed to risk: either their pay-off would be positive (if they win the auction) or negative (if they lose the auction). Now, given that level-1 players choose to bid 0 when they are risk neutral, it is obvious that they must also choose to bid 0 if they are assumed to be risk averse. Thus, introducing risk aversion makes no difference to the model’s predictions – and also (for this reason) makes

²⁷All categorisations were independently checked by another researcher who was not involved in the study. The reader is invited to examine the responses themselves which are provided (along with all other non-demographic data generated by the experiment) in the supplementary materials.

²⁸While the bids submitted by the subjects did not exactly conform to any of the levels they were very close, with the subjects bidding a mix of 2’s and 3’s at relatively low valuations.

²⁹Of course, expected utility maximisers should be approximately risk-neutral over such small stakes (Arrow, 1970; Rabin, 2000). Thus, the ‘risk aversion’ discussed here should really be thought of as capturing utility curvature within a reference dependent model (Bleichrodt et al., 2019).

no difference to the predictions for high levels (provided that they are sufficiently low to be characterised by bidding strategy β_{LK}).

While risk aversion cannot improve the model’s performance in the all-pay auction, matters are more complicated in the first-price auction. In general, holding the bidding strategies of one’s opponents fixed, introducing risk aversion must (weakly) increase one’s optimal bid. In our setting, this means that introducing risk aversion can make positive bidding by a level-1 type optimal, thereby bringing the model’s predictions closer in line with the observed data. Note, however, that introducing risk aversion *also* makes equilibrium bids more aggressive (Holt, 1980), which might be expected to increase the predicted accuracy of our benchmark. It is therefore unclear at the onset whether introducing risk aversion improves the relative performance of the level- k model.

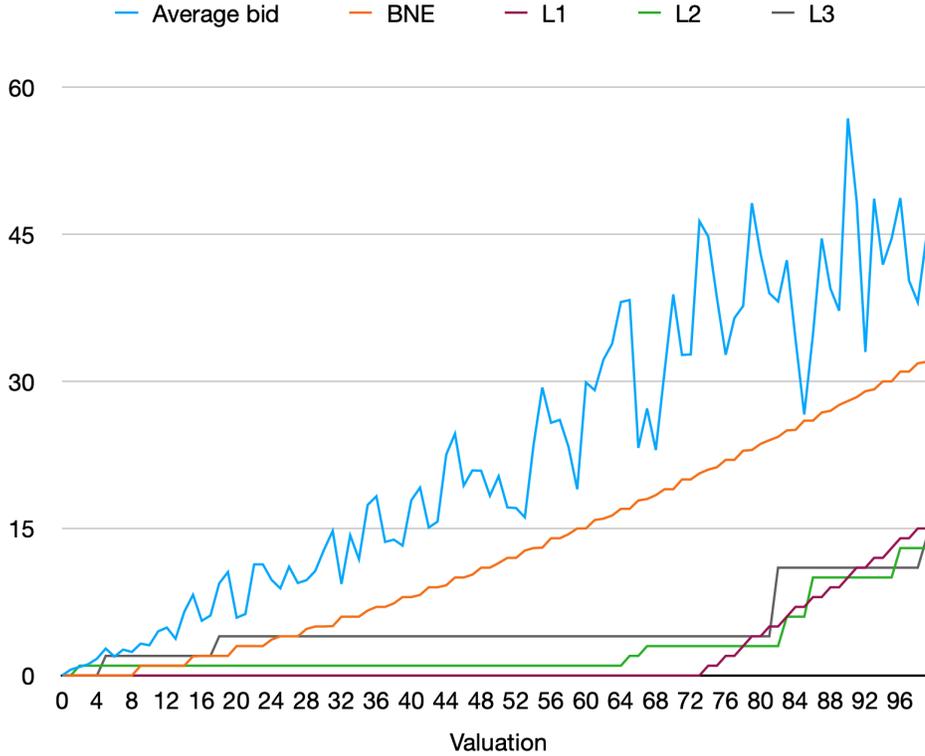
To study this issue more carefully, we assume that every individual i has constant relative risk aversion (i.e. $u_i(x) = x_i^\alpha$ where x_i is the individual’s pay-off and $\alpha_i > 0$ determines their level of risk aversion.) We then estimate every individual’s parameter α_i using their choice in the “bomb” risk elicitation task (Crosetto and Filippin, 2013). We then re-calculate the level- k model’s predictions for every player using their estimated parameter α_i (along with their estimated level as in Section 4.1). We also re-calculate the equilibrium predictions on the assumption that every individual has the average estimated risk aversion parameter of the subject pool ($\alpha = 0.72$).³⁰

Figure 7 displays the results for the main treatment (see also 11 for the case where bids are multiples of five). As can be seen, both the level- k and equilibrium predictions are now somewhat improved and considerably closer to the actual data. Overall, however, it is clear that equilibrium still outperforms level- k dramatically – a conclusion confirmed by the prediction error rates reported in Table 7.

Cognitive hierarchy. We now ask if any conclusions might change if individuals are assumed to best respond to a *distribution* over levels, as in Camerer et al. (2004)’s cognitive hierarchy model. To this end, we assume (following Camerer et al. (2004)) that levels are Poisson distributed with mean τ . Using standard maximum likelihood methods, we then estimate τ for the participants in our experiments using their choices in the 11-20 game, finding that $\tau \approx 1.9$. Finally, we compute the predictive errors of the cognitive hierarchy model and compare those with the associated errors for equilibrium. As shown in Table 7,

³⁰The code can be viewed here: https://github.com/Itzhak95/risk_aversion/tree/master. The $\alpha = 0.72$ is obtained after dropping the two outliers who collected 61 or more out of the 62 boxes. If such subjects are included, we obtain $\alpha = 1.8$ (risk-seeking), in which case allowing non-linear utility is unable to improve the performance of the level- k model.

Figure 7: Risk aversion in the first-price auction



Notes: This figure plots the average observed bid for every possible valuation, along with the associated equilibrium and level- k predictions. Equilibrium predictions are derived under the assumption that each individual has constant relative risk aversion with $\alpha = 0.72$. Level- k predictions are derived by fitting each individual with the α_i parameter that rationalises their choice in the bomb risk elicitation task.

moving from level- k to cognitive hierarchy makes essentially no difference to the prediction errors. This is not a surprise: obviously, best responding to a mix of levels between 1 and 3 yields (in our context) very similar optimal bids to best responding to level-2.

Including dominated bids. As noted earlier, our main analysis excluded the small number of subjects who submitted at least one dominated bid on the grounds that such subjects may have misunderstood the rules of the auction. However, excluding such subjects is a debatable choice; and we re-ran all analyses after including such subjects. This did not change any of our substantive conclusions (see Table 7).

Changing the tie-breaking rule. In the previous analysis, we assumed that level- k players always chose the lowest optimal bid if multiple bids were optimal. While this allowed the model to avoid predicting that individuals choose dominated strategies, it is natural to ask whether any of our conclusions hinge on this assumption. To investigate this, we now make the opposite assumption – that level- k players always choose the highest bid out of the

bids that they believe are optimal. Re-running our main analyses, we see that this slightly worsens the model’s fit with the data (although the change is small). This is not a surprise: the main change is that the model now predicts some dominated bids, and such bids are very rare in the data.

Dropping the second round. As noted in Section 3, the level- k model is supposed to be a model of initial play. In order to test the model in its natural domain, we therefore blunted subject feedback and restricted the number of rounds to two. Nonetheless, one might worry that even two rounds is too many to count as ‘initial play’. To test this possibility, we re-ran the analyses after dropping the second round. While this does somewhat improve level- k ’s performance, the model remains clearly outperformed by equilibrium (see Table 5).

Changing the L_0 specification. The level- k auction model, as formulated by Crawford and Iriberri (2007), assumes that level-0 players either bid their value (the ‘truthful’ specification) or randomise uniformly over \mathbb{X} . Since the predictions of such a model seem strongly disconfirmed by the data, one might wonder whether the model might be improved by choosing an alternative level-0 specification. While our data do not allow us to absolutely rule this possibility out, there are various reasons to think that this is not likely to be a fruitful theoretical exercise.

First, as stressed by Crawford et al. (2013), one must be careful to ensure that any L0 specification represents a ‘strategically naive’ assessment of how others will play the game. This is most obviously important if one wishes to connect the model to its standard interpretation as a model of iterated reasoning given some naive but initially plausible starting point. In addition, it is vital to restrict the range of possible level-0 strategies if one wishes the model to make anything like definite predictions in strategic situations; and to be clearly disentangled from rival models. For instance, if an analyst were allowed to specify that level-0 play coincided with the symmetric equilibrium, it would be impossible to disentangle the equilibrium from the predictions of the level- k model: for every $k \geq 1$, the level- k prediction would simply be the symmetric equilibrium strategy.

Thus, if one did want to change the L_0 specification, one would need to make sure that it represents a psychologically plausible way in which a player might naively expect others to play. This brings us to our second point, however: there simply do not appear to be a large number of such L_0 specifications. Other than the specifications proposed by Crawford and Iriberri (2007), we are only able to think of one alternative: bidding uniformly between 0 and one’s value v (modifying their ‘random’ specification so as to remove dominated bids). However, opting for this L0 specification would render the model unable to match even the

general trend of bidding in the most experimentally studied of all auction formats – the first-price auction with independent, uniform values (and no possibility of cancelled bids).³¹

Finally, we note that the level- k predictions as formulated in Section 2 are robust to a very wide range of level-0 specifications. For example, in the all-pay auction, level- k bids will be characterised by Proposition 2 for *any* L_0 specification that satisfies

$$F(b) \leq \left(\frac{b}{x+1} \right)^{\frac{1}{n-1}} \quad (13)$$

where $F(b)$ is the probability that the level-0 type bids b or lower (i.e. the CDF of b). In particular, this means that Proposition (5) will continue to hold provided that the CDF $F(b)$ that is bounded by the uniform CDF from above. Similarly, in the first-price auction, level-1 players will bid 0 (as in Proposition 5) for any level-0 specification that satisfies

$$F(b) \leq \left(\frac{n-1}{n} \right) \left(\left(\frac{x}{x-b} \right)^{n-1} - 1 \right) \quad (14)$$

for all $b \in \mathbb{X}$. Moreover, while these inequalities are sufficient for our level- k characterisation to hold, they are absolutely not necessary. For example, in the all-pay auction, the level-1 bidding strategy turns out to be a best response to the empirical bid distribution observed in the experiment, even though that distribution does not satisfy inequality (13).

In light of the reasons above, we think that it is unlikely that one can find a level-0 specification that is simultaneously non-strategic, psychologically plausible and generates predictions that match the data in both this and other auction experiments. Ultimately though, this conclusion is somewhat more speculative than the central conclusion of this study: that the level- k model as specified by Crawford and Iriberry (2007) cannot predict or rationalise the observed bidding behaviour.

³¹In this setting, a large number of auction experiments report overbidding relative to the risk-neutral equilibrium $\beta(v) = (n-1)v/n$ (see Kagel and Levin (1995) for a review). However, it is easy to see the level-0 specification suggested here would predict bids that are substantially lower than the risk-neutral equilibrium. As a result, it would result in level- k predictions that are far too low (at least, given a plausible calibration of the levels).

Note also that other apparently plausible L_0 specifications, such as thinking that one's opponents bid half their value, are excluded on the grounds that L_0 reasoning is supposed to be non-strategic (whereas bidding a fraction of one's value quite clearly reflects an attempt to strategically balance one's chance of winning with one's payoff if one does win). Moreover, such specifications imply that level- k bidders should submit a large mass of bids just above the maximum value bid by L_0 (for instance L_1 players in the all-pay auction should bid 50 for all $v \geq 51$). We do not observe such masses in this experiment nor in the other auctions experiments of which we are aware.

6 Concluding remarks

In this paper, we have designed and implemented an experiment aimed at testing the level- k model of auctions. Overall, the evidence would appear to strongly reject the model. When plausibly calibrated, the model produces bids that are at least an order of magnitude too low. Moreover, fitting the model without such a plausibility constraint results in estimated levels in the 30–35 range; and these levels in turn fail to bear any relation to levels inferred from the 11-20 game. In addition, changing the bid discretisation has little effect on the observed bids (in contrast to the model’s predictions); and subjects almost never cite iterated reasoning when asked to explain why they bid in the way they did. Thus, while one might be able to disagree with any of the pieces of evidence separately, it seems that the evidence collectively casts strong doubt on the model’s ability to predict or even rationalise the observed data.

This finding might be seen as surprising given that the experimental environment was carefully chosen to be the level- k model’s natural domain of applicability. Individual feedback was severely restricted and (following Crawford and Iriberri (2007)) the number of rounds was restricted to just two. However, the level- k model fared poorly despite the initial play setting.

The findings also might be seen as surprising in light of evidence that the level- k model very accurately captures facets of human behaviour in other settings, whether ‘beauty contests’ (Nagel, 1995), co-ordination games (Costa-Gomes et al., 2009), strategic communication (Crawford, 2003; Cai and Wang, 2006) or zero sum betting (Brocas et al., 2014). We thus close with some speculations as to why the model appears to do so badly at predicting bids in auctions despite its successes in predicting behaviour in other areas.

First, recall that the level- k model requires a level-0 anchor, i.e. some sort of naive starting point that forms the basis of the iterated reasoning. In many settings, such an anchor is easy to provide. For example, one’s natural first thought (in a strategic communication setting) is that one’s opponents will tell the truth; and a natural first thought (in the context of guessing games) is that all possible guesses are equally likely. In contrast, it is not at all clear how the level-0 ought to be specified in the setting of auctions. Crawford and Iriberri (2007) select what appear to be the two best proposals — that subjects either bid their valuation or randomise uniformly over the strategy space. However, it is not obvious whether either represents a natural first thought for how one’s opponents will play. More to the point, it does not seem that there exists more plausible L0 specifications than those proposed by Crawford and Iriberri (2007) (see Section 5 for discussion). Thus, the failure of the level- k model in this context may well be down to the lack of a salient and psychologically plausible

anchor.

Second, iterated reasoning is cognitively much more taxing in Bayesian games like auctions than the simpler settings to which the level- k model is normally applied. To calculate one's optimal action, one must combine one's conjecture about one's opponent's strategy with their distribution over types to arrive at their distribution over actions. One must then solve an (often non-trivial) optimisation problem to find the optimal action given this distribution. Thus, even if a plausible level 0 anchor were available, it is unclear whether individuals would be able to correctly solve for the L1 strategy – and even less clear whether they would be able to correctly identify the higher order strategies.

If this is right, it follows that the level- k model should be applied only in situations which possess both an intuitively appealing 'first thought' along with easily computable best response dynamics. Rigorously testing whether these features are indeed necessary for the predictive success of the model — along with identifying strategic settings which possess these apparently necessary features — would seem to be an important task for future research.

References

- Alaoui, L. and Penta, A. (2016). Endogenous depth of reasoning. *The Review of Economic Studies*, 83(4):1297–1333.
- An, Y. (2017). Identification of first-price auctions with non-equilibrium beliefs: A measurement error approach. *Journal of econometrics*, 200(2):326–343.
- Arad, A. and Rubinstein, A. (2012). The 11-20 money request game: A level- k reasoning study. *American Economic Review*, 102(7):3561–73.
- Arrow, K. J. (1970). Essays in the theory of risk-bearing. Technical report.
- Bleichrodt, H., Doctor, J. N., Gao, Y., Li, C., Meeker, D., and Wakker, P. P. (2019). Resolving rabin’s paradox. *Journal of Risk and Uncertainty*, 59(3):239–260.
- Bosch-Domenech, A., Montalvo, J. G., Nagel, R., and Satorra, A. (2002). One, two,(three), infinity,...: Newspaper and lab beauty-contest experiments. *American Economic Review*, 92(5):1687–1701.
- Brocas, I., Carrillo, J. D., Wang, S. W., and Camerer, C. F. (2014). Imperfect choice or imperfect attention? understanding strategic thinking in private information games. *Review of Economic Studies*, 81(3):944–970.
- Cai, H. and Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.
- Cheng, S.-F., Reeves, D. M., Vorobeychik, Y., and Wellman, M. P. (2004). Notes on equilibria in symmetric games.
- Costa-Gomes, M. A., Crawford, V. P., and Iriberri, N. (2009). Comparing models of strategic thinking in van huyck, battalio, and beil’s coordination games. *Journal of the European Economic Association*, 7(2-3):365–376.
- Costa-Gomes, M. A. and Shimoji, M. (2015). A comment on “can relaxation of beliefs rationalize the winner’s curse?: An experimental study”. *Econometrica*, 83(1):375–383.

- Crawford, V. P. (2003). Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *American Economic Review*, 93(1):133–149.
- Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62.
- Crawford, V. P. and Iriberri, N. (2007). Level- k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770.
- Crawford, V. P., Kugler, T., Neeman, Z., and Pauzner, A. (2009). Behaviorally optimal auction design: Examples and observations. *Journal of the European Economic Association*, 7(2-3):377–387.
- Crosetto, P. and Filippin, A. (2013). The “bomb” risk elicitation task. *Journal of Risk and Uncertainty*, 47(1):31–65.
- De Clippel, G., Saran, R., and Serrano, R. (2019). Level- k mechanism design. *The Review of Economic Studies*, 86(3):1207–1227.
- Filiz-Ozbay, E. and Ozbay, E. Y. (2007). Auctions with anticipated regret: Theory and experiment. *American Economic Review*, 97(4):1407–1418.
- Galavotti, S., Moretti, L., and Valbonesi, P. (2018). Sophisticated bidders in beauty-contest auctions. *American Economic Journal: Microeconomics*, 10(4):1–26.
- Georganas, S., Healy, P. J., and Weber, R. A. (2015). On the persistence of strategic sophistication. *Journal of Economic Theory*, 159:369–400.
- Gillen, B. (2009). Identification and estimation of level- k auctions. *Available at SSRN 1337843*.
- Harsanyi, J. C. (1967). Games with incomplete information played by “bayesian” players, i–iii part i. the basic model. *Management Science*, 14(3):159–182.
- Harstad, R. M., Kagel, J. H., and Levin, D. (1990). Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters*, 33(1):35–40.
- Holt, C. A. (1980). Competitive bidding for contracts under alternative auction procedures. *Journal of political Economy*, 88(3):433–445.

- Hortaçsu, A., Luco, F., Puller, S. L., and Zhu, D. (2019). Does strategic ability affect efficiency? evidence from electricity markets. *American Economic Review*, 109(12):4302–42.
- Ivanov, A., Levin, D., and Niederle, M. (2010). Can relaxation of beliefs rationalize the winner’s curse?: An experimental study. *Econometrica*, 78(4):1435–1452.
- Kagel, J. and Levin, D. (1995). Auctions: A survey of experimental research,” handbook of experimental economics, edited by j. kagel and a. roth.
- Kirchkamp, O. and Reiß, J. P. (2011). Out-of-equilibrium bids in first-price auctions: Wrong expectations or wrong bids. *The Economic Journal*, 121(557):1361–1397.
- Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of economic surveys*, 13(3):227–286.
- Krishna, V. (2009). *Auction theory*. Academic press.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326.
- Noussair, C. and Silver, J. (2006). Behavior in all-pay auctions with incomplete information. *Games and Economic Behavior*, 55(1):189–206.
- Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292.
- Rasooly, I. and Gavidia-Calderon, C. (2020). The importance of being discrete: on the (in-) accuracy of continuous approximations in auction theory. *arXiv preprint arXiv:2006.03016*.
- Stahl, D. O. and Wilson, P. W. (1994). Experimental evidence on players’ models of other players. *Journal of economic behavior & organization*, 25(3):309–327.

Appendices

A Proofs

Proof of Proposition 1. Since our game has a symmetric and finite normal form, it must have a symmetric equilibrium (see Harsanyi (1967) and Cheng et al. (2004) for details). Turning to uniqueness, we now generalise three lemmas from Rasooly and Gavidia-Calderon (2020) to the case of mixed strategies. Let σ denote an arbitrary symmetric equilibrium strategy; and for any $v \in \mathbb{X}$, let $s(v) \subseteq \mathbb{X}$ denote the set of bids that are submitted with positive probability at value v (i.e. the support at v).

Lemma 1 *In any symmetric equilibrium σ , $s(0) = \{0\}$.*

Proof. Consider a player with $v = 0$. If they bid 0, they get (regardless of their opponents' strategies) a payoff $\pi(v = 0, b = 0) = 0 \times \mathbb{P}(\text{win}|b = 0) - 0 = 0$. On the other hand, if they submit any bid $b \geq 1$, they get (given any opponent strategy profile) $\pi(v = 0, b) = 0 \times \mathbb{P}(\text{win}|b) - b = -b < 0$. Thus, bidding 0 strictly dominates bidding any $b \geq 1$; which means that $\mathbb{P}(b = 0|v = 0) = 1$ in any equilibrium (and in any symmetric equilibrium). \square

Lemma 2 (Monotonicity) *Consider any symmetric equilibrium σ and consider any two values $v, v' \in \mathbb{X}$ with $v > v'$. Then for any $b \in s(v)$ and $b' \in s(v')$, $b \geq b'$.*

Proof. For contradiction, let us suppose that $b < b'$. Since b is submitted with positive probability when the value is v , bidding b must be weakly optimal. In particular, it must be weakly better than bidding b' :

$$\begin{aligned} \pi(v, b) &\geq \pi(v, b') \\ \iff v\mathbb{P}(\text{win}|b) - b &\geq v\mathbb{P}(\text{win}|b') - b' \\ \iff b' - b &\geq v(\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)). \end{aligned} \tag{15}$$

Similarly, since b' is submitted with positive probability when the value is v' , it must be weakly better than b :

$$\begin{aligned} \pi(v', b') &\geq \pi(v', b) \\ \iff v'\mathbb{P}(\text{win}|b') - b' &\geq v'\mathbb{P}(\text{win}|b) - b \\ \iff v'(\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)) &\geq b' - b. \end{aligned} \tag{16}$$

Inequalities (15) and (16) jointly imply that

$$v' (\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)) \geq v (\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)). \quad (17)$$

Since $b' > b$, and b is a symmetric equilibrium bid, it must be that $\mathbb{P}(\text{win}|b') > \mathbb{P}(\text{win}|b)$ (for example, one wins with a bid of b' , but not with a bid of b , if all of one's opponents bid b – and this happens with positive probability). Since $\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b) > 0$, inequality 17 then yields $v' \geq v$, which contradicts our initial assumption that $v' < v$. This establishes that $b \geq b'$ as claimed. \square

Lemma 3 (No gaps) *In any symmetric equilibrium σ , the bids that are submitted with positive probability are a set of consecutive integers.*

Proof. Suppose for contradiction that this were false. Then there must exist two bids $b_{high} \geq b_{low} + 2$ that are submitted with positive probability even though no bid in between b_{high} and b_{low} is submitted with positive probability. Now if types were to deviate from bidding b_{high} to bidding $b_{low} + 1$, their payment would fall. However, their probability of winning would remain unchanged: both before and after the deviation, they win the auction if and only if all opponents bid b_{low} or lower. So this deviation would be strictly profitable, implying that this could not have been an SE. \square

We now introduce the concept of jump form.

Definition A.1 *For every bid $i \in B^\sigma$, we define the i th jump j_i as*

$$j_i = v_i + \mathbb{P}(b < i | v_i)$$

where v_i is the minimum value $v \in \mathbb{X}$ such that $\mathbb{P}(b = i | v) > 0$. We refer to the vector of jumps $j = (j_0, j_1, \dots, j_m)$ as a jump vector and say that it is increasing if $j_{i+1} > j_i$ for every $i \in \{0, 1, \dots, m-1\}$.

Under the lemmas outlined previously, this is an equivalent representation of behavioural strategies.

Lemma 4 *There is a bijection between the set of gapless and monotone strategies that satisfy $s(0) = 0$ and the set of increasing jump vectors.*

Proof. To establish the first claim, take any gapless and monotone strategy that specifies $\mathbb{P}(b = 0 | v = 0) = 1$ and convert it into a jump vector j . To show that j is increasing, consider any two consecutive jump points $j_i, j_{i+1} \in j$. If the strategy is monotone, every value that submits a bid $b = i + 1$ must be weakly larger than the every value that submits a bid $b = i$.

So in particular, $v_{i+1} \geq v_i$. This yields two possibilities: either $v_{i+1} > v_i$ or $v_{i+1} = v_i$. In the first instance ($v_{i+1} > v_i$),

$$\begin{aligned}
j_{i+1} &\equiv v_{i+1} + \mathbb{P}(b < i + 1 | v_{i+1}) \\
&\geq v_i + 1 + \mathbb{P}(b < i + 1 | v_{i+1}) \\
&\geq v_i + 1 \\
&> v_i + \mathbb{P}(b < i | v_i) \equiv j_i
\end{aligned} \tag{18}$$

where the final inequality holds since $\mathbb{P}(b < i | v_i) < 1$ (which in turn holds since necessarily $\mathbb{P}(b = i | v_i) > 0$). So in that case, $j_{i+1} > j_i$. In the second instance ($v_{i+1} = v_i$), we have

$$\begin{aligned}
j_{i+1} &\equiv v_{i+1} + \mathbb{P}(b < i + 1 | v_{i+1}) \\
&= v_i + \mathbb{P}(b < i + 1 | v_i) \\
&= v_i + \mathbb{P}(b < i | v_i) + \mathbb{P}(b = i | v_i) \\
&> v_i + \mathbb{P}(b < i | v_i) \equiv j_i
\end{aligned} \tag{19}$$

where the inequality again holds since $\mathbb{P}(b = i | v_i) > 0$. Either way, then, we see that monotonicity implies that $j_{i+1} > j_i$. Since this true for all $i = 2, 3, \dots, m$, this then implies that that vector of jumps is increasing. So every monotone behavioural strategy generates an increasing vector of jumps.

We now argue that any increasing vector of jumps $j = (j_1, \dots, j_m)$ is consistent with a unique monotone and gapless behavioural strategy that satisfies $\mathbb{P}(b = 0 | v = 0)$. First, use j to compute the minimum values that submit each of the possible bids v_i using the fact that $v_i = \lfloor j_i \rfloor$ for all $i = 1, \dots, m$. (Notice that, since $j_{i+1} > j_i$ for every i , $v_{i+1} = \lfloor j_{i+1} \rfloor \geq v_i = \lfloor j_i \rfloor$, i.e. the v_i are weakly increasing.) We will begin by showing how this uniquely determines the behavioural strategy for all values other than the v_i .

First, consider all $v \in \mathbb{X}$ with $v < v_1$. Recall that v_1 is the minimum $v \in \mathbb{X}$ such that $\mathbb{P}(b = 1 | v) > 0$. So $\mathbb{P}(b = 1 | v_1) > 0$; and $\mathbb{P}(b = 1 | v) = 0$ for any value $v < v_1$. Since $\mathbb{P}(b = 1 | v_1) > 0$, monotonicity implies that $\mathbb{P}(b > 1 | v) = 0$ for any $v < v_1$. Since $\mathbb{P}(b = 1 | v) = 0$ and $\mathbb{P}(b > 1 | v) = 0$ for all $v < v_1$, $\mathbb{P}(b = 0 | v) = 1$ for all $v < v_1$, i.e. all such values bid 0 with probability 1.

Next, consider now all $v \in \mathbb{X}$ with $v > v_m$ (where m is the largest bid submitted). By definition, v_m is the minimum value v such that $\mathbb{P}(b = m | v) > 0$. So $\mathbb{P}(b = m | v_m) > 0$. By monotonicity, this means that $\mathbb{P}(b \geq m | v) = 1$ for all $v > v_m$. But m is the largest bid submitted with positive probability. Hence, $\mathbb{P}(b = m | v) = 1$ for all $v > v_m$, i.e. all such

values bid m with probability 1.

Finally, consider all remaining values v such that $v \neq v_i$ for all $i = 1, \dots, m$. For any such value, we can find a bid i such that $v_i < v < v_{i+1}$. Using the previous two arguments, one can show that $\mathbb{P}(b \geq i|v) = 1$ and $\mathbb{P}(b \leq i|v) = 1$. But then $\mathbb{P}(b = i|v) = 1$, i.e. all such values bid i with probability 1.

It remains to consider the values $v \in \mathbb{X}$ such that $v = v_i$ for some $i = 1, 2, \dots, m$. Consider then any such value and let B^v denote the set of bids i such that $v_i = v$. Since the v_i are weakly increasing (in the bid i), B^v must be a set of consecutive integers (or a set containing a single integer). Label these integers $b', b' + 1, \dots, b' + k$. For all such bids, v is the minimum value such that $\mathbb{P}(b|v) > 0$. So all such bids must belong to the support of the strategy at value v .

We now argue that $\mathbb{P}(b \geq b' - 1|v) = 1$, i.e. bids strictly lower than $b' - 1$ cannot belong to the support. To see this, note that, since the v_i are increasing, $v_{b'} \geq v_{b'-1}$. However, $v_{b'} = v \neq v_{b'-1}$: otherwise, the bid $b' - 1$ would belong to B^v . Hence, $v_{b'} = v > v_{b'-1}$. So the minimum value that bids $b' - 1$ is strictly smaller than v . By monotonicity, this means that the value v must bid at least $b' - 1$ with probability 1.

Next, we argue that $\mathbb{P}(b > b' + k|v) = 0$. The argument is similar. Since the v_i are increasing, $v_{b''} \geq v_{b'+k}$ for any higher bid $b'' > b' + k$. However, $v_{b''} \neq v_{b'+k} = v$: otherwise, the bid b'' would be in B^v . So $v_{b''} > v_{b'+k}$. So the smallest value that bids b'' is strictly larger than v . As a result, $\mathbb{P}(b = b''|v) = 0$.

Given the previous arguments, we know that the support at v must contain $B^v = \{b', b' + 1, \dots, b' + k\}$. In addition, the only other bid that may belong to the support is $b' - 1$. So, for every bid $i \in B^v$, we can write

$$j_i = v_i + \mathbb{P}(b < i|v) = v + \sum_{j=b'-1}^{i-1} \mathbb{P}(b = j|v) \quad (20)$$

Thus, we have the system of equations

$$\begin{aligned} j_{b'} - v &= \mathbb{P}(b = b' - 1|v) \\ j_{b'+1} - v &= \mathbb{P}(b = b' - 1|v) + \mathbb{P}(b = b'|v) \\ &\vdots \\ &\vdots \\ j_{b'+k} - v &= \mathbb{P}(b = b' - 1|v) + \mathbb{P}(b = b'|v) + \dots + \mathbb{P}(b = b' + k - 1|v) \end{aligned} \quad (21)$$

Given the vector j , the first equation uniquely determines $\mathbb{P}(b = b' - 1|v)$ (which may equal zero). Using the value obtained for $\mathbb{P}(b = b' - 1|v)$, the second equation uniquely determines $\mathbb{P}(b = b'' - 1|v)$. Continuing in this manner, the system of equations uniquely pins down $\mathbb{P}(b = i|v)$ for $i = b' - 1, b', \dots, b' + k - 1$. Given the previous arguments, the only other bid in the support is $b = b' + k$ (the maximum bid). This is then also uniquely determined using

$$\mathbb{P}(b = b' + k|v) = 1 - \sum_{i=b''-1}^{b'-1} \mathbb{P}(b = i|v) \quad (22)$$

The argument above shows how a vector of jumps uniquely pins down a probability distribution over bids for any value $v \in \mathbb{X}$. Repeating the argument for every such value, we see that it is consistent with a unique behavioural strategy. \square

Having recast strategies in jump form, we can now define our algorithm:

Definition A.2 Let $\hat{j} = (\hat{j}_1, \dots, \hat{j}_m)$ denote the output of the following algorithm:

1. Impose $j_0 = 0$.
2. Starting at $i = 1$, find the minimum $j_i \in (j_{i-1}, S]$ such that $\pi(v = \lfloor j_i \rfloor, b = i) - \pi(v = \lfloor j_i \rfloor, b = i - 1) \geq 0$.
3. If there is no such j_i , the algorithm terminates.
4. There there does exist such a j_i , repeat (from step 2) for $i + 1$.

We then observe the following.

Lemma 5 The vector \hat{j} is the only possible symmetric equilibrium of the all-pay auction.

Proof. Suppose for contradiction that there were some symmetric equilibrium $j \neq \hat{j}$. If $j \neq \hat{j}$, there must be some jump at which they differ. Let j_i denote the first such jump (i.e. i is the smallest number such that $j_i \neq \hat{j}_i$). There are two possibilities: $j_i < \hat{j}_i$ or instead $j_i > \hat{j}_i$. We will argue that each leads to a contradiction.

To see why the first case is impossible, first define $\pi^{\hat{j}}(v, b)$ and $\pi^j(v, b)$ as the payoff of a player with a value v who bids b given that their opponents all bid according to \hat{j} and j respectively. Now recall that, by construction, \hat{j}_i is the *minimum* jump such that

$$\pi^{\hat{j}}(v = \lfloor \hat{j}_i \rfloor, b = i) \geq \pi^{\hat{j}}(v = \lfloor \hat{j}_i \rfloor, b = i - 1) \quad (23)$$

Since $j_i < \hat{j}_i$, this means that

$$\pi^j(v = \lfloor j_i \rfloor, b = i) < \pi^j(v = \lfloor j_i \rfloor, b = i - 1) \quad (24)$$

But then, given that one's opponents bid according to j , one should never bid $b = i$ with a value $v = \lfloor j_i \rfloor$. That is, $\mathbb{P}(b = i | v = \lfloor j_i \rfloor) = 0$. But then $v = \lfloor j_i \rfloor$ cannot be the smallest value that bids i (contradiction).

To see why the second case is impossible, recall again that every \hat{j}_i satisfies

$$\begin{aligned} \pi^{\hat{j}}(v = \lfloor \hat{j}_i \rfloor, b = i) &\geq \pi^{\hat{j}}(v = \lfloor \hat{j}_i \rfloor, b = i - 1) \iff \\ \lfloor \hat{j}_i \rfloor \mathbb{P}^{\hat{j}}(\text{win} | b = i) - 1 &\geq \lfloor \hat{j}_i \rfloor \mathbb{P}^{\hat{j}}(\text{win} | b = i - 1) \end{aligned} \quad (25)$$

As a preliminary, note that since the RHS is non-negative, $\lfloor \hat{j}_i \rfloor \mathbb{P}^{\hat{j}}(\text{win} | b = i) - 1 \geq 0$ and so $\lfloor \hat{j}_i \rfloor \geq 1$ for every i . Let us record this fact.

Turning to the main argument, fix the value at \hat{j}_i but consider now changing opponent bidding from \hat{j} to j . If $j_i > \hat{j}_i$, then $\mathbb{P}^j(\text{win} | b = i) > \mathbb{P}^{\hat{j}}(\text{win} | b = i)$. However, since j and \hat{j} agree for all previous jumps, $\mathbb{P}^j(\text{win} | b = i - 1) = \mathbb{P}^{\hat{j}}(\text{win} | b = i - 1)$. Since $\lfloor \hat{j}_i \rfloor \geq 1$ (established earlier), this all means that

$$\begin{aligned} \lfloor \hat{j}_i \rfloor \mathbb{P}^j(\text{win} | b = i) - 1 &> \lfloor \hat{j}_i \rfloor \mathbb{P}^{\hat{j}}(\text{win} | b = i - 1) \iff \\ \pi^j(v = \lfloor \hat{j}_i \rfloor, b = i) &> \pi^{\hat{j}}(v = \lfloor \hat{j}_i \rfloor, b = i - 1) \end{aligned} \quad (26)$$

If j is an SE, this implies that $\mathbb{P}^j(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) = 0$. However, this is impossible: since $j_i > \hat{j}_i$, $\mathbb{P}^j(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) = \min\{j_i, \lfloor \hat{j}_i \rfloor + 1\} - \max\{\hat{j}_{i-1}, \lfloor \hat{j}_i \rfloor\} > \hat{j}_i - \max\{\hat{j}_{i-1}, \lfloor \hat{j}_i \rfloor\} = \mathbb{P}^{\hat{j}}(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) \geq 0$ and so $\mathbb{P}^j(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) > 0$. \square

By the argument above, there is no possible symmetric equilibrium except for \hat{j} . At we noted earlier, however, our game must possess some symmetric equilibrium. From this, it follows that it has exactly one symmetric equilibrium (namely, \hat{j}). This concludes the proof. \square

Proof of Proposition 2. See main text. \square

Proof of Proposition 3. A proof of this result can be easily reconstructed from the proof of Proposition 1 (and is available from the author upon request). \square

Proof of Proposition 4. We argue by induction. First, we show that the desired pattern of bidding holds when $k = 1$. Next, we show that if it holds for any level $k - 1 \in \mathbb{K}$, then it also holds (provided that x is sufficiently large) for any level $k \in \mathbb{K}$.

Starting with the base case ($k = 1$), note that level-1 bidders choose a bid $b \in \mathbb{X}$ to maximise

$$\pi^1(v, b) = (v - b) \left(\frac{n-1}{n} + \frac{b}{n(x+1)} \right)^{n-1} \quad (27)$$

(inserting $p = 1/n$ into 5). Considering this function's continuous extension, observe that

$$\frac{\partial \ln[\pi^1(v, b)]}{\partial b} = \frac{(n-1)(v - (x+1)) - nb}{(v-b)(b + (n-1)(x+1))} \quad (28)$$

Since we are searching for the optimal bid, we can just consider the range $b \leq v$. In this range, the denominator of (28) is non-negative. Moreover, since $v \leq x$ for all $v \in \mathbb{X}$, the numerator is negative. Thus, $\partial \ln[\pi^1(v, b)]/\partial b \leq 0$ and so $\partial \pi^1(v, b)/\partial b \leq 0$. Since this is true of the function's continuous extension, it follows that $\pi(v, 0) \geq \pi(v, b)$ for every value $v \in \mathbb{X}$ and every positive integer bid $b \in \mathbb{X}$. That is, bidding zero is optimal (as claimed).

Next, we show that, if the desired pattern of bidding holds for level $k-1 \in \mathbb{K}$, then it also holds (provided that x is sufficiently large) for level $k \in \mathbb{K}$. Suppose then that all players of level $k-1$ bid as proposed. Then they bid 0 when $v \leq v^*(k-1)$, i.e. at $\lfloor v^*(k-1) + 1 \rfloor$ distinct values; and otherwise, they bid $k-2$. Hence, for a level- k player,

$$\mathbb{P}^k(\text{win}|b) = \begin{cases} 1 & \text{if } b > k-2 \\ \left(1 - p + p \left(\frac{\lfloor v^*(k-1) + 1 \rfloor}{x+1}\right)\right)^{n-1} & \text{if } b \in \{1, \dots, k-2\} \\ (1-p)^{n-1} & \text{if } b = 0 \end{cases} \quad (29)$$

Given this, it is immediate that $b^* \in \{0, 1, k-1\}$, i.e. there are at most three optimal bids. Moreover, $\pi^k(v, b=0) - \pi^k(v, b=1) \rightarrow v(1-p)^{n-1} - (v-1)(1-p)^{n-1} < 0$ as $x \rightarrow \infty$, and so $b^* \in \{0, k-1\}$ if x is sufficiently large. Finally, observe that the latter is optimal iff

$$\begin{aligned} \pi(v, b=k-1) > \pi(v, b=0) &\iff v - (k-1) > v(1-p)^{n-1} \\ &\iff v > \frac{k-1}{1 - (1-p)^{n-1}} \end{aligned} \quad (30)$$

i.e. iff $v > v^*(k)$ (as claimed). □

B Illustrating the algorithm

In this section, we illustrate the algorithm using two examples. In both, values are uniform so (in any symmetric strategy profile)

$$\begin{aligned}
 \mathbb{P}(b < i) &= \mathbb{P}(v < \lfloor j_i \rfloor) + \mathbb{P}(v = \lfloor j_i \rfloor)(j_i - \lfloor j_i \rfloor) \\
 &= \frac{\lfloor j_i \rfloor}{S} + \frac{1}{S}(j_i - \lfloor j_i \rfloor) \\
 &= \frac{j_i}{S}
 \end{aligned} \tag{31}$$

where S is the number of possible valuations.

Example 1. Consider an all-pay auction with uniform values and $n = 2$. To find the first jump point, we look for the minimum $j_1 \in [2, S]$ such that

$$\pi(v = \lfloor j_1 \rfloor, b = 1) - \pi(v = \lfloor j_1 \rfloor, b = 0) \geq 0 \tag{32}$$

Clearly, $\pi(v = \lfloor j_1 \rfloor, b = 0) = 0$ since $\mathbb{P}(\text{win}|b = 0) = 0$. In addition, $\mathbb{P}(\text{win}|b = 1) = \mathbb{P}(b = 0) = j_1/S$ and so $\pi(v = \lfloor j_1 \rfloor, b = 1) = \lfloor j_1 \rfloor j_1/S - 1$. So we look for the minimum j_1 such that

$$\frac{\lfloor j_1 \rfloor j_1}{S} - 1 \geq 0 \tag{33}$$

Of course, the solution will depend on S . If $S = 100$, $j_1 = 10$ (and similarly $j_1 = \sqrt{S}$ for any square S). If $S = 99$, the solution remains at $j_1 = 10$. However, if $S = 101$, the solution becomes $j_1 = 10.1$, implying randomisation. (One can prove that, as $S \rightarrow \infty$, the proportion of solutions that are integer solutions converges to $1/2$.) Let us suppose that $S = 101$ so that the first jump is $j_1 = 10.1$.

To find the next jump, we look for the smallest j_2 that satisfies

$$\begin{aligned}
 &\pi(v = \lfloor j_2 \rfloor, b = 2) - \pi(v = \lfloor j_2 \rfloor, b = 1) \geq 0 \\
 \iff &\lfloor j_2 \rfloor (\mathbb{P}(\text{win}|b = 2) - \mathbb{P}(\text{win}|b = 1)) \geq 1 \\
 \iff &\lfloor j_2 \rfloor (\mathbb{P}(b \leq 1) - \mathbb{P}(b = 0)) \geq 1 \\
 \iff &\lfloor j_2 \rfloor \left(\frac{j_2}{101} - \frac{10.1}{101} \right) \geq 1
 \end{aligned} \tag{34}$$

One can check that the solution is $j_2 = 16.4125$, so the equilibrium vector of jumps starts with $(0, 10.1, 16.4125, \dots)$. Returning to behavioural strategies, this means that all values $v \in \{0, 1, \dots, 9\}$ bid 0 (with probability 1), value $v = 10$ randomises between $b = 0$ and $b = 1$

with probabilities 0.1 and 0.9, values $v \in \{11, \dots, 15\}$ bid 1 (with probability 1), and finally the value $v = 16$ bids $b = 1$ with probability 0.4125. (Since we haven't computed j_3 , we cannot technically determine the probability with which value $v = 16$ bids $v = 2$; but in fact $j_3 \geq 17$ so this probability is $1 - 0.4125$.)

Example 2. Consider now the first price auction with cancelled bids. To find the first jump j_1 , we look for the smallest $j_1 \in (0, S]$ such that

$$\begin{aligned}
& \pi(v = \lfloor j_1 \rfloor, b = 1) \geq \pi(v = \lfloor j_1 \rfloor, b = 0) \\
& \iff (\lfloor j_1 \rfloor - 1)\mathbb{P}(\text{win}|b = 1) \geq \lfloor j_1 \rfloor\mathbb{P}(\text{win}|b = 0) \\
& \iff (\lfloor j_1 \rfloor - 1)(1 - p + p\frac{j_1}{S})^{n-1}p \geq \lfloor j_1 \rfloor(1 - p)^{n-1}p \\
& \iff (\lfloor j_1 \rfloor - 1)(1 - p + p\frac{j_1}{S})^{n-1} \geq \lfloor j_1 \rfloor(1 - p)^{n-1}
\end{aligned} \tag{35}$$

For instance, suppose that $p = 1/2$, $n = 2$ and $S = 101$. Then (35) reduces to

$$\lfloor j_1 \rfloor j_1 - 101 - j_1 \geq 0 \tag{36}$$

which is almost quadratic in j_1 . One can check that this inequality is satisfied by $j_1 = 11$: then $\lfloor j_1 \rfloor j_1 - 101 - j_1 = 11^2 - 101 - 11 = 9 > 0$. On the other hand, it won't hold if $j_1 < 11$: for then $\lfloor j_1 \rfloor j_1 - 101 - j_1 \leq 10 \times 11 - 101 - 11 = -2 < 0$. So the minimum j_1 that satisfies the inequality is 11; this is the first jump. This means that all values smaller than 11 must bid zero (with probability 1); but the value $v = 11$ never bids zero with any probability.

C Optimising p

To find the probability $p^* \in [0, 1]$ that maximises $d(p)$, recall that $p^* > 0$ (from Proposition 5). Moreover, it is obvious that $p^* < 1$: if $p = 1$, then $d(p) = 0$ (whereas $d(p) > 0$ is plainly possible). Hence, $p^* \in (0, 1)$, which means that the optimal p must satisfy the first-order condition

$$\frac{\partial}{\partial p} \int_0^x |\beta(v) - \beta^1(v)| dv = 0. \quad (37)$$

One can check that $\beta(v) \geq \beta^1(v)$ for all $v \in [0, x]$. So (37) is equivalent to

$$\frac{\partial}{\partial p} \int_0^x \beta(v) dv = \frac{\partial}{\partial p} \int_0^x \beta^1(v) dv. \quad (38)$$

When $n = 2$, standard calculations reveal that

$$\begin{aligned} \frac{\partial}{\partial p} \int_0^x \beta(v) dv &= \frac{\partial}{\partial p} \int_0^x \frac{pv^2}{2(pv - p + 1)} dv \\ &= \frac{\partial}{\partial p} \frac{x^2(p(3p - 2) - 2(p - 1)^2 \ln(1 - p))}{4p^2} \\ &= \frac{x^2((2 - p)p - 2(p - 1) \ln(1 - p))}{2p^3} \end{aligned} \quad (39)$$

Similarly, when $n \geq 3$, we find that

$$\begin{aligned} \frac{\partial}{\partial p} \int_0^x \beta(v) dv &= \frac{\partial}{\partial p} \int_0^x \left(\frac{n-1}{n} \right) v - \frac{x(1-p)}{np} \left[1 - \left(\frac{1-p}{1-p+p(v/x)} \right)^{n-1} \right] dv \\ &= \frac{\partial}{\partial p} \frac{x^2(2 + np((n-1)p - 2)) - 2(1-p)^n}{2(n-2)np^2} \\ &= \frac{x^2(p(n(1-p)^n - 2(1-p)^n + n + 2) + 2((1-p)^n - 1) - np^2)}{(n-2)n(p-1)p^3} \end{aligned} \quad (40)$$

Turning to the level-1 bidding function, we find that

$$\begin{aligned} \frac{\partial}{\partial p} \int_0^x \beta^1(v) dv &= \frac{\partial}{\partial p} \int_{v^*}^x \left(\frac{n-1}{n} \right) v - \left(\frac{1-p}{p} \right) \frac{x}{n} dv \\ &= \frac{\partial}{\partial p} \frac{x^2(np - 1)^2}{2(n-1)np^2} \\ &= \frac{x^2(np - 1)}{(n-1)np^3} \end{aligned} \quad (41)$$

where $v^* = (1-p)x/(n-1)p$.

To find the optimal probability when $n = 2$, we now substitute (39) and (41) into the first-order condition (38), obtaining

$$\frac{(2-p)p - 2(p-1)\ln(1-p)}{2p^3} = \frac{np-1}{(n-1)np^3} \quad (42)$$

Similarly, when $n \geq 3$, we substitute (40) and (41) and into the first-order condition, yielding

$$\frac{p(n(1-p)^n - 2(1-p)^n + n + 2) + 2((1-p)^n - 1) - np^2}{(n-2)n(p-1)p^3} = \frac{np-1}{(n-1)np^3} \quad (43)$$

Either way, we obtain an equation that can be solved numerically to reveal candidates for p^* . It is then a straightforward matter to check which is of these candidates generate a larger distance. Since our problem must have a global maximiser (by the Weierstrass theorem), and the global maximiser must satisfy the first-order condition, the p satisfying the first-order condition that generates the largest distance must be the global maximiser.

Table 2: Optimal cancellation probabilities (rounded)

n	2	3	4	5	6	7	8
p^*	0.536	0.343	0.256	0.204	0.170	0.145	0.127

Table 2 outlines for the solutions for $n \in \{1, \dots, 8\}$. As can be seen, the solution is generally quite close to $1/n$ — so the bound identified in Proposition 5 comes close to identifying the exact solution. Note also that, as might be expected, the optimal probability p does not depend on the scale parameter x .

D Tables and figures

Figure 8: The first-price auction when bids are multiples of 5

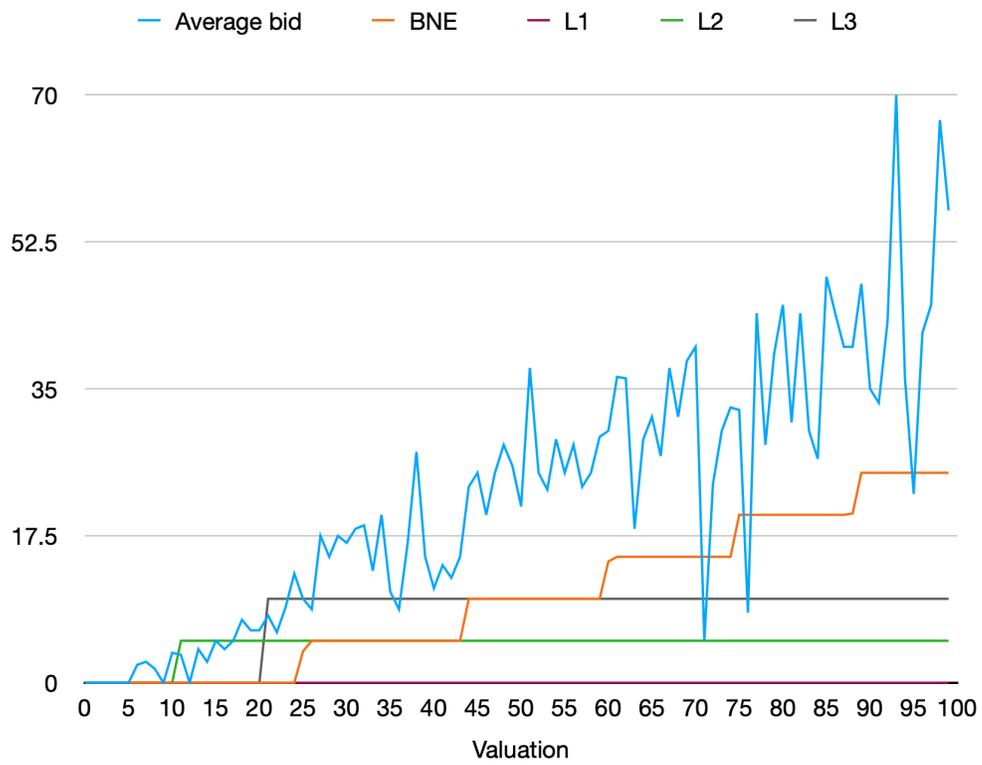


Figure 9: The all-pay auction when bids are multiples of 5

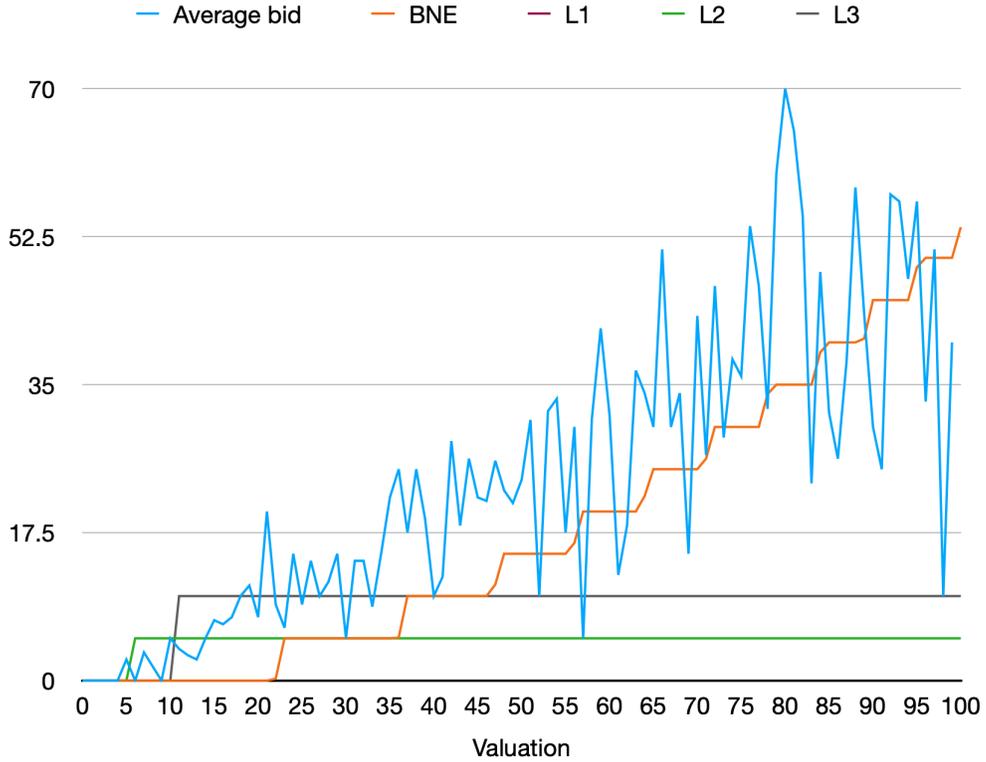


Table 3: Balance Table

	Integer bids	Multiples of five	p -value
Mean age	31.4	32.0	0.29
Share male	0.54	0.43	0.33
Share maths	0.39	0.38	0.95
Share humanities	0.24	0.14	0.36
Share social-sciences	0.30	0.29	0.93

Notes: This table displays how demographics (age, sex and subject) vary between the two treatment groups. The column ‘ p -value’ reports t -tests of the hypothesis that the relevant variable is on average equal across the groups.

Table 4: Average bids

	First-price			All-pay		
	T1	T2	<i>p</i> -value	T1	T2	<i>p</i> -value
Average bid	23.6	23.0	0.85	22.2	25.4	0.12
Average value	50.9	49.6	0.27	49.4	50.8	0.89
Ratio	0.46	0.46		0.45	0.50	

Notes: This table reports the average bids across two treatments and auction structures. The column ‘*p*-value’ reports *t*-tests of the hypothesis that the relevant variable is on average equal across the groups.

Table 5: Comparing equilibrium and level-*k*

		T1FP	T1AP	T2FP	T2AP
BNE	LL	-4407.9	-4332.7	-1024.7	-1042.2
	BIC	8820.5	8670.2	2053.1	2088.1
L1	LL	-4530.4	-4501.2	-1088.8	-1128.6
	BIC	9065.4	9007.1	2181.3	2260.9
L1-L2	LL	-4530.4	-4501.2	-1085	-1126.9
	BIC	9070.1	9011.7	2177.4	2261.2
L1-L3	LL	-4530.1	-4501.2	-1071.1	-1120.7
	BIC	9074.3	9016.4	2153.4	2252.6

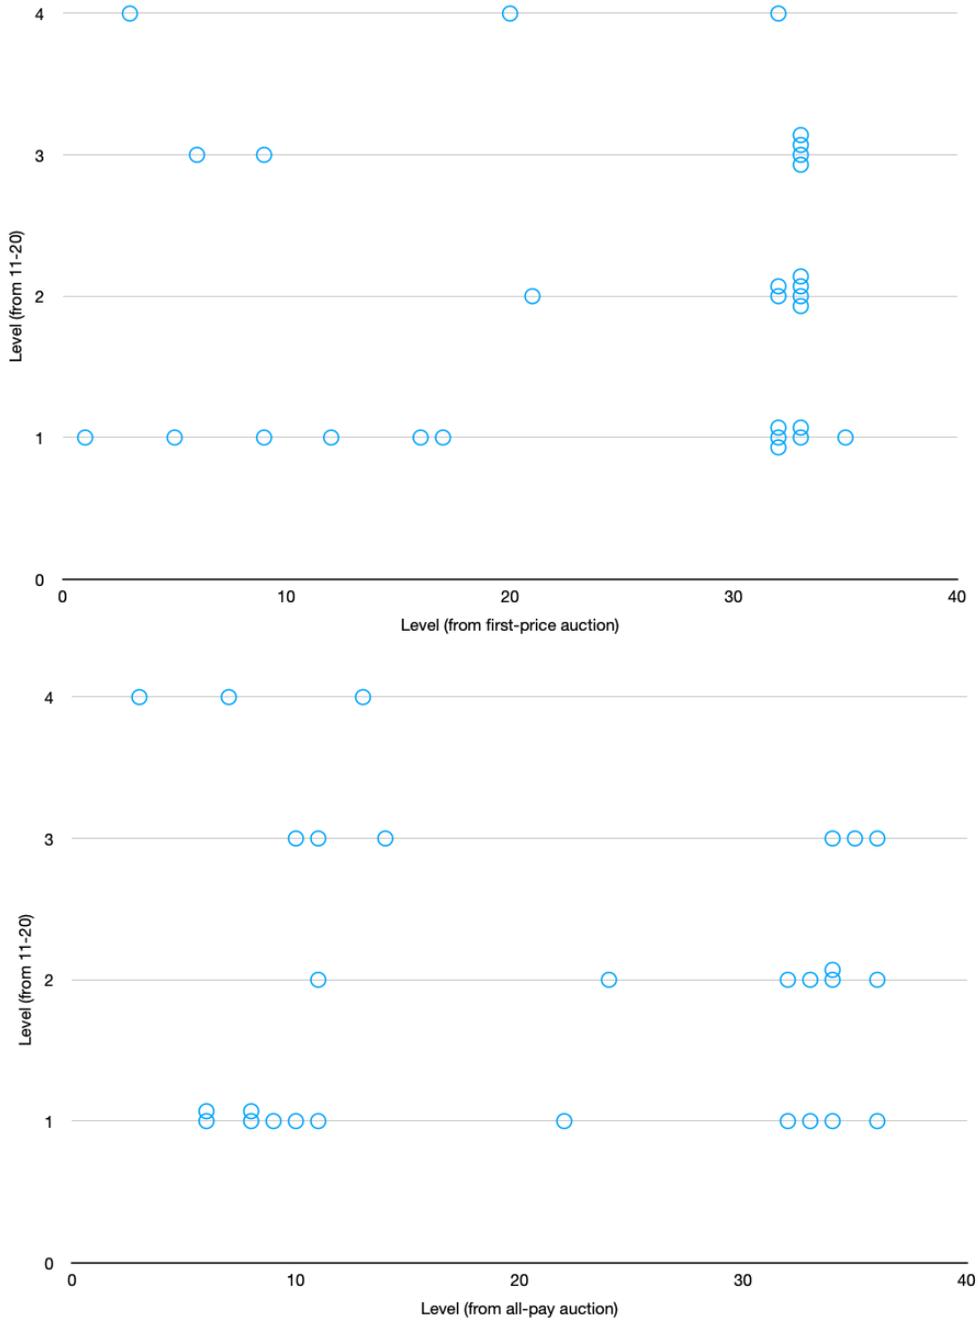
Notes: This table reports the log-likelihoods and associated Bayes-Information Criteria values for four different structural models. In the first model, equilibrium is the only type. The subsequent models are populated by level-1 types, level-1 and 2 types, and level-1, 2 and 3 types respectively.

Table 6: The hybrid model

	T1FP		T1AP		T2FP		T2AP	
	Subject	Round	Subject	Round	Subject	Round	Subject	Round
p_0	0.758 (0.059)	0.760 (0.046)	0.815 (0.053)	0.757 (0.06)	0.905 (0.066)	0.858 (0.054)	0.741 (0.149)	0.788 (0.083)
p_1	0.036 (0.025)	0.046 (0.02)	0.000 (0.0)	0.010 (0.01)	0.095 (0.066)	0.044 (0.034)	0.048 (0.048)	0.071 (0.04)
p_2	0.133 (0.051)	0.027 (0.016)	0.019 (0.019)	0.055 (0.022)	0.000 (0.0)	0.000 (0.0)	0.048 (0.048)	0.024 (0.024)
p_3	0.073 (0.041)	0.167 (0.041)	0.167 (0.051)	0.178 (0.057)	0.000 (0.0)	0.098 (0.048)	0.164 (0.141)	0.117 (0.073)
σ_0	27.3 (1.9)	27.3 (1.8)	19.9 (1.6)	20.2 (1.5)	23.2 (3.0)	23.7 (2.4)	21.7 (2.8)	21.3 (2.0)
σ_1	3.8 (2.0)	1.8 (0.3)	-	2.3 (0.4)	5.4 (2.2)	6.3 (0.7)	2.7 (0.0)	2.4 (0.4)
σ_2	9.2 (0.6)	4.2 (0.5)	1.1 (0.0)	1.0 (0.0)	-	-	2.7 (0.0)	1.0 (0.0)
σ_3	13.3 (1.1)	10.5 (1.2)	5.3 (0.6)	7.9 (2.1)	-	4.1 (0.9)	14.5 (7.0)	6.5 (4.1)
LL	-4241.5	-4217.1	-4043.8	-4004.1	-987.9	-982.6	-962.9	-940.9

Notes: The ‘subject’ columns report analyses in which subjects’ levels are fixed; the ‘round’ columns report analyses in which levels are allowed to vary between rounds. The symbols p_i and σ_i respectively denote the probability and noise parameter of type $i \in \mathbb{K}$ (equilibrium is indexed by $i = 0$). LL abbreviates ‘log-likelihood’ and jack-knife standard errors are in parentheses. When $p_i = 0$, the associated parameter σ_i is unidentified and therefore omitted from the table.

Figure 10: Comparing estimated levels



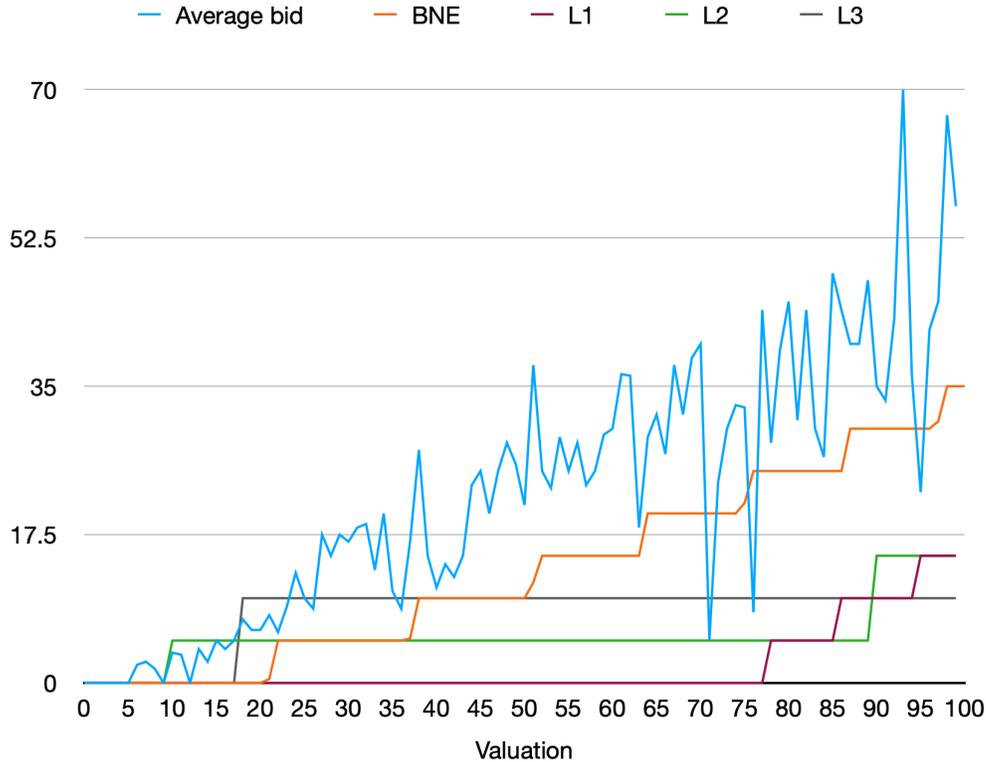
Notes: The first panel plots levels as estimated from the first-price auction against levels as inferred from the 11-20 game. The estimated correlation between these variables is 0 (95% confidence interval: -0.37 to 0.37). The second panel plots levels as estimated from the all-pay auction against levels as inferred from the 11-20 game. The estimated correlation between these variables is -0.09 (95% confidence interval: -0.44 to 0.30).

Table 7: Robustness checks

		T1FP	T1AP	T2FP	T2AP
<i>Risk aversion</i>	BNE	19.4	-	17.9	-
	Level- <i>k</i>	26.9	-	22.7	-
<i>Cognitive hierarchy</i>	BNE	21.7	18.2	19.7	18.8
	Level- <i>k</i>	31.3	23.8	26.6	29.6
<i>Dominated bids</i>	BNE	22.6	19.4	19.5	19.0
	Level- <i>k</i>	30.8	29.9	22.7	26.9
<i>Tie-breaking</i>	BNE	21.7	18.2	19.7	18.8
	Level- <i>k</i>	30.3	29.5	23.4	26.6
<i>Dropping round 2</i>	BNE	18.4	18.2	18.3	17.4
	Level- <i>k</i>	26.5	29.0	22.2	24.0

Notes: This table reports the root-mean-square prediction errors of the relevant theories following a series of robustness checks. In every case, level-*k* predictions are obtained by assigning each datapoint the level from the 1 – 3 range that minimises the model’s prediction error (a procedure that will tend to overstate the model’s predictive performance).

Figure 11: Risk aversion (multiples of 5 treatment)



Competitive equilibrium and the double auction

Abstract

In this paper, we revisit the common claim that double auctions necessarily generate competitive equilibria. We begin by observing that competitive equilibrium has some counterintuitive implications: specifically, it predicts that monotone shifts in the value distribution can leave prices unchanged. Using experiments, we then test whether these implications are borne out by the data. We find that in double auctions with stationary value distributions, the resulting prices can be far from competitive equilibria. However, we also show that double auctions display stronger equilibrating tendencies once traders can leave without replacement as time progresses. Taken together, these findings suggest that the exit of traders over time is a key driver of equilibrium prices in double auctions.

KEYWORDS: double auction, competitive equilibrium, Marshallian path

JEL CODES: C92, D01, D02, D90

1 Introduction

In his ‘Introduction to Economic Science’, Fisher (1910) wrote that ‘If you want to make a first-class economist, catch a parrot and teach him to say “supply and demand” in response to every question you ask him.’ Apparently, this joke was considered dated even in 1910 — it is attributed to a critic of economics writing from a ‘long time ago’ — and it thus serves to illustrate the dominance of supply and demand in the century before Fisher’s publication. However, supply and demand analysis remained popular in Fisher’s time; and indeed one of the purposes of Fisher’s book was to expound and refine such analysis.

More than a century later, the idea of supply and demand — or ‘competitive equilibrium’, as we might now call it — remains pervasive throughout economic theory. Its extension to multiple markets in the form of general equilibrium theory has provided a basis for celebrated welfare theorems (Arrow, 1951; Debreu, 1954), as well as the foundation for much of modern macroeconomics (starting with Lucas (1977), Kydland and Prescott (1982), etc.). In addition, partial equilibrium models have been applied to areas as diverse as discrimination (Becker, 1957), marriage (Becker, 1973, 1974) and location choice (Glaeser, 2007). Thus, although the notion of competitive equilibrium may not have quite the dominance that it did in Fisher’s time — supplanted in part by the rise of non-cooperative game theory — it surely remains one of the foundational concepts of economic theorising.

In part, the pervasiveness of competitive equilibrium may be due to the perception that its predictions have been experimentally vindicated by a series of double auction experiments starting with Smith (1962). As Plott (1981) puts it in an early review of the literature: ‘The overwhelming result [from these experiments] is that these markets converge to the competitive equilibrium even with very few traders’. In a more recent study of two thousand classroom experiments, Lin et al. (2020) reach a similar conclusion, declaring that competitive equilibrium convergence in double auctions appears to be ‘as close to a culturally universal, highly reproducible outcome as one is likely to get in social science’. They add that such convergence should be considered ‘as reproducible as the kinds of experiments that are done in a college chemistry laboratory to demonstrate universal chemistry principles’.¹

In this paper, we revisit this conclusion. Our starting point is that competitive equilibrium can make highly counterintuitive, but previously unstudied, predictions. To see the basic

¹Indeed, existing experimental results appear to be surprisingly robust to changing the number of bidders (Smith, 1965), changing the cultural context (Kachelmeier and Shehata, 1992), and even introducing ‘extreme earnings inequality’ at equilibrium (Holt et al., 1986; Smith and Williams, 2000; Kimbrough and Smyth, 2018). See also Ikica et al. (2018) for a large-scale replication of competitive equilibrium convergence in standard double auction environments.

idea, consider a market populated by a group of buyers and a group of sellers. Given the distribution of buyer valuations and seller costs, it is straightforward to calculate the set of ‘equilibrium’ prices that equate total demand with total supply. Now imagine that we decrease the costs of all sellers whose costs were initially below the minimum equilibrium price. Since these sellers were willing to sell before the change (at any equilibrium price), and remain willing to sell after the change (at any of these equilibrium prices), such a change cannot alter the set of equilibrium prices. Intuitively, however, one might expect the change to decrease the prices that are realised in the market — both by allowing sellers to profitably submit lower offers, and also by inducing them to accept lower offers from buyers.

In Section 2, we begin by generalising the example just given. That is, we identify a broad class of downward shifts to the distributions of buyer and seller valuations which preserve the set of competitive equilibrium prices. As discussed, it is highly counterintuitive that such downward shifts would in fact leave market prices unaltered. As a result, such shifts provide a challenging test for competitive price theory.

We then conduct experimental double auctions to investigate whether such shifts do in fact depress observed prices. One important feature of our initial set of double auction experiments is that they hold the value distributions fixed through use of a queue of buyers and sellers: every time a buyer or seller exits the market, a new buyer or seller takes their place (see Brewer et al. (2002) for a similar methodology). We use this queueing procedure for two reasons. First, it can be justified on grounds of realism: in actual markets, trade does not end after a couple of periods once all willing buyers and sellers have traded. Instead, the market is continually replenished by a new stock of traders. Second, and much more importantly, our design necessarily holds competitive equilibrium fixed, thereby allowing us to rigorously study if double auctions outcomes converge to *the* set of competitive equilibrium prices. In contrast, standard designs do not possess this feature: every time a pair of traders drop out of the market, the supply and demand schedules shift, something which may (or may not) change the set of competitive equilibrium prices.²

Our initial set of experiments yields three main findings. First, contrary to the predictions of competitive equilibrium theory, our downward shifts in the valuations do markedly decrease observed prices. Second, and partly as a result of the first finding, prices are almost never at

²Remarkably, this issue was noticed in the first-ever experimental study of double auctions: see footnote 6 of Smith (1962). Despite conceding that ‘the supply and demand functions continually alter as the trading process occurs’, Smith asserts that it is ‘the *initial* [supply and demand] schedules prevailing at the opening of each trading period’ that are of interest to ‘the theorist’. However, Smith does not give us any reason for privileging the initial demand and supply schedules over any others; and in the absence of such a reason, such privileging would appear to be entirely arbitrary.

the competitive equilibrium. Remarkably, this is the case even for our symmetric treatment, which might have been expected to yield competitive equilibrium prices. Third, prices show very little sign of converging to competitive equilibrium. Taken together, these findings imply that competitive equilibrium is not a good description of double auctions with stationary value distributions.

We then discuss why decreasing valuations and costs might decrease observed prices. While we do not attempt to identify one model as definitively ‘correct’, we do identify a number of models that can rationalise our finding. To take one example, within a zero intelligence model (Gode and Sunder, 1993), decreasing valuations leads to stochastically lower distributions of bids and asks, thereby lowering expected prices. More interestingly, perhaps, this effect is also generated by optimising models, including those based both on myopic pay-off maximisation (Gjerstad and Dickhaut, 1998) and more sophisticated optimal stopping (Friedman, 1991).

Our first set of experiments establishes that, if value distributions are held fixed, double auctions need not produce competitive equilibria; and that double auction outcomes are sensitive to the kind of value shifts described previously. It is natural to wonder, however, whether our shifts are still able to move observed prices in more standard double auction environments in which players are allowed to drop out of the market (without replacement) as trade progresses. To investigate this question, we also run a series of double auctions without queues, including some very long sessions involving over an hour’s trading in order to give the auctions the best possible chance of equilibrating.

Our findings from these more conventional experiments are more mixed. On the one hand, there is still some evidence that our shifts depress the observed prices; and some of the sessions we run fail to equilibrate even after many periods of trading. On the other hand, there is now a marked tendency towards equilibrium, and some of our sessions do converge to equilibrium despite the marked asymmetry in the designs. Thus, these sessions reveal both the power of our shifts as well as the equilibrating forces first observed in Smith (1962).

Our central pair of findings — that one can easily ‘break’ competitive equilibrium in environments with a queue of buyers and sellers but much less easily in an environment without a queue — suggests that the question of whether the value distribution is held fixed is of critical importance. We provide a theoretical explanation as to why this should be the case. We assume that trade follows a ‘Marshallian path’ (Brewer et al., 2002), which means that (i) at any point in time, trade takes place between the active buyer with the highest valuation and the active seller with the lowest cost, and (ii) trade occurs if and only if it is mutually beneficial (we provide a formal definition in Section 6). We prove that in standard double

auction formats, a Marshallian path implies that final trades must take place at equilibrium prices; and we discuss why this result might also extend to non-final trades. Importantly, this result does *not* extend to double auctions with fixed value distributions, which can help explain why such auctions lack the standard equilibrating tendencies. Therefore, we identify the Marshallian path dynamic as a key driver of equilibration in standard double auctions, thus helping to solve the ‘scientific mystery’ introduced by Vernon Smith 60 years ago.

The remainder of this article is structured as follows. Section 2 generalises and formalises the idea of competitive equilibrium preserving shifts. Section 3 outlines the design of experiments aimed at testing the impacts of such shifts; Section 4 contains the results of such experiments; and Section 5 discusses which models can rationalise our findings. Section 6 uses a combination of theory and further experimentation to study whether our findings change once value distributions are no longer held stationary. Finally, Section 7 concludes with an outline of some new areas of research opened up by this work.

2 CE preserving shifts

In this section, we generalise and formalise the example from the introduction in order to obtain a better understanding of its structure. To this end, let us consider a unit mass of buyers, each indexed by $i \in [0, 1]$.³ Each buyer has a valuation $v_i \in [0, \bar{v}_b]$ (where $\bar{v}_b > 0$ is the maximum buyer valuation) and chooses to buy (exactly) one unit of the good if and only if their valuation is at least the market price p . The distribution of buyer valuations is described by the cumulative distribution function $F: \mathbb{R} \rightarrow [0, 1]$. For simplicity, we assume that (i) F has full support on the interval $[0, \bar{v}_b]$ (ii) F is continuous.

Let $d(p)$ denote market demand at price p . Then

$$d(p) = \int_0^1 \mathbb{1}(v_i \geq p) di = \mathbb{P}(v_i \geq p) = 1 - F(p) \quad (1)$$

where $\mathbb{1}(v_i \geq p)$ is an indicator function. Since, $d(p) = 1 - F(p)$, $d(0) = 1 - F(0) = 1$ and $d(\bar{v}_b) = 1 - F(\bar{v}_b) = 0$. That is, demand starts at 1 (at a price of 0) and eventually falls to 0 (at a price of \bar{v}_b). In addition, observe that d is continuous (since it inherits the continuity of F) and strictly decreasing over the interval $[0, \bar{v}_b]$ (since F has full support on this interval). In summary, then, we obtain a continuous and strictly decreasing demand function which starts at 1 before falling to 0 when the price equals the maximum valuation.

³While we work with a continuum of buyers and sellers for convenience, similar results are available in the discrete case.

We treat sellers entirely symmetrically. That is, we have a unit mass of sellers, indexed by $i \in [0, 1]$; and each seller has a valuation (or ‘cost’) $v_i \in [0, \bar{v}_s]$ (where $v_s > 0$ is the maximum seller valuation). Seller i chooses to sell their one unit of the good if and only if $v_i \leq p$. The distribution of seller valuations is described by the cumulative distribution function $G: \mathbb{R} \rightarrow [0, 1]$. As before, we assume that (i) G has full support on the interval $[0, \bar{v}_s]$ (ii) G is continuous.

Let $s(p)$ denote market supply at price p . Then

$$s(p) = \int_0^1 \mathbb{1}(v_i \leq p) di = \mathbb{P}(v_i \leq p) = G(p) \quad (2)$$

where $\mathbb{1}(v_i \leq p)$ is again an indicator function. Since $s(p) = G(p)$, observe that $s(0) = G(0) = 0$ and $s(\bar{v}_s) = G(\bar{v}_s) = 1$. Moreover, s is continuous (since G is continuous) and strictly increasing over the interval $[0, \bar{v}_s]$ (since G has full support on that interval).

We define a *competitive equilibrium price* as a price $p^* \in \mathbb{R}^+$ such that $d(p^*) = s(p^*)$. According to competitive price theory, this is the price which will prevail in a market; and the associated quantity traded will be $d(p^*) = s(p^*)$. As an aside, we notice that, although competitive price theory gives us a clear prediction as to the price which should prevail in a market, it does not provide us with an explanation as to *why* such a price should arise. Whether such an explanation can be provided is itself an interesting question.⁴

As our first observation, let us note that, under our assumptions, there is exactly one competitive equilibrium price p^* . To see this, define excess demand by $e(p) = d(p) - s(p)$ and note that excess demand is positive at a price of zero (specifically, $e(0) = d(0) - s(0) = 1 - 0 = 1$). Meanwhile, $e(\bar{v}_b) = d(\bar{v}_b) - s(\bar{v}_b) = 0 - G(\bar{v}_b)$, so excess demand becomes negative at a price of v_b . Given that excess demand inherits the continuity of d and s , this means that there must be some p^* such that $e(p^*) = 0$, i.e there exists an equilibrium price. Moreover, since e is strictly decreasing, we see that p^* is unique.

We now define the central concept of this section.

Definition 1 *A competitive equilibrium preserving demand contraction is a transformation $T_b: [0, \bar{v}_b] \rightarrow \mathbb{R}^+$ such that*

1. $T(V_b) \leq V_b$ for all $V_b \leq p^* - \epsilon^-$
2. $T(V_b) = V_b$ for all $V_b \in (p^* - \epsilon^-, p^* + \epsilon^+)$

⁴Aumann (1964) and Cripps and Swinkels (2006) are particularly influential attempts to provide a foundation to competitive equilibrium theory. The result we prove in Section 6 provides a rather different (and much simpler) foundation, albeit one that is most directly applicable to final period trades.

3. $T(V_b) \in [p^* + \epsilon^+, V_b]$ for all $V_b \geq p^* + \epsilon^+$

for some $\epsilon^+, \epsilon^- > 0$.

As stated above, competitive equilibrium preserving demand contractions are downward shifts to the distribution of buyer valuations that satisfy three conditions. First, we require that low valuations (specifically those below $p^* - \epsilon^-$) are weakly decreased. Second, we require that there is some (possibly asymmetric) ϵ -ball around p^* at which valuations remain unchanged. Finally, we require that high valuations (those above $p^* + \epsilon^+$) are reduced, but not reduced so much that they are brought below $p^* + \epsilon^+$. Altogether, this amounts to a stochastic reduction to the distribution of buyer valuations, although one that leaves the ranking (i.e. percentile) of valuations in a neighbourhood of p^* unchanged.

We define a competitive equilibrium preserving decrease in the seller valuations analogously.

Definition 2 *A competitive equilibrium preserving supply expansion is a transformation $T_s: [0, \bar{v}_s] \rightarrow \mathbb{R}^+$ such that*

1. $T(V_s) \leq V_s$ for all $V_s \leq p^* - \epsilon^-$
2. $T(V_s) = V_s$ for all $V_s \in (p^* - \epsilon^-, p^* + \epsilon^+)$
3. $T(V_s) \in [p^* + \epsilon^+, V_s]$ for all $V_s \geq p^* + \epsilon^+$

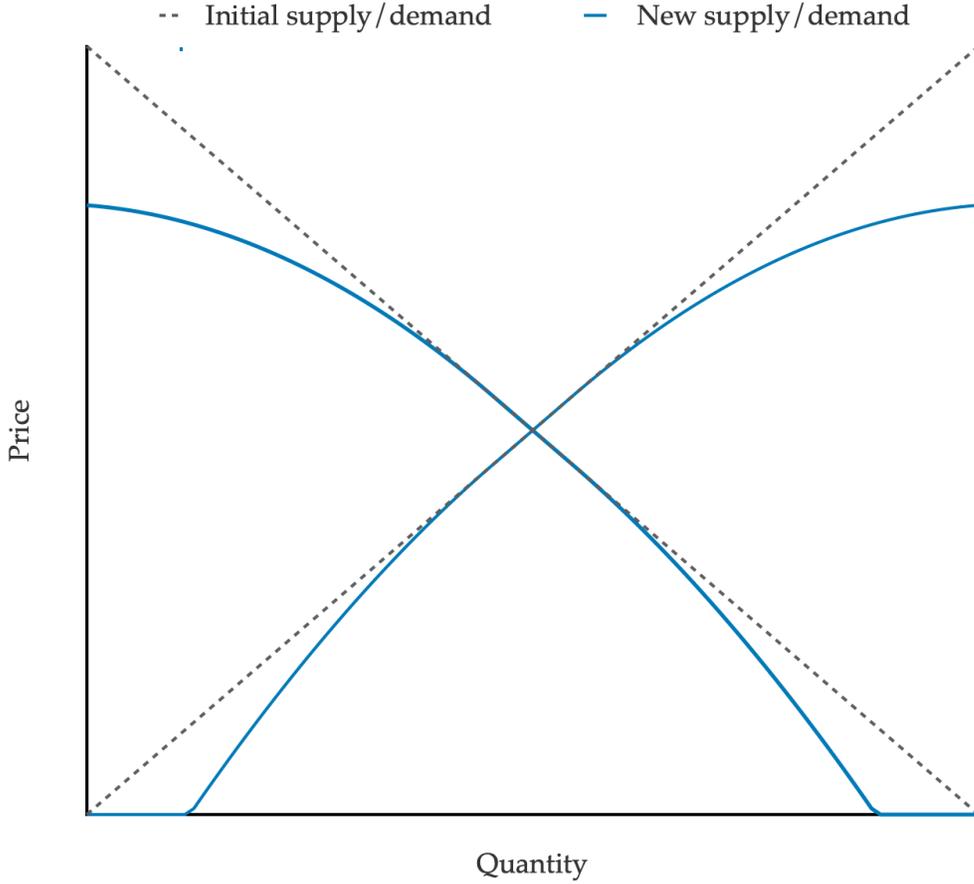
for some $\epsilon^+, \epsilon^- > 0$.

Figure 1 plots a competitive equilibrium preserving demand contraction and supply expansion. As can be seen, we have decreased the buyer and seller valuations when the values were previously low. This corresponds to a downward shift in the demand and supply functions in the leftward portion of the diagram. In addition, we have left the buyer and seller valuations around the equilibrium price unchanged, which means that the demand and supply schedules remain unchanged in a neighbourhood of p^* . Finally, we have decreased valuations that were initially very high, corresponding to a downward shift in supply and demand in the rightward portion of Figure 1. As can be seen from the diagram, the unique equilibrium price remains at p^* despite these downward shifts.⁵ We now verify that this result holds in general.⁶

⁵As will be clear from Figure 1, we hold valuations fixed within an ϵ -ball of p^* purely to preserve the uniqueness of the competitive equilibrium price. In fact, only one half of the ball is required for this purpose, although we retain the full ball for simplicity of exposition.

⁶All proofs are collected in Appendix A.

Figure 1: A CE preserving shift



Proposition 1 *Let p^* denote the competitive equilibrium price when buyer and seller valuations are distributed according to V_b and V_s respectively. Then if T_b (respectively, T_s) is a competitive equilibrium preserving demand contraction (supply expansion), p^* remains the unique competitive equilibrium price when buyer and seller valuations are distributed according to $V'_b = T_b(V_b)$ and $V'_s = T_s(V_s)$.*

We have thus seen that there is a wide class of downward shifts to the buyer and seller valuations that leave competitive equilibrium predictions unaffected.⁷ Intuitively, one might expect such shifts to depress prices, either because they encourage buyers to offer lower prices or accept lower offers from sellers, or otherwise because they encourage sellers to offer lower prices or accept lower prices from buyers. That such shifts should leave the predictions of competitive equilibrium unchanged may thus come as a surprise. Whether

⁷Unsurprisingly, one can also define an analogous class of *upward* shifts to the buyer and seller valuations that leave equilibria unaffected. More generally, any shift that preserves the ranking of valuations in a neighbourhood of p^* unaffected will preserve the unique equilibrium price p^* (although we focus on everywhere upward or everywhere downward shifts since these generate the most counterintuitive implications).

these counterintuitive predictions are borne out by the data is a topic that we take up in the next section.

3 Experimental design

In order to examine the effect of competitive equilibrium preserving shifts, we ran a series of double auction experiments in Oxford in early June 2022.⁸ The basic idea of the experiment was straightforward. Buyers and sellers were first endowed with their own (private) values and costs using the technique of induced valuations (see Smith (1976) for discussion and elaboration). They then participated in a series of double auctions. In such auctions, buyers may, at any point in time, make ‘bids’ to purchase at a particular price or accept offers that have been made by sellers. Similarly, at any point in time, sellers may offer to sell at a particular price (‘making an ask’) or accept a bid made by a buyer.

We opted to conduct a series of oral double auctions, which means that subjects made bids/asks verbally instead of submitting them electronically. While this resulted in somewhat slower data collection than one would have obtained from a computerised experiment, it yielded several important advantages. First, based on some pilot experiments, it seemed that subjects found oral double auctions more engaging, and also found the structure of oral double auctions rather easier to understand. Second, using oral double auctions ensured that our results were maximally comparable to classic studies like Smith (1962), Smith (1965) and so forth. For this reason, we kept as close as possible to classical experimental economics protocols (see Plott (1991) for a helpful document outlining how such experiments were run and Kimbrough and Smyth (2018) for a more recent experiment that adheres closely to such protocols).

In this initial set of experiments, we employed a queue in order to keep supply and demand schedules stable over time. This meant that every time a trade was executed, a new buyer and seller entered the market with the valuation and cost of the just departed buyer and seller. As discussed in the introduction, we did this for two reasons. First, such a design is arguably more realistic: actual markets do not typically dissolve after several trades have occurred (as in standard double auction experiments). Instead, they are continually replenished by a steady flow of new buyers and sellers. Second and much more importantly, our queue ensured that the set of competitive equilibria remained fixed over time, allowing us to rigorously study whether prices approached *the* competitive equilibrium set. Standard experiments do not

⁸The experiments received IRB Approval from the University of Oxford (ECONCIA21-22-44) and were pre-registered on the AEA registry: <https://www.socialsciceregistry.org/trials/9547>

necessarily have this property.

To implement the queue experimentally, we recruited a group of four buyers and four sellers for every session who began each trading period as ‘inactive’ (so they could not engage in market activity). As trade progressed, buyers and sellers were successively drawn from the queue into the main trading area, and explicitly adopted the value and cost of the buyer and seller who had just departed (by sitting down in their place and inspecting the back of the value/cost card that had been left on the table). This was done in full view of the other experimental participants so as to emphasise that the distribution of the values and costs had remained unchanged.

To study the impact of competitive equilibrium preserving shifts, we used two different treatments, each with five ‘active’ buyers and five ‘active’ sellers (in addition to the eight initially inactive traders in the queue). In the ‘symmetric’ treatment, buyer valuations were £12, £32, £52, £72, £92; and seller valuations were £8, £28, £48, £68, £88. One can check that this yields an interval of competitive equilibrium prices from £48 to £52.⁹ Our second treatment (the ‘low values’ treatment) was obtained by decreasing the valuations and costs in the symmetric treatment as aggressively as possible in a way that preserves the set of competitive equilibrium prices. To this end, we changed the buyer valuations of £12 and £32 to £0 and reduced the buyer valuations of £72 and £92 to £52. Meanwhile, we reduced the seller valuations of £8 and £28 to £0 and reduced the seller valuations of £68 and £58 to £52. This yielded the new vector of buyer and seller valuations, namely £0, £0, £52, £52, £52 for the buyers and £0, £0, £48, £52, £52 for the sellers. One can verify directly that these distributions yield the very same set of equilibrium prices, namely £48 – £52.¹⁰

To control for subject fixed effects, we conducted both treatments sequentially within every experimental session. To get a handle on order effects, we conducted two sessions (Sessions 1 and 2) and varied the order of the treatments within these sessions. We ran the symmetric treatment first in Session 1; and ran the low values treatment first in Session 2. In total, we

⁹The competitive equilibrium prices of £48 and £52 are only *weak* competitive equilibria: at such prices, there exists at least one trader who lacks any strict incentive to act in a way that clears the market. However, prices of £49, £50 and £51 are strict equilibria.

¹⁰Since the low values treatment yields most of the surplus to sellers at the competitive equilibrium, it is reminiscent of the extreme earnings inequality design that has been studied in the literature (see, for example, Smith and Williams (2000)). However, it differs in at least two important respects. First, our design has an equal number of buyers and sellers: in contrast, the classic design generates extreme earnings inequality through an imbalance in the number of buyers and sellers. Second, our design generates an interval of strict competitive equilibrium prices (namely, £49, £50 and £51). In contrast, the standard extreme earnings design leads to the non-existence of strict competitive equilibria (the weak competitive equilibrium is computed by finding the price at which one side of the market is indifferent between trading and not trading).

recruited 18 subjects for each of these initial two experimental sessions (see Table B1 for an overview of all experimental sessions).

At the start of each session, the auction rules were presented in written form (see Appendix C for the rules with which subjects were presented). The rules were then further emphasised through an extensive oral quiz; and subjects were asked if they had any outstanding questions about the auction rules. Finally, subjects were asked to engage in a mock round of trading (which would not be used to calculate payoffs) for didactic purposes. As a result of these measures, subjects’ understanding of the rules seemed to be excellent. As one indication of this, 99.5% of bids and asks made in our experiments were ‘individually rational’ in the sense that they would have made a (non-negative) profit if accepted. Moreover, only 3% of rounds involved a trade that violated individual rationality. This compares favourably to existing auction datasets: for example, the dataset used by Lin et al. (2020) involves individual rationality violations in 90% of rounds (see Arifovic et al. (2022) for discussion).¹¹

Once the illustrative trading round had concluded, the real trading began. Within each round of trading, active buyers and sellers were free to make or accept offers (‘bids’ or ‘asks’) at any time. All offers were repeated by the auctioneer and recorded on a whiteboard.¹² Trading used the standard improvement rule, which meant that bids needed to get successively higher and asks needed to get successively lower until a transaction occurred (at which point everything was reset and all bids and asks became permissible). Each round continued until the queue had been exhausted (i.e. until four trades had occurred); and we conducted five rounds of trading for each of the two sets of value distributions. In line with recent recommendations (Charness et al., 2016; Azrieli et al., 2018) and double auction experiments (Ikica et al., 2018), subjects were only paid for one randomly chosen round within a session.¹³

4 Results

We begin by examining the transactions that occurred in our first two experimental sessions. Panel A of Figure 2 displays the buyer valuation (top line), price (middle line) and

¹¹In our main analysis, we drop the handful of bids and asks that violate individual rationality from our dataset. However, our results are almost entirely unaffected if we include such data points.

¹²To make or accept an offer, a buyer would say (for example) ‘Buyer 2 bids 30’ or ‘Buyer 2 accepts 60’; and this would be duly repeated by the auctioneer. An analogous comment applies to sellers. All offers and acceptances were recorded by a research assistant and double-checked using an audio recording of the experimental sessions.

¹³On average, subjects received £18.36, with a mean absolute deviation of £12.09. In the experiments reported in Section 6, average earnings were £16.98 with a mean absolute deviation of £11.35. In general, sessions took around 1.5 hours, about half an hour of which was devoted to carefully explaining the rules to subjects, with the remaining time devoted to trading.

cost (bottom line) associated with each of the transactions in Session 1. The band of equilibrium prices (£48–£52) is indicated by the dotted lines, and the end of each of a round is indicated by a break. The left half of Panel A shows the transactions from the first half of the experiment (i.e., the symmetric treatment); whereas the right half of Panel A shows the transactions from the second half of the experiment (i.e. the low values treatment). Analogously, Panel B of Figure 2 displays the valuations, prices and costs associated with the transactions from Session 2, which began with the low values treatment and proceeded to the symmetric treatment.

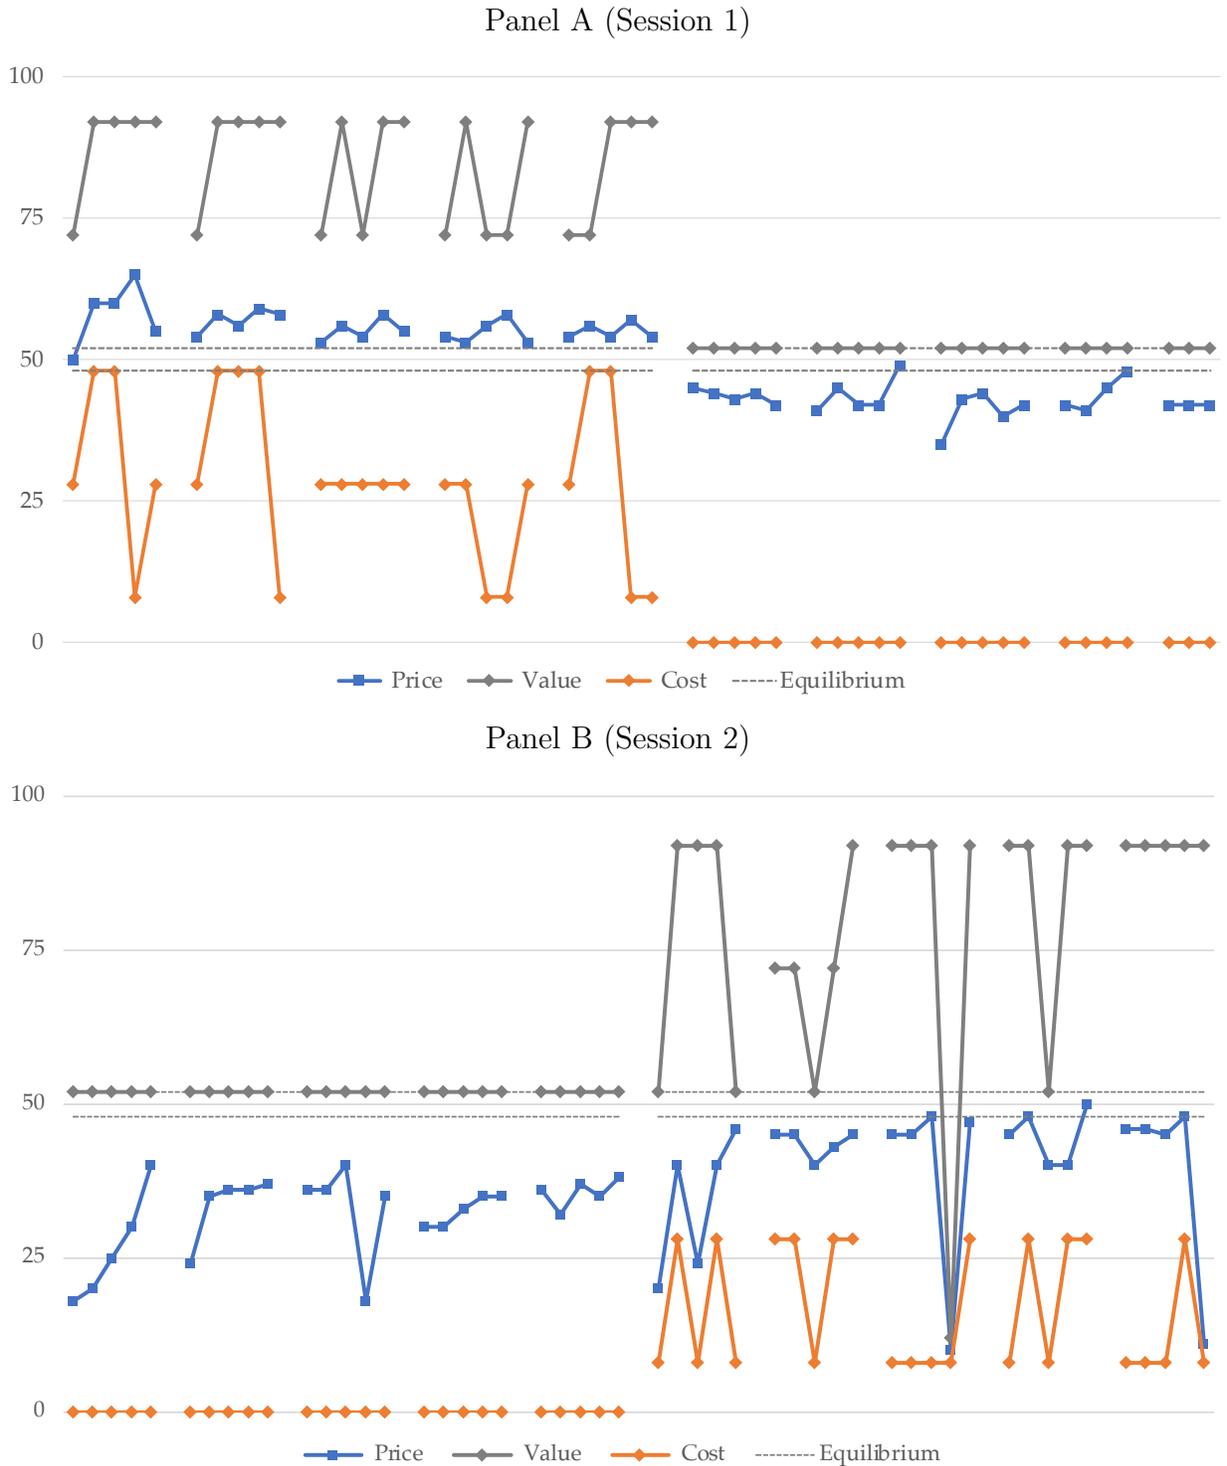
Three results are apparent. First, it is clear that shifting valuations and costs downward lowers observed prices, in violation of competitive equilibrium. Comparing the first halves of the separate sessions — which is perhaps the cleanest comparison since it is uncomplicated by order effects — we see that average prices are £56.0 (in the symmetric treatment) as opposed to £32.3 (in the low value treatment) ($p < 0.0001$).¹⁴ We also see a similar trend within sessions. In the first session, shifting values and costs downward reduces average bids from £56.0 to £42.9 ($p < 0.0001$). In the second session, shifting values and costs upward increased average bids from £32.3 to £40.0 ($p < 0.01$). Therefore, we obtain strong evidence that competitive equilibrium preserving shifts do in fact shift prices, though the effects are substantially larger when comparing across sessions than when comparing within sessions (as one might expect given the ‘price stickiness’ observed in the data).

While Figure 2 solely displays data for the transactions, we can also see a similar pattern when examining the data on bids and asks. Comparing the first halves of the two sessions, we see that average bids/asks are £47.0/£67.9 in the symmetric treatment, as opposed to £26.6/£58.9 in the low values treatment ($p < 0.0001$, $p = 0.03$).¹⁵ In the first session, shifting values downward reduces average bids from £47.0 to £38.0 ($p < 0.01$) and reduces average asks from £67.9 to £48.5 ($p < 0.0001$). In the second session, shifting values up increases average bids from £26.6 to £31.7 ($p = 0.03$) and increases average asks from £59.0 to £153.2 ($p = 0.02$). We therefore conclude that, just as shifting valuations and costs downward reduces observed transaction prices, it also tends to reduce the bids and asks

¹⁴The p -values in this and the next paragraph are generated by unpaired t -tests of the hypothesis of equal means. To account for potential serial correlation in the data, we also regressed the price associated with transaction t on a treatment dummy variable along with the price associated with transaction $t - 1$ (thus allowing each price to be partly driven by the price associated with the previous transaction). As shown by Table B2, this yields similar results.

¹⁵In all our analyses of the bid and ask data, we remove the handful of (rather optimistic) asks that exceed £1,000 (e.g. we drop one participant’s offer to sell for £1 million). Including these asks only strengthens our conclusions.

Figure 2: Valuations, prices, and costs



Notes: This figure reports the buyer valuation (top line), price (middle line), and seller cost (bottom line) associated with each of the transactions. The band of equilibrium prices (£48–£52) is indicated by the dotted lines, and the end of a round is indicated by a break. Panel A displays the transactions from Session 1 (symmetric treatment then low values treatment), whereas Panel B displays the transactions from Session 2 (low values treatment then symmetric treatment).

made by traders.¹⁶

Our second main finding, which partially although not entirely follows from the first, is that prices are almost never at competitive equilibrium. In the first session (again, see Figure 2), prices start consistently above equilibrium, and then fall persistently below it. Moreover, not merely does competitive equilibrium fail as a literal description of the observed prices, but we can also reject a stochastic version of competitive equilibrium that allows for independent errors in every period ($p < 0.0001$, $p < 0.0001$).¹⁷ The prices in our second session are also almost never at equilibrium, but in a different way. Now, prices are persistently below even the lowest competitive equilibrium price. Again, not only can we reject a rather literal interpretation of competitive equilibrium, but we can also reject a stochastic version that allows for independent errors ($p < 0.0001$, $p < 0.01$).

Our third main result is that prices do not seem to be converging to competitive equilibrium over time. In Session 1, there is little indication that prices are trending downwards (in the first half) or upwards (in the second half). Indeed, the average price changes are close to zero (£0.17 and -£0.14 in the first and second half respectively) and neither are statistically different from zero ($p = 0.84$, $p = 0.89$). In Session 2, there is a little more indication of an upward drift in prices, but again this is very weak: average changes are again close to zero (£0.83, £0.24) and again statistically insignificant ($p = 0.61$, $p = 0.94$). Even more strikingly, in neither session does competitive equilibrium seem to be an absorbing state. For example, although the price starts at equilibrium in Session 1 (with a first transaction price of £50), prices quickly drift upwards away from the equilibrium set. Similarly, although prices hit competitive equilibrium briefly in the third round of the second half of Session 2, they again move away from it. Thus, there is little appearance of convergence in about an hour's worth of trading.

In summary, we see that in double auctions with stationary value distributions (a property ensured through our use of a queue), resulting prices can remain far from competitive equilibrium even after long periods of trading, and show little sign of converging to equilibrium. Consistent with this, our competitive equilibrium preserving shifts substantially depress observed prices. In the next section, we turn to the question of what might explain this phenomenon.

¹⁶The statistical tests whose results we report in this paragraph implicitly assume that the data are i.i.d. See Table B3 for analogous results that cluster standard errors at the individual level.

¹⁷To test this, we evaluate the null hypothesis that the price data were i.i.d. draws from a normal distribution with a mean of the closest competitive equilibrium price and a variance to be estimated from the data. Observe that both of these choices — choosing the closest competitive equilibrium price along with allowing the variance to be fit ex post — substantially stacks the deck in competitive equilibrium's favour, which then makes the clear rejection of competitive equilibrium even more striking.

5 Understanding the monotonicity

There are a number of double auction models which can rationalise the monotonicity documented in the previous section (i.e., that lower values/costs generate lower prices). To start with, consider the zero intelligence (ZI) model introduced by Gode and Sunder (1993): buyers bid uniformly between 0 and their valuation, sellers bid uniformly between their valuation and some maximum, and trade occurs when the market bid and market ask cross (at a price equal to the earlier of the two offers). Under such assumptions, decreasing buyer and seller valuations leads to buyer bid distributions and seller ask distributions that are stochastically lower (in the sense of first-order stochastic dominance). As a result, it will tend to decrease observed prices.

To see how this works quantitatively, we conducted an extensive simulation of ZI trading under both of our experimental treatments. To operationalise ZI trading, one needs to specify a maximum ask; and we set this maximum at £100. We also assumed that at every point in time, one trader (either a buyer or seller) is chosen randomly to make an offer; and we then simulated a sequence of 10 million such offers (leading to about 800,000 market prices).¹⁸ The simulation reveals that, as one would expect, average ZI prices in the symmetric treatment are around £50. Meanwhile, average ZI prices in the low values treatment are around £35. Thus, the ZI model predicts that shifting values and costs down in the way done in our experiments should very substantially depress average prices.

While ZI can rationalise the monotonicity we observe in our data, it is doubtful that postulating random bids and asks can be said to explain the source of the monotonicity in any meaningful way. Fortunately, however, such monotonicity is also generated by optimising models. For instance, consider the model developed in Gjerstad and Dickhaut (1998): buyers and sellers choose bids and asks in a way that maximises this period's expected pay-off. Ignoring integer constraints, the optimal bid/ask satisfies a first order condition, inspection of which reveals that optimal bids/asks are increasing in valuation/costs. Thus, the Gjerstad/Dickhaut model also predicts the monotonicity observed in our data.

Finally, we observe that this monotonicity also arises in more complicated optimising models. For example, consider the model of Friedman (1991), in which buyers and sellers optimally choose reservation prices so as to balance the benefit of waiting for better bids/offers against the costs of running out of time. As Friedman remarks (p. 60), optimal reservation prices are monotone in valuations: this means, for example, that buyers with lower valuations are

¹⁸If the value distributions are held fixed as in our experiment, then nothing changes under ZI trading at the conclusion of a round. Thus, it is easiest to simply simulate a very large number of offers (and examine the resulting prices when these offers lead to trade) instead of simulating a large number of rounds.

happy to accept lower offers. As a result, Friedman’s optimal stopping logic also predicts the monotonicity that we observe.

In this section, our goal is not to select one model which is the ‘true’ explanation for the observed monotonicity; and still less to discuss which model can best explain all aspects of double auction experiments (for efforts in this direction, see Cason and Friedman (1996) and Arifovic et al. (2022)). Rather, our goal is simply to argue that the observed monotonicity is nothing very mysterious: indeed, it is a simple consequence of both non-optimising as well as optimising double auction models.

6 The Marshallian path

Although the previous results demonstrate that double auctions need not generate equilibrium prices, one might suspect that this has something to do with the queuing procedure used in order to ensure that the value distribution remains fixed. Indeed, if traders are allowed to drop out without replacement as time progresses, then one may expect that prices approach competitive equilibrium due to a Marshallian path dynamic. While this route to equilibration has been discussed informally (see e.g. Brewer et al. (2002)), we now formalise the dynamic in order to obtain a more rigorous understanding of its properties.

To this end, return to the environment in Section 2, recalling that F and G denote the distribution of buyer and seller valuations respectively, and that p^* denotes the unique competitive equilibrium price. Consider now a sequence of trades, indexed by $t \in [0, T]$. Let $v_b(t)$, $v_s(t)$ and $p(t)$ denote the buyer valuation, seller valuation, and price associated with trade t ; and (with some abuse of notation) denote the corresponding functions by v_b , v_s and p . We can now formalise the concept of a Marshallian path.

Definition 3 *A Marshallian path is a triple (v_b, v_s, p) such that*

1. *For all $t \in [0, T]$, $v_b(t) = F^{-1}(1 - t)$ and $v_s(t) = G^{-1}(t)$.*
2. *For all $i \in [0, 1]$, $i \in [0, T]$ if and only if $F^{-1}(1 - i) \geq G^{-1}(i)$.*
3. *For all $t \in [0, T]$, $v_s(t) \leq p(t) \leq v_b(t)$.*

As stated above, a Marshallian path is a sequence of trades that satisfies three conditions. To understand the first condition, start with the simpler equation $v_s(t) = G^{-1}(t)$ and invert it to get $t = G(v_s(t))$. This says that the fraction of values that are below $v_s(t)$ is t ; so at time t , the valuation of the seller about to engage in trade is at the t -th percentile of the

seller value distribution.¹⁹ Likewise, the other equation reads $v_b(t) = F^{-1}(1 - t)$, which may be inverted to yield $t = 1 - F[v_b(t)]$. This says that the fraction of buyer valuations that are *above* $v_b(t)$ is t ; so at time t , the valuation of the buyer about to engage in trade is at the $(1 - t)$ -th percentile of the buyer value distribution. Taken together, these assumptions say that trade takes place ‘in order’: at every point in time, trade occurs between the buyer with the highest valuation and the seller with the lowest valuation out of those ‘active’ traders who remain in the market.

The second condition says that for any possible unit, that unit is traded if and only if the buyer’s valuation for that unit exceeds the seller’s valuation for that unit (assuming that these units are traded ‘in order’, in line with condition 1). In other words, trades take place if and only if there is some price at which they would be mutually beneficial. We will discuss the plausibility of this claim when we analyse the experimental results of this section.

The final condition says that, in actual fact, trade is mutually beneficial: the price of every trade lies between the valuation of the relevant buyer and the valuation (i.e. ‘cost’) of the relevant seller. Observe that, although there is just one path of trades satisfying conditions 1 and 2, there are multiple mutually beneficial price paths. Thus, as far as prices are concerned, the Marshallian path is indeterminate; which is why we may, for reasons of pedantry, choose to speak of ‘a’ rather than ‘the’ Marshallian path.²⁰

Proposition 2 *If (v_b, v_s, p) is a Marshallian path, then $p(T) = p^*$.*

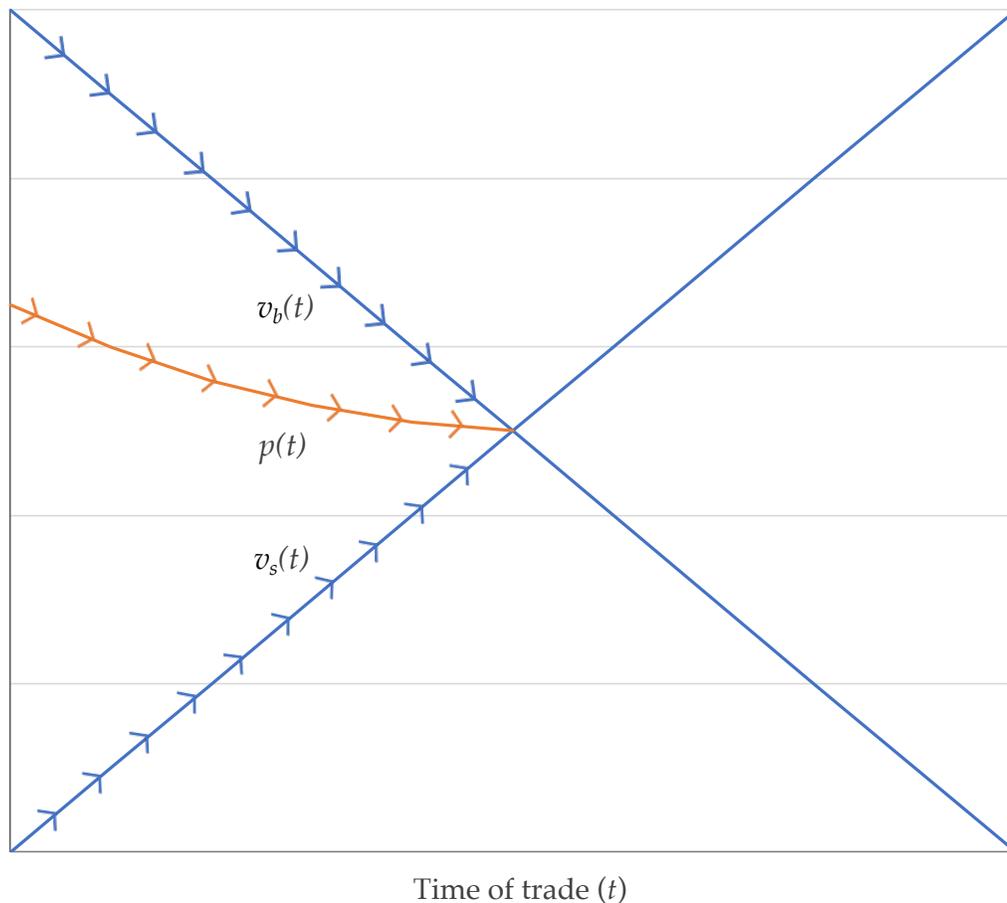
Proposition 2 says that, if trade follows a Marshallian path, then the final trade must be transacted at the equilibrium price. To see why this is true, consider Figure 3. Recall that the demand curve can be interpreted as the valuation of the marginal buyer (once buyers have been sorted from highest valuation to lowest), and the supply curve can be interpreted as the valuation of the marginal seller (assuming sellers have been sorted in order from lowest to highest valuation). As a result, to assume a Marshallian path is simply to assume that trade takes place in the rightward direction, starting with the buyer with the highest valuation and the seller with the lowest cost (condition 1). At every point in time, the price must lie between the buyer and seller valuation (condition 3), which means that the price path must lie between the demand and supply functions. Furthermore, trade occurs if and only if it could be mutually beneficial (condition 2), which means that the final trade is at the intersection of the demand and supply curves. At such a trade, the only mutually

¹⁹Our usage of the word ‘percentile’ differs from the normal usage by a factor of 100: for example, we say ‘0.2-th percentile’ to mean the 20th percentile.

²⁰On the other hand, there is a single path of *trades* which counts as Marshallian, which is presumably why one finds discussion of ‘the’ Marshallian path in prior literature (Brewer et al., 2002; Plott et al., 2013).

beneficial price is $p(T) = p^*$, so we obtain final period equilibration.²¹

Figure 3: A Marshallian path



The previous result establishes that if trade follows a Marshallian path, then final period prices must be at competitive equilibrium. However, it does not address why trade should take place in exactly the order required for a Marshallian path (condition 1). While we will not address this issue in detail, there are two reasons why it might seem reasonable to make this assumption. First, prior literature shows that at least as far as experimental double auctions go, condition 1 is indeed a rough approximation of reality (for instance, see Plott et al. (2013)). Indeed, as we later show, this assumption matches up very well with our own experimental data. Second, condition 1 should arise for precisely the reasons discussed in Section 5: it is, after all, simply an assertion that offers are monotone in valuations.

²¹Incidentally, it is unclear whether this argument for equilibration actually appears anywhere in the work of Marshall. Marshall (1890) does present a theory of equilibration (see Book V, Chapter 3, Section 6), but this theory involves a quantity adjustment at disequilibrium prices. At the risk of giving Marshall credit for yet another idea that he did not originate (see Humphrey (1996)), we retain the phrase ‘Marshallian path’ in deference to previous literature on this issue (Brewer et al., 2002; Plott et al., 2013).

We should perhaps also emphasise that this result does *not* extend to settings with stationary value distributions (like those studied experimentally in Section 4). In such settings, Marshallian path logic would suggest that trade should take place between the buyer with the highest value (here, £52) and the seller with the lowest cost (here, £0): and indeed this is precisely what we observe in the low values treatment.²² However, as far as the logic of Proposition 2 goes, there is no reason why they should trade at equilibrium prices if they are continually replaced once they leave the market. Moreover, under anything like an equal division of the surplus, we would expect prices to be considerably below equilibrium (as documented in Section 4). Thus, not only does this result provide us with a reason to expect final period equilibration, but it provides us with a reason that is conspicuously absent in the case of stationary value distributions.

Motivated by our result, we now report the results of a series of additional experimental sessions held in order to determine whether the Marshallian path is indeed crucial for equilibrium convergence. These experiments followed exactly the same procedure as those outlined in Section 3, with the exception that we removed the queue of buyers and sellers (so traders could drop out without replacement as time progressed). As in classical double auction experiments, each round now concluded once there were no longer any active traders who wanted to make or accept an offer.²³

As before, we ran two sessions (Sessions 3 and 4), varying the order in which the sessions were presented. Session 3 began with the symmetric treatment, whereas Session 4 began with the low values treatment. Panels A and B of Figure 4 display the main results for Session 3 and Session 4 respectively. As usual, the buyer valuation, price, and seller cost associated with each transaction are denoted by the top, middle, and bottom lines respectively. The conclusion of a round is indicated by a break; and the set of equilibrium prices is indicated by the dotted lines.

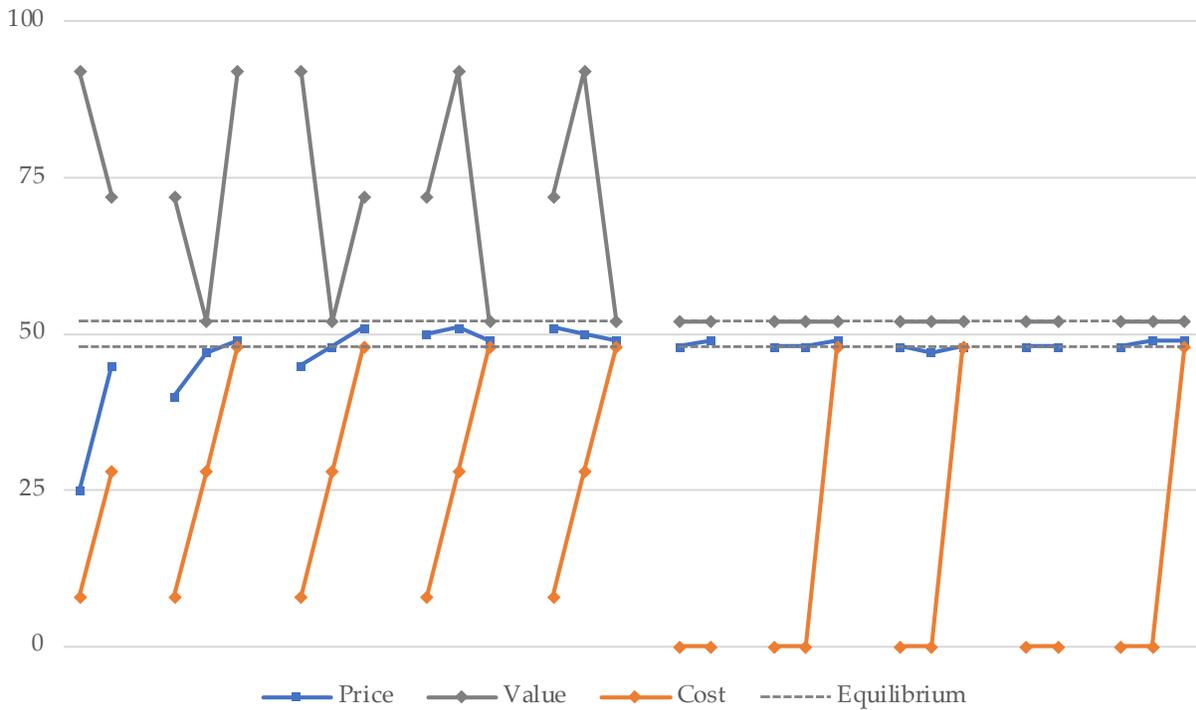
Two features of the panels are especially noteworthy. First, we still obtain evidence that our shifts depress prices, although the evidence is now more mixed. Comparing across the first halves of each session, which is again the cleanest comparison since it is uncomplicated by order effects, we see that average bids are £46.4 in the symmetric treatment as opposed to £35.3 in the low values treatment ($p < 0.001$). Thus, shifting the values downward appears to substantially depress prices, at least in the short run. Looking within periods, we see that

²²This phenomenon is also observed, although a little less strikingly, in the symmetric treatments.

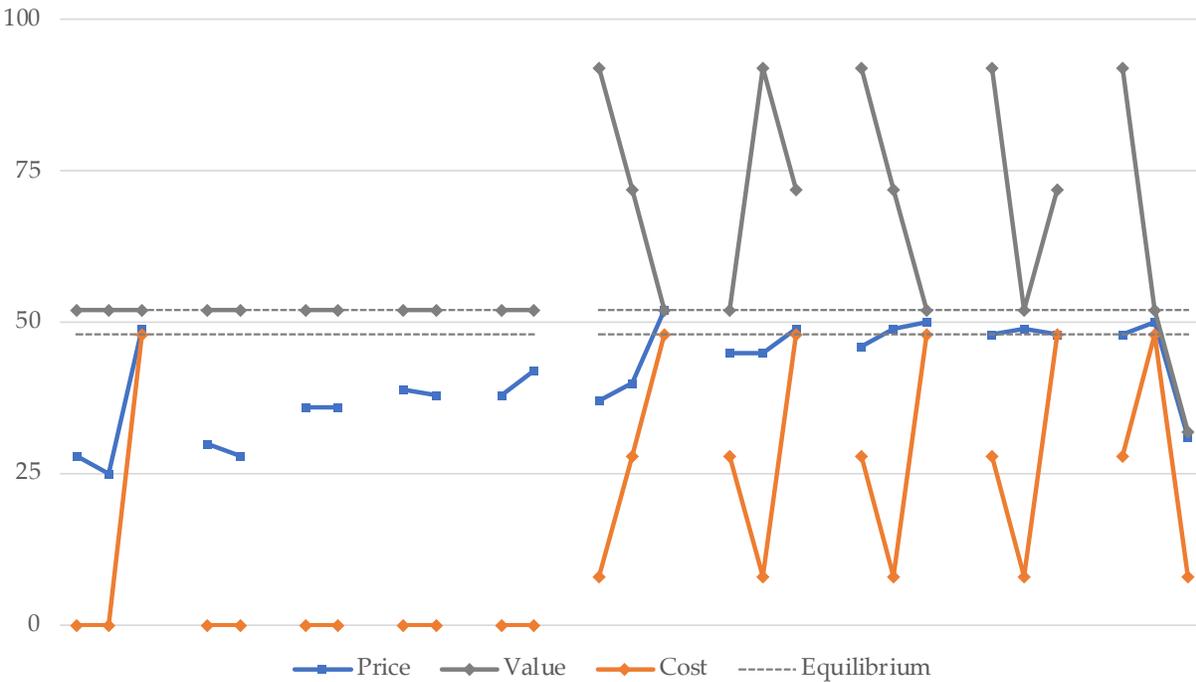
²³Similarly to Kimbrough and Smyth (2018), the auctioneer asked: ‘Sellers, would any of you like to make any offers or accept the current market bid? Buyers, would any of you like to make any bids or accept the current market ask? Going once... going twice... the round has concluded’. See Appendix C for the full instructions for the experiment.

Figure 4: Valuations, prices, and costs

Panel A (Session 3)



Panel B (Session 4)



Notes: This figure reports the buyer valuation (top line), price (middle line), and seller cost (bottom line) associated with each of the transactions. The band of equilibrium prices (£48–£52) is indicated by the dotted lines, and the end of a round is indicated by a break. Panel A displays the transactions from Session 3 (symmetric treatment then low values treatment), whereas Panel B displays the transactions from Session 4 (low values treatment then symmetric treatment).

shifting the values and costs upward increased the average bids in Session 4, from £35.4 to £45.8 ($p < 0.001$). On the other hand, no such effect is observed within Session 3: average prices are actually a little higher in the low value treatment, although the difference is not statistically significant ($p = 0.36$).

Our second finding, which again is related to the first, is that the double auctions now demonstrate strong equilibrating tendencies. In Session 3, prices reach equilibrium by the last trade of the second round of the symmetric treatment and remain within the band of equilibrium prices for the rest of the session (even after valuations are shifted downward). In Session 4, prices failed to equilibrate despite half an hour's trading in the low values treatment. However, they do show a clear upward drift, and we do obtain an equilibrium price for one transaction. Moreover, switching to the symmetric treatment leads prices to more or less equilibrate by the end of the second round.

Given the previous findings, it is natural to conjecture that the low values treatment might equilibrate (even without the help of the symmetric treatment) if it were simply given enough time to do so. To check this conjecture, we ran two further experimental sessions (Sessions 5 and 6) which solely studied the low value treatment.²⁴ To give our auctions the best chance of equilibration, we conducted nine rounds within each session (which corresponded to around one hour of trading time per session). To prevent subjects with unfavourable valuations from becoming bored by their lack of profitable trading opportunities, we re-shuffled the values and costs twice (once at the beginning of the fourth round, and once at the beginning of the seventh). We did not explicitly inform subjects that the distribution of values and costs had remained the same, though we also gave them no indication that it had changed.

Panels A and B of Figure 5 display the results of our final two sessions. Several features of the data are evident. First, trade very clearly follows the order assumed by the Marshallian path. We find that trade follows the exact Marshallian order in 17 out of the 18 rounds (the corresponding figure is 47 out of 48 rounds if one includes the low valuation treatments from Sessions 3 and 4). These results are consistent with previous evidence on the order of trade within double auctions (see, e.g., Cason and Friedman (1996), Plott et al. (2013), Lin et al. (2020) and Sherstyuk et al. (2021)).²⁵

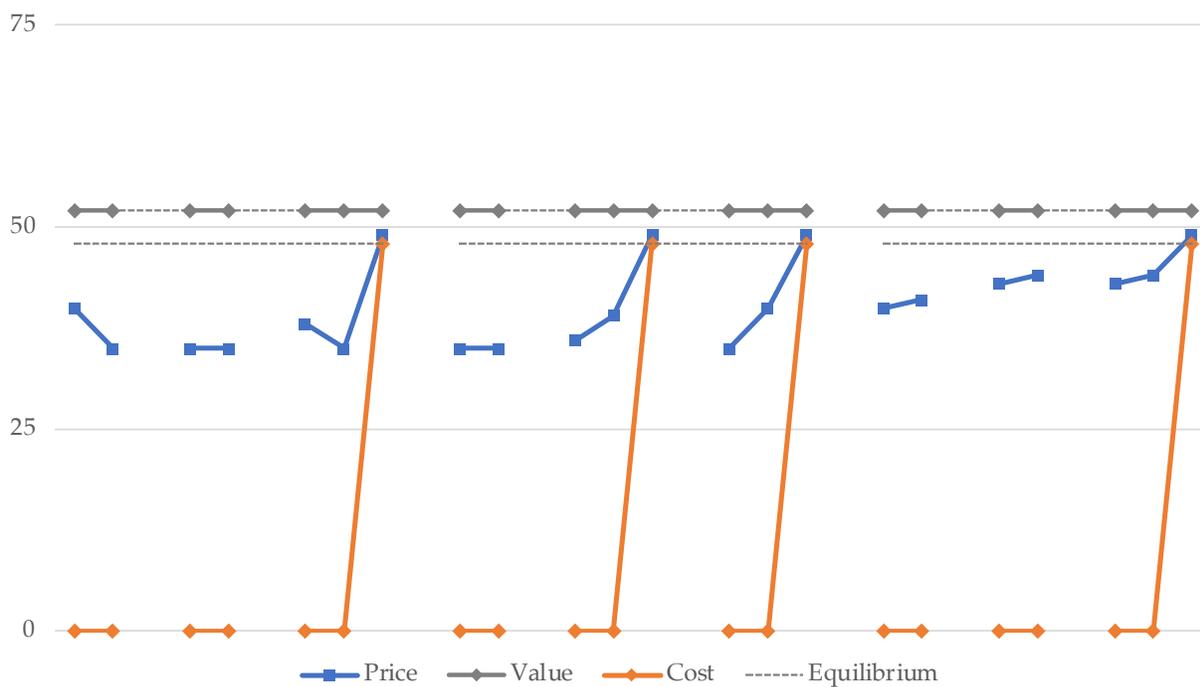
Second, we observe striking evidence of third period equilibration. As predicted by our theory, *all* third trades occur at a competitive equilibrium price. Moreover, whilst not all

²⁴These experimental sessions were not contained in our original pre-registration.

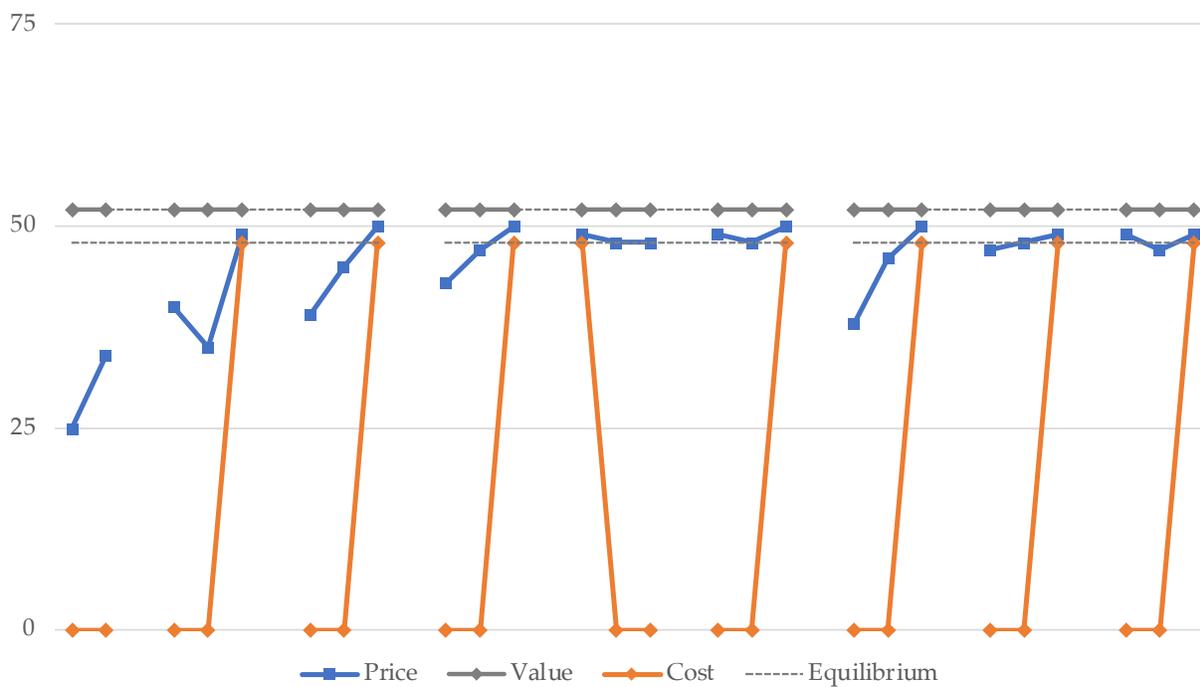
²⁵While other studies observe a similar phenomenon in terms of trading order, their results on trading order are rather less pronounced. Presumably, this is due to the fact that the two highest valuations are the same, and the two lowest costs are also the same, making it much easier to obtain a Marshallian trading order.

Figure 5: Valuations, prices, and costs

Panel A (Session 5)



Panel B (Session 6)



Notes: This figure reports the buyer valuation (top line), price (middle line), and seller cost (bottom line) associated with each of the transactions. The band of equilibrium prices (£48–£52) is indicated by the dotted lines, and the end of a round is indicated by a break. Panel A displays the transactions from Session 5, whereas Panel B displays the transactions from Session 6.

rounds lead to a third trade, a good portion of them do. As a result, we obtain final period equilibration in 12 out of 18 of the rounds (this corresponding figure is 16 out of 38 rounds if one includes the low values treatments from Sessions 3 and 4). Consistent with our theory, prices are much less likely to be at equilibrium for non-third period trades. For example, in Session 5, we *never* observe equilibrium prices for first and second period trades.

Third, we also obtain evidence that equilibration is possible even for non-final periods. This is certainly not evident in Session 5, where (as noted) we obtain equilibrium prices only for third period trades. On the other hand, in Session 6, we obtain equilibrium prices for all trades in the fifth, sixth, eighth and ninth rounds. We discuss why this may be below.

Altogether, our results are highly consistent with the theoretical result presented above. In essentially all periods, the first trades happen between buyers with the highest valuations (here £52) and sellers with the lowest cost (here £0): this validates condition 1 of the Marshallian path. Once these two trades have been executed, the only remaining mutually beneficial traders are a buyer with a valuation of £52 and a seller with a cost of £48. Thus, if all mutually beneficial trades are going to occur, a deal must be struck in the equilibrium price range of £48 – £52. This is exactly what we see in our data (again, see Figure 5): it is as if a force, namely a large jump in the valuation of the marginal seller, is ‘squeezing’ the realised price into the equilibrium set.

While our results are broadly consistent with our theory, they do leave at least two unanswered questions. The first question is why not all mutually beneficial trades are made within certain rounds. From a mechanical point of view, one can see that the obstacle is essentially on the buyers’ side. In all but one of the third period trades, seller offers fall to the equilibrium range, and it is the market bid which potentially fails to rise high enough for a deal to be struck (see Figure B1 for the evolution of the market bids and asks in such situations). The question is then why third period buyers with a valuation of £52 refuse to trade at equilibrium prices, thereby ‘leaving money on the table’.

While we do not attempt to definitively answer this question, at least two explanations are suggested by comments made by buyers in post-experimental discussion. First, buyers may have had a strategic motive: by refusing to pay equilibrium prices, they may have been attempting to signal a reluctance to pay such prices in future rounds, thereby encouraging lower future offers from sellers.²⁶ Second, buyers may have refused to accept sellers’ offers of equilibrium prices on the mistaken belief that such offers would deliver almost all the surplus to the relevant seller. In rounds that generated only two trades, prices tended to be in the

²⁶Ironically, this would be a reason why repetition of rounds might actually hinder equilibration: such an effect would not be present in a single round double auction.

£35 – £40 range, with an average of about £37. As a result, buyers might have reasonably formed the expectation that sellers' values are rather low (about £22 under the assumption of equal surplus sharing) and therefore become quite frustrated when the remaining seller started to demand competitive equilibrium prices. In such a situation, a buyer might be willing to forfeit a few pounds in profit in return for the pleasure of punishing what they take to be greediness on a seller's part (see Rabin (1993) and Fehr and Schmidt (1999) for formal models in this vein).

A final question is why prices can potentially become 'stuck' at equilibrium, even for non-third period trades. For concreteness, let us consider Session 6 (see panel B of Figure 5) since this phenomenon does not arise in Session 5. By Marshallian path logic, one would expect all third period trades to be at equilibrium, and indeed this is precisely what we observe. However, there is no special reason, at least so far as the logic of Proposition 2 goes, to expect equilibrium prices for earlier trades; and under anything like equal surplus sharing, one would expect the realised prices to be substantially lower. Strikingly, however, we see that once prices have been pushed to equilibrium repeatedly by the force of the Marshallian path, they may persist at equilibrium (or may not, as in Session 5). The question is then what might explain this.

While we do not attempt to answer this question formally, it is not too hard to see how the Marshallian path logic might extend to trades that do not occur in the final period. Suppose that the first two trades occur at reasonably low prices, say £35, and the third period price jumps (by the logic of Proposition 2) to a competitive equilibrium price. Over time, traders should notice this pattern — and indeed many traders reported noticing this in post-experimental discussions. Given that they expect the price to rise in the final period, buyers should be especially keen to trade earlier than their competitors. This in turn should induce them to submit higher bids and to accept higher asks for early period offers. Similarly, given that the price is expected to jump in the third period, sellers should be keen to trade late, which should induce them to only offer very high asks or accept very high bids in the early periods. Taken together, these forces explain how the final period equilibration delivered by the Marshallian path might eventually drag earlier round prices up to equilibrium levels.²⁷

²⁷Importantly, a similar argument can apply even in rounds which lack a third period trade. Such third periods typically involve a substantial increase in buyer bids and seller asks relative to previous periods. As a result, even those periods that do not lead to a trade can facilitate equilibration.

7 Concluding remarks

In this paper, we revisit the link between competitive equilibrium and the double auction. We begin by showing that competitive equilibrium generates a set of highly counterintuitive predictions. Specifically, it predicts that prices can remain unchanged following a certain class of downward shifts to buyer and seller valuations. We find that in double auctions with stationary value distributions, these predictions are strongly falsified; and more generally that competitive equilibrium is a poor description of outcomes in such auctions. On the other hand, we also find that the effectiveness of our counterexamples is blunted in double auctions in which traders may drop out (without replacement) as time progresses. Taken together, this pair of findings imply that whether value distributions are held stationary is a crucial determinant of whether prices converge to competitive equilibrium; which in turn suggests that the Marshallian path is a key driver of equilibration in double auctions.

Despite the long history of double auction research, we believe that our findings open up several new avenues for investigation. First, further experimental work on the potential importance of the stationarity of the value distribution is clearly needed. Indeed, we are aware of only one other paper on this issue (Brewer et al., 2002); and that paper reaches the rather contrasting conclusion that double auctions can converge even with stationary value distributions.²⁸ Given the paucity of studies on this issue and the apparently mixed nature of the evidence, it is vital that further experiments are conducted to verify whether the issue of stationarity is indeed as critical as we claim.

On the theoretical front, it may also be useful to explicitly develop models of double auction behaviour in stationary environments. Classical double auction models, e.g. Easley and Ledyard (1993), Friedman (1991) and Gjerstad and Dickhaut (1998), explicitly address the situation in which traders drop out without replacement as time progresses, and it would be interesting to investigate the extent to which their insights carry over to the stationary setting. Indeed, the stationary setting would appear to be rather more tractable from a modelling perspective, since it may be viewed (at least very roughly) as a repeated version of the much simpler k -double auction. As a result, the insights of Chatterjee and Samuelson (1983), Jackson and Swinkels (2005), Reny and Perry (2006), Fudenberg et al. (2007) and others may prove highly relevant.

²⁸In part, the apparent difference between this paper and ours may be somewhat illusory. Examining Figures 7 – 9 in Brewer et al. (2002), one sees that prices are persistently above equilibrium, and then (following the shifts) persistently below equilibrium. While we would not classify this as ‘convergence to equilibrium’, it does count as convergence under the relatively undemanding definition used by Brewer et al. (2002) (see p. 191 for details).

Finally, it might be valuable to further study what precisely hinders equilibration in the ‘low values treatment’ investigated in this paper. We have informally sketched two alternatives: that buyers reject equilibrium offers on the (mistaken) belief that they allocate almost all the surplus to sellers, or otherwise that buyers reject these offers strategically as a means of generating higher offers in subsequent rounds. While we have not attempted to learn which of these is the key driver, this would seem straightforward to check experimentally. For example, the second channel is possible only when rounds are repeated, so can be easily ‘turned off’ by conducting double auctions with a single incentivised round.

References

- Arifovic, J., Donmez, A., Ledyard, J., and Tjandrasuwita, M. (2022). Individual evolutionary learning and zero-intelligence in the continuous double auction. *Unpublished manuscript*.
- Arrow, K. J. (1951). An extension of the basic theorems of classical welfare economics. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 2, pages 507–533. University of California Press.
- Aumann, R. J. (1964). Markets with a continuum of traders. *Econometrica*, 32:39–50.
- Azrieli, Y., Chambers, C. P., and Healy, P. J. (2018). Incentives in experiments: A theoretical analysis. *Journal of Political Economy*, 126(4):1472–1503.
- Becker, G. S. (1957). *The economics of discrimination*. University of Chicago Press.
- Becker, G. S. (1973). A theory of marriage: Part I. *Journal of Political Economy*, 81(4):813–846.
- Becker, G. S. (1974). A theory of marriage: Part II. *Journal of Political Economy*, 82(2, Part 2):S11–S26.
- Brewer, P. J., Huang, M., Nelson, B., and Plott, C. R. (2002). On the behavioral foundations of the law of supply and demand: Human convergence and robot randomness. *Experimental Economics*, 5(3):179–208.
- Cason, T. N. and Friedman, D. (1996). Price formation in double auction markets. *Journal of Economic Dynamics and Control*, 20(8):1307–1337.
- Charness, G., Gneezy, U., and Halladay, B. (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131:141–150.
- Chatterjee, K. and Samuelson, W. (1983). Bargaining under incomplete information. *Operations Research*, 31(5):835–851.
- Cripps, M. W. and Swinkels, J. M. (2006). Efficiency of large double auctions. *Econometrica*, 74(1):47–92.
- Debreu, G. (1954). Valuation equilibrium and Pareto optimum. *Proceedings of the National Academy of Sciences*, 40(7):588–592.

- Easley, D. and Ledyard, J. O. (1993). Theories of price formation and exchange in double oral auctions. In *The Double Auction Market Institutions, Theories, and Evidence*, pages 63–98. Routledge.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868.
- Fisher, I. (1910). *Introduction to economic science*. Macmillan.
- Friedman, D. (1991). A simple testable model of double auction markets. *Journal of Economic Behavior & Organization*, 15(1):47–70.
- Fudenberg, D., Mobius, M., and Szeidl, A. (2007). Existence of equilibrium in large double auctions. *Journal of Economic theory*, 133(1):550–567.
- Gjerstad, S. and Dickhaut, J. (1998). Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29.
- Glaeser, E. L. (2007). The economics approach to cities. Technical report, National Bureau of Economic Research.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.
- Holt, C. A., Langan, L. W., and Villamil, A. P. (1986). Market power in oral double auctions. *Economic Inquiry*, 24(1):107–123.
- Humphrey, T. (1996). Marshallian cross diagrams and their uses before alfred marshall: The origins of supply and demand. In *Alfred Marshall: Critical Assessments*, volume 8, page 224. Taylor & Francis.
- Ikica, B., Jantschgi, S., Nax, H. H., Nunez Duran, D., and Pradelski, B. (2018). Competitive market behavior: Convergence and asymmetry in the experimental double auction. *Available at SSRN 3131004*.
- Jackson, M. O. and Swinkels, J. M. (2005). Existence of equilibrium in single and double private value auctions 1. *Econometrica*, 73(1):93–139.
- Kachelmeier, S. J. and Shehata, M. (1992). Culture and competition: a laboratory market comparison between China and the West. *Journal of Economic Behavior & Organization*, 19(2):145–168.

- Kimbrough, E. O. and Smyth, A. (2018). Testing the boundaries of the double auction: The effects of complete information and market power. *Journal of Economic Behavior & Organization*, 150:372–396.
- Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370.
- Lin, P.-H., Brown, A. L., Imai, T., Wang, J. T.-y., Wang, S. W., and Camerer, C. F. (2020). Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments. *Nature Human Behaviour*, 4(9):917–927.
- Lucas, R. E. (1977). Understanding business cycles. In *Carnegie Rochester Conference Series on Public Policy (Vol. 5)*. Oxford University Press, Oxford.
- Marshall, A. (1890). Principles of economics: An introductory volume.
- Plott, C., Roy, N., and Tong, B. (2013). Marshall and walras, disequilibrium trades and the dynamics of equilibration in the continuous double auction market. *Journal of Economic Behavior & Organization*, 94:190–205.
- Plott, C. R. (1981). Theories of industrial organization as explanations of experimental market behavior. *Social Science Working Paper, 388. California Institute of Technology*.
- Plott, C. R. (1991). An introduction to laboratory experimental procedures for political economy. *Unpublished manuscript*.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5):1281–1302.
- Reny, P. J. and Perry, M. (2006). Toward a strategic foundation for rational expectations equilibrium. *Econometrica*, 74(5):1231–1269.
- Sherstyuk, K., Phankitnirundorn, K., and Roberts, M. J. (2021). Randomized double auctions: gains from trade, trader roles, and price discovery. *Experimental Economics*, 24(4):1325–1364.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–137.
- Smith, V. L. (1965). Experimental auction markets and the Walrasian hypothesis. *Journal of Political Economy*, 73(4):387–393.

Smith, V. L. (1976). Experimental economics: Induced value theory. *The American Economic Review*, 66(2):274–279.

Smith, V. L. and Williams, A. W. (2000). The boundaries of competitive price theory: Convergence, expectations, and transaction costs. In *Bargaining and market behavior: essays in experimental economics*, pages 286–319. Cambridge University Press.

Appendices

A Proofs

Proof of Proposition 1. We begin by arguing that p^* remains *an* equilibrium price when valuations are distributed according to $T_b(V_b) = V'_b$ and $T_s(V_s) = V'_s$. To this end, define $\mathcal{B} = (p^* - \epsilon^-, p^* + \epsilon^+)$ and fix some $p \in \mathcal{B}$. By the law of total probability,

$$\begin{aligned} P(V'_b \leq p) &= P(V'_b \leq p | V_b \leq p^* - \epsilon^-)P(V_b \leq p^* - \epsilon^-) + P(V'_b \leq p | V_b \in \mathcal{B})P(V_b \in \mathcal{B}) \\ &\quad + P(V'_b \leq p | V_b \geq p^* + \epsilon^+)P(V_b \geq p^* + \epsilon^+) \end{aligned} \quad (3)$$

Since $T_b(V_b) = V'_b$, where T_b is a CE preserving demand contraction,

- If $V_b \leq p^* - \epsilon^-$, then $V'_b \leq V_b \leq p^* - \epsilon^- \leq p$, so $P(V'_b \leq p | V_b \leq p^* - \epsilon^-) = 1$
- If $V_b \in \mathcal{B}$, then $V'_b = V_b$ and so $P(V'_b \leq p | V_b \in \mathcal{B}) = P(V_b \leq p | V_b \in \mathcal{B})$
- If $V_b \geq p^* + \epsilon^+$, then $V'_b \geq p^* + \epsilon^+ > p$, so $P(V'_b \leq p | V_b \geq p^* + \epsilon^+) = 0$

Inserting these equalities into (3), we obtain

$$\begin{aligned} P(V'_b \leq p) &= P(V_b \leq p^* - \epsilon^-) + P(V_b \leq p | V_b \in \mathcal{B})P(V_b \in \mathcal{B}) \\ &= P(V_b \leq p^* - \epsilon^-) + P(V_b \leq p \wedge V_b \in \mathcal{B}) \\ &= P(V_b \leq p^* - \epsilon^-) + P(p^* - \epsilon^- < V_b \leq p) \\ &= P(V_b \leq p) \end{aligned} \quad (4)$$

Let F' and G' denote the distributions of V'_b and V'_s ; let $d'(p)$ and $s'(p)$ denote the demand and supply functions generated by these distributions; and define $e'(p) \equiv d'(p) - s'(p)$.

The argument just given reveals that $F'(p) = F(p)$ for any $p \in \mathcal{B}$. Therefore, $d'(p) = 1 - F'(p) = 1 - F(p) = d(p)$ (for $p \in \mathcal{B}$). In particular, $d'(p^*) = d(p^*)$. Likewise, if we replace F with G in the previous argument, we see that $G'(p) = G(p)$ for any $p \in \mathcal{B}$. Hence, $s'(p) = G'(p) = G(p) = s(p)$ at such prices; and in particular, $s'(p^*) = s(p^*)$. Since $d(p^*) = s(p^*)$ by definition, this means that $d'(p^*) = s'(p^*)$, i.e. p^* remains a CE.

To see that the CE remains unique, fix some $p < p^*$ and consider the following cases:

Case 1: $p \in \mathcal{B}$. Then $e'(p) = e(p)$ (shown above), and $e(p) > e(p^*) = 0$ (since $p < p^*$ and e is strictly decreasing). This means that $e'(p) > 0$, i.e. no such prices clear the market.

Case 2: $p \notin \mathcal{B}$. Since $p < p^*$, this means that $p \leq p^* - \epsilon^-$. Now define $\bar{p} \equiv p^* - 0.5\epsilon^-$. Plainly, $\bar{p} \in \mathcal{B}$ and $p < p^*$, so $e'(\bar{p}) = e(\bar{p}) > 0$. Moreover, since $e'(p) = 1 - F'(p) - G'(p)$ where F' and G' are CDFs, e' is weakly decreasing over \mathbb{R}^+ . Thus, if $p \leq p^* - \epsilon^- < \bar{p}$, then $e'(p) \geq e'(\bar{p}) = e(\bar{p}) > 0$, so no such prices can clear the market either.

Since these cases exhaust the possibilities, we see that no prices $p < p^*$ can be CE. By an analogous argument, one can also rule out prices $p > p^*$. We conclude that p^* must remain the unique competitive equilibrium price. \square

Proof of Proposition 2. First, we show that $v_b(T) = v_s(T)$. To see this, recall condition 2 of the Marshallian path definition: for every $i \in [0, 1]$, $i \in [0, T]$ if and only if $F^{-1}(1 - i) \geq G^{-1}(i)$. Define $\phi(i) = F^{-1}(1 - i) - G^{-1}(i)$, so our inequality is $\phi(i) \geq 0$. Observe that $\phi(0) = F^{-1}(1) - G^{-1}(0) = \bar{v}_b$ and $\phi(1) = F^{-1}(0) - G^{-1}(1) = -\bar{v}_s$. Also, ϕ is continuous and strictly decreasing in i . Therefore, letting i^* denote the unique root of ϕ , the set of $i \in [0, 1]$ such that $\phi(i) \geq 0$ is the set $[0, i^*]$. Hence, $T = i^*$ and so $\phi(T) = 0$, or $F^{-1}(1 - T) = G^{-1}(T)$. Using condition 1, this finally yields $v_b(T) = v_s(T)$.

By condition 3, we see that $v_b(T) \geq p(T) \geq v_s(T)$. However, $v_b(T) = v_s(T)$. Therefore, $p(T) = v_b(T) = v_s(T)$. Finally, invert condition 1 to get $F[v_B(T)] = 1 - T$ and $G[p(T)] = T$. Using $p(T) = v_b(T) = v_s(T)$, we infer that $F[p(T)] = 1 - T$ and $G[p(T)] = T$, and so

$$G[p(T)] = 1 - F[p(T)] \tag{5}$$

But this is precisely the equation whose unique solution is p^* ! Hence, $p(T) = p^*$. \square

B Additional tables and figures

Table B1: Overview of the experimental sessions

Session	Queue?	Treatments	Rounds	Participants
1	Yes	Symmetric then low values	10	18
2	Yes	Low values then symmetric	10	18
3	No	Symmetric then low values	10	10
4	No	Low values then symmetric	10	10
5	No	Low values only	9	10
6	No	Low values only	9	10

Notes: This table describes the differences between our experimental treatments. The second column specifies whether a treatment used a queue of buyers and sellers, and the third column specifies which treatments were run. The fourth and fifth columns specify the number of rounds and participants in the session.

Table B2: The effect of shifting values on prices

	Session 1	Session 2	Session 3	Session 4
Low values treatment	-13.69*** [2.077]	-7.401*** [2.364]	-0.254 [0.818]	-10.21*** [3.474]
Lagged price	-0.0243 [0.145]	-0.0215 [0.148]	0.204*** [0.057]	-0.0435 [0.301]
Constant	57.61*** [7.966]	40.97*** [5.470]	38.64*** [2.800]	47.82*** [14.140]
n	46	49	26	25
R^2	0.853	0.143	0.205	0.388

Notes: This table reports the results of regressing observed transaction prices on a dummy variable indicating treatment assignment (low values treatment vs symmetric treatment) and the price of the previous transaction. Regressions are conducted separately for each session; and robust standard errors are reported in parentheses (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

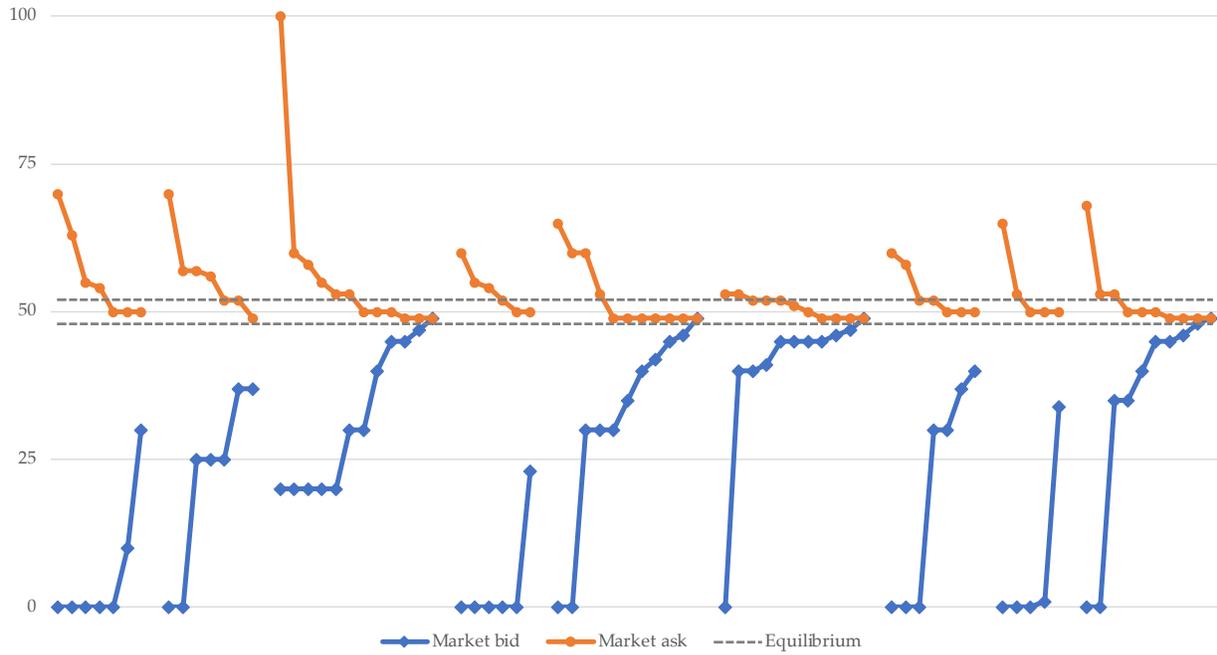
Table B3: The effect of shifting values on offers

	Sessions 1 and 2		Sessions 3 and 4	
	Bids	Asks	Bids	Asks
Low values treatment	-6.782** [3.227]	-53.60** [23.19]	-2.318 [4.765]	-19.85 [14.96]
Session	-13.56*** [3.346]	48.58* [25.49]	-5.796 [5.010]	26.24 [16.18]
Constant	59.71*** [5.401]	33.27 [21.67]	54.15*** [17.49]	-2.888 [51.10]
n	174	166	297	243
R^2	0.305	0.073	0.056	0.085

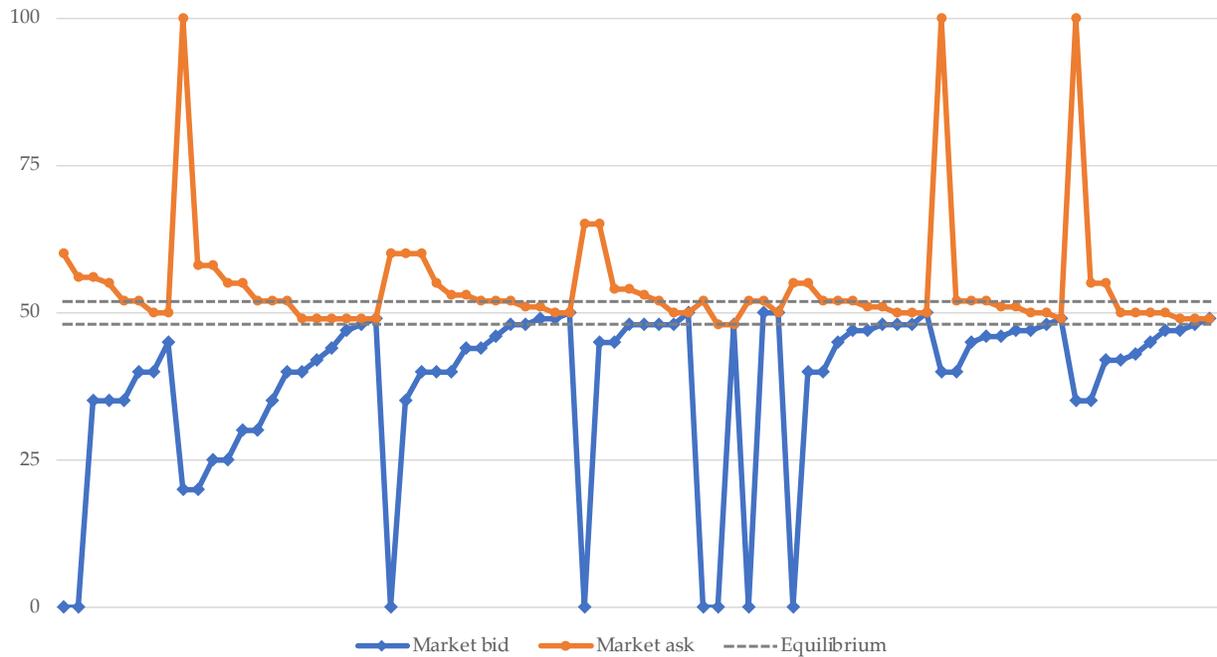
Notes: This table reports the results of regressing bids/asks on a dummy variable indicating treatment assignment (low values treatment vs symmetric treatment), along with a session fixed effect. The first two columns report the results for Sessions 1 and 2 (stationary value distributions); and the next two columns report the results for Sessions 3 and 4 (non-stationary value distributions). Robust standard errors are clustered at the individual level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Figure B1: Market bids and asks

Panel A (Session 5)



Panel B (Session 6)



Notes: This figure reports the market ask and bid in negotiation sequences following a second trade. For simplicity, we discard market asks above £100 and set the market ask at £100 if no asks have yet been made. Similarly, we set the market bid at £0 in the absence of any bids. When a buyer accepts the market ask, we set the market bid equal to the market ask (and likewise for seller acceptances of the market bid). Thus, trade occurs when the market bid and ask meet. The end of a round is denoted by a break.

C Experimental instructions

C.1 Instruction for Buyers (Sessions 1 and 2)

Welcome to the experiment! Please read the following instructions as carefully as possible.

Preliminaries

- Please do not talk to your fellow participants at any stage. Talking may result in a loss of experimental earnings.
- You will have just received an ID card specifying your buyer number. B1 means ‘Buyer 1’, B2 means ‘Buyer 2’, and so forth. Please wear your ID card visibly at all times.

Roles

- There are two types of buyers in this experiment: active buyers and pending buyers. If your buyer number is between 1 and 5 inclusive, then you are an active buyer and will be sitting in the main trading area. On the other hand, if your buyer number is 6 or higher, then you are a pending buyer and will be sitting in the queue.
- Active buyers are free to trade from the very start of a trading period. However, pending buyers may only trade after active buyers have made a trade and dropped out of the market (see elaboration below).

Valuations

- If you are an active buyer, you will have a card marked ‘Valuation’ in front of you. Examine the number on the back of this card, being careful not to let any other participants see this number. This number is your ‘valuation’ for the fictitious commodity to be traded. Memorise your valuation, turn your ID card back face-down (so the valuation is hidden), and do not reveal your valuation to anybody else!
- If you are a pending buyer, then you have not yet been assigned a valuation. However, you will acquire a valuation as soon as an active buyer has dropped out and you have taken that buyer’s place (and valuation).
- If you manage to make a trade, you will receive your valuation minus the price that you agreed to pay (this assumes that the trade ‘counts’ — see discussion below). For example, if your valuation is £55 and you agree to a price of £20, your net earnings are £35.
- Notice that your valuation represents the most that you should be conceivably willing

to pay to make a trade. For example, if your valuation is £40, then you should never pay more than £40 to buy a unit: doing so would just lose you money!

Trading

- At any point, active buyers may offer to buy at a particular price – this is called making a ‘bid’. To make a bid, raise your hand. Once the auctioneer has pointed at you, state your identity along with how much you want to bid. For example, if you are buyer 3 and you want to bid £40, say ‘buyer 3 bids 40’.
- In the sequence of market activity leading up to a trade, each bid must be higher than the current bid. For example, if one buyer bids £20, then all subsequent buyers need to bid more than £20.
- All bids need to be whole numbers. For example, while you can bid £30, you cannot bid £30.14.
- Just as active buyers may make bids (at any point in time), active sellers can make ‘asks’ (at any point in time). For example, if a seller ‘asks’ for £70, that means that she is willing to sell for £70. Each new ask must be lower than the current ask, so asks must ‘improve’ over time.
- At any stage, active buyers may accept an ask that has been made by a seller. To do this, raise your hand and wait until the auctioneer points at you. Once this has occurred, state your identity and that you want to accept the current ask. For example, if you are buyer 2 and want to accept an ask of £70, say ‘buyer 2 accepts 70’.
- If multiple buyers want to make a bid or accept an ask, priority will be given to the buyer who has raised their hand first.
- If you have made a trade, you become inactive and cannot make any further trades in that round. At this point, you should move to the inactive area and allow your place to be taken by the buyer at the front of the queue (or go to the back of the queue if you were previously a pending buyer). That new buyer will acquire your valuation (so should examine the card on the desk to see what that valuation is).
- While each bid needs to be higher than the previous bid, everything is reset following a trade. In other words, once a trade has been made, active buyers are free to submit any bid that they choose – even if that bid is lower than previously submitted bids.

Example

- Buyer 3 wants to bid £40 so raises her hand. Once she is pointed at by the auctioneer, she says ‘buyer 3 bids 40’.
- Buyer 1 wants to outbid her, so raises his hand. Once he is pointed at by the auctioneer, he says ‘buyer 1 bids 43’.
- Seller 4 wants to offer to sell for £95, so raises her hand. Once she is pointed at by the auctioneer, she says ‘seller 4 asks 95’.
- Buyer 2 wants to accept the ask, so raises his hand. Once he is pointed at by the auctioneer, he says ‘buyer 2 accepts 95’. A trade has occurred (at a price of £95).
- Since buyer 2 has just traded with seller 4, they both become inactive and should move to the inactive area. Two traders from the queue take their place and acquire their valuation and cost respectively. Participants then continue to bargain over prices. Since a trade has just occurred, subsequent bids are no longer constrained to be above £43 (the previous leading bid).

Rounds

- Trade will continue in this fashion until the auctioneer chooses to end the round.
- Within each round, you are only allowed to purchase (at most) one unit of the commodity. However, each round starts afresh: so even if you have made a trade in a particular round, you are free to make trades in subsequent rounds.
- After several rounds of trading have concluded, you will be given a new valuation for the fictitious commodity. So even if your current valuation is rather low, you might be luckier later on!
- At the end of the experiment, one round will be randomly selected to ‘count’ for calculating your earnings. Since any of the rounds could turn out to be the one that counts, you should do your best to maximise your net earnings within each round.

Summary

- There are two types of buyers: active buyers (IDs 1-5) and pending buyers (IDs 6 or higher). Once active buyers have made a trade, the pending buyer at the front of the queue takes their place and their valuation.
- Active buyers may make bids or accept asks at any point in time. Similarly, active sellers may make asks or accept bids at any point in time.
- If you want to make a bid or accept an ask, you need to first raise your hand.

- Bids must be whole numbers; and new bids must be greater than previous bids (until a trade is made).
- If you manage to make a trade, you earn your valuation minus the price you agreed to pay (assuming that this trade is selected to ‘count’ for your earnings).
- You can purchase up to one unit within every round.

C.2 Instructions for Sellers (Sessions 1 and 2)

Welcome to the experiment! Please read the following instructions as carefully as possible.

Preliminaries

- Please do not talk to your fellow participants at any stage. Talking may result in a loss of experimental earnings.
- You will have just received an ID card specifying your seller number. S1 means ‘Seller 1’, S2 means ‘Seller 2’, and so forth. Please wear your ID card visibly at all times.

Roles

- There are two types of sellers in this experiment: active sellers and pending sellers. If your seller number is between 1 and 5 inclusive, then you are an active seller and will be sitting in the main trading area. On the other hand, if your seller number is 6 or higher, then you are a pending seller and will be sitting in the queue.
- Active sellers are free to trade from the very start of a trading period. However, pending sellers may only trade after active sellers have made a trade and dropped out of the market (see elaboration below).

Costs

- If you are an active seller, you will have a card marked ‘Cost’ in front of you. Examine the number on the back of this card, being careful not to let any other participants see this number. This number is how much it would cost you to produce and sell a unit of the fictitious commodity to be traded. Memorise your cost, turn your ID card back face-down (so the cost is hidden), and do not reveal your cost to anybody else!
- If you are a pending seller, then you have not yet been assigned a cost. However, you will acquire a cost as soon as an active seller has dropped out and you have taken that seller’s place (and cost).

- If you manage to make a trade, you will receive the price paid by the buyer minus your cost (this assumes that the trade ‘counts’ — see discussion below). For example, if you sell for a price of £55 and your cost is £20, your net earnings are £35.
- Notice that your cost represents the least that you should be conceivably willing to sell for. For example, if your cost is £40, then you should never sell for less than £40: doing so would just lose you money!

Trading

- At any point, active sellers may offer to sell at a particular price – this is called making an ‘ask’. To make an ask, raise your hand. Once the auctioneer has pointed at you, state your identity along with how much you are asking for. For example, if you are seller 3 and you want to ask for £60, say ‘seller 3 asks 60’.
- In the sequence of market activity leading up to a trade, each ask must be lower than the current ask. For example, if one seller asks £80, then all subsequent sellers need to ask for less than £80.
- All asks need to be whole numbers. For example, while you can ask £70, you cannot ask £70.14.
- Just as active sellers may make asks (at any point in time), active buyers can make bids (at any point in time). For example, if a buyer bids £70, that means that she is willing to pay £70. Each new bid must be higher than the current bid, so bids must ‘improve’ over time.
- At any stage, active sellers may accept a bid that has been made by a buyer. To do this, raise your hand and wait until the auctioneer points at you. Once this has occurred, state your identity and that you want to accept the current bid. For example, if you are seller 2 and want to accept a bid of £70, say ‘seller 2 accepts 70’.
- If multiple sellers want to make an ask or accept a bid, priority will be given to the seller who has raised their hand first.
- If you have made a trade, you become inactive and cannot make any further trades in that round. At this point, you should move to the inactive area and allow your place to be taken by the seller at the front of the queue (or go to the back of the queue if you were previously a pending seller). That new seller will acquire your cost (so should examine the card on the desk to see what that cost is).
- While each ask needs to be lower than the previous ask, everything is reset following

a trade. In other words, once a trade has been made, you are free to submit any ask that you choose – even if that ask is higher than previously submitted asks.

Example

- Seller 3 wants to ask for £60 so raises her hand. Once she is pointed at by the auctioneer, she says ‘seller 3 asks 60’.
- Seller 1 wants to undercut her, so raises his hand. Once he is pointed at by the auctioneer, he says ‘seller 1 asks 57’.
- Buyer 4 wants to bid £5, so raises her hand. Once she is pointed at by the auctioneer, she says ‘buyer 4 bids 5’.
- Seller 2 wants to accept the bid, so raises his hand. Once he is pointed at by the auctioneer, he says ‘seller 2 accepts 5’. A trade has occurred (at a price of £5).
- Since seller 2 has just traded with buyer 4, they both become inactive and should move to the inactive area. Two traders from the queue take their place and acquire their valuation and cost respectively. Participants then continue to bargain over prices. Since a trade has just occurred, subsequent asks are no longer constrained to be below £57 (the previous lowest).

Rounds

- Trade will continue in this fashion until the auctioneer chooses to end the round.
- Within each round, you are only allowed to sell (at most) one unit of the commodity. However, each round starts afresh: so even if you have made a trade in a particular round, you are free to make trades in subsequent rounds.
- After several rounds of trading have concluded, you will be given a new cost for the fictitious commodity. So even if your current cost is rather high, you might be luckier later on!
- At the end of the experiment, one round will be randomly selected to ‘count’ for calculating your earnings. Since any of the rounds could turn out to be the one that counts, you should do your best to maximise your net earnings within each round.

Summary

- There are two types of sellers: active sellers (IDs 1-5) and pending sellers (IDs 6 or higher). Once active sellers have made a trade, the pending seller at the front of the queue takes their place and their cost.

- Active sellers may make asks or accept bids at any point in time. Similarly, active buyers may make bids or accept asks at any point in time.
- If you want to make an ask or to accept a bid, you need to first raise your hand.
- Asks must be whole numbers; and new asks must be lower than previous asks (until a trade is made).
- If you manage to make a trade, you earn the price paid by the buyer minus your cost (assuming that this trade is selected to ‘count’ for your earnings).
- You can sell up to one unit within every round.

C.3 Instructions for Buyers (Sessions 3 – 6)

Welcome to the experiment! Please read the following instructions as carefully as possible.

Preliminaries

- Please do not talk to your fellow participants at any stage. Talking may result in a loss of experimental earnings.
- You will have just received an ID card specifying your buyer number. B1 means ‘Buyer 1’, B2 means ‘Buyer 2’, and so forth. Please keep your ID card visible at all times.

Valuations

- Please examine the number of the back of your ID card, being careful not to let any other participants see this number. This number is your ‘valuation’ for the fictitious commodity to be traded. Memorise your valuation, turn your ID card back face-down (so the valuation is hidden), and do not reveal your valuation to anybody else!
- If you manage to make a trade, you will receive your valuation minus the price that you agreed to pay (this assumes that the trade ‘counts’ — see discussion below). For example, if your valuation is £55 and you agree to a price of £20, your net earnings are £35.
- Notice that your valuation represents the most that you should be conceivably willing to pay to make a trade. For example, if your valuation is £40, then you should never pay more than £40 to buy a unit: doing so would just lose you money!

Trading

- At any point, you may offer to buy at a particular price – this is called making a ‘bid’. To make a bid, raise your hand. Once the auctioneer has pointed at you, state your identity along with how much you want to bid. For example, if you are buyer 3 and you want to bid £40, say ‘buyer 3 bids 40’.
- In the sequence of market activity leading up to a trade, each bid must be higher than the current bid. For example, if one buyer bids £20, then all subsequent buyers need to bid more than £20.
- All bids need to be whole numbers. For example, while you can bid £30, you cannot bid £30.14.
- Just as buyers may make bids (at any point in time), sellers can make ‘asks’ (at any point in time). For example, if a seller ‘asks’ for £70, that means that she is willing to sell for £70. Each new ask must be lower than the current ask, so asks must ‘improve’ over time.
- At any stage, you may accept an ask that has been made by a seller. To do this, raise your hand and wait until the auctioneer points at you. Once this has occurred, state your identity and that you want to accept the current ask. For example, if you are buyer 2 and want to accept an ask of £70, say ‘buyer 2 accepts 70’.
- If multiple buyers want to make a bid or accept an ask, priority will be given to the buyer who has raised their hand first.
- Once you have made a trade, you become inactive and cannot make any further trades in that round. Likewise, each seller is only able to make at most one trade within a round – so it is as if they possess just one unit of the fictitious commodity.
- While each bid needs to be higher than the previous bid, everything is reset following a trade. In other words, once a trade has been made, you are free to submit any bid that you choose – even if that bid is lower than previously submitted bids.

Example

- Buyer 3 wants to bid £40 so raises her hand. Once she is pointed at by the auctioneer, she says ‘buyer 3 bids 40’.
- Buyer 1 wants to outbid her, so raises his hand. Once he is pointed at by the auctioneer, he says ‘buyer 1 bids 43’.
- Seller 4 wants to offer to sell for £95, so raises her hand. Once she is pointed at by the auctioneer, she says ‘seller 4 asks 95’.

- Buyer 2 wants to accept the ask, so raises his hand. Once he is pointed at by the auctioneer, he says ‘buyer 2 accepts 95’. A trade has occurred (at a price of £95).
- Since buyer 2 has just traded with seller 4, they both become inactive and remain silent for the rest of the round. Meanwhile, other participants continue to bargain over prices. Since a trade has just occurred, subsequent bids are no longer constrained to be above £43 (the previous leading bid).

Rounds

- Trade will continue in this fashion until there is a pause in market activity for roughly 20 seconds. At that point, the auctioneer will ask buyers if they would like to make any new bids, or would like to accept the current market ask. The auctioneer will then ask sellers if they would like to make any new asks, or would like to accept the current market bid. If all traders remain silent, then the auctioneer will close the market and conclude the round.
- Within each round, you are only allowed to purchase (at most) one unit of the commodity. However, each round starts afresh: so even if you have made a trade in a particular round, you are free to make trades in subsequent rounds.
- After several rounds of trading have concluded, you will be given a new valuation for the fictitious commodity. So even if your current valuation is rather low, you might be luckier later on!
- At the end of the experiment, one round will be randomly selected to ‘count’ for calculating your earnings. Since any of the rounds could turn out to be the one that counts, you should do your best to maximise your net earnings within each round.

Summary

- There is a group of buyers, and a group of sellers. Buyers may make bids or accept asks at any point in time. Similarly, sellers may make asks or accept bids at any point in time.
- If you want to make a bid or accept an ask, you need to first raise your hand.
- Bids must be whole numbers; and new bids must be greater than previous bids (until a trade is made).
- If you manage to make a trade, you earn your valuation minus the price you agreed to pay (assuming that this trade is selected to ‘count’ for your earnings).

- You can purchase up to one unit within every round.

C.4 Instructions for Sellers (Sessions 3 – 6)

Welcome to the experiment! Please read the following instructions as carefully as possible.

Preliminaries

- Please do not talk to your fellow participants at any stage. Talking may result in a loss of experimental earnings.
- You will have just received an ID card specifying your seller number. S1 means ‘Seller 1’, S2 means ‘Seller 2’, and so forth. Please keep your ID card visible at all times.

Costs

- Please examine the number on the back of your ID card, being careful not to let any other participants see this number. This number is how much it would cost you to produce and sell one unit of the fictitious commodity to be traded. Memorise your cost, turn your ID card back face-down (so the cost is hidden), and do not reveal your cost to anybody else!
- If you manage to make a trade, you will receive the price paid by the buyer minus your cost (this assumes that the trade ‘counts’ — see discussion below). For example, if you sell for a price of £55 and your cost is £20, your net earnings are £35.
- Notice that your cost represents the least that you should be conceivably willing to sell for. For example, if your cost is £40, then you should never sell for less than £40: doing so would just lose you money!

Trading

- At any point, you may offer to sell at a particular price – this is called making an ‘ask’. To make an ask, raise your hand. Once the auctioneer has pointed at you, state your identity along with how much you are asking for. For example, if you are seller 3 and you want to ask for £60, say ‘seller 3 asks 60’.
- In the sequence of market activity leading up to a trade, each ask must be lower than the current ask. For example, if one seller asks £80, then subsequent sellers need to ask for less than £80.
- All asks need to be whole numbers. For example, while you can ask for £70, you cannot ask for £70.14.

- Just as sellers may make asks (at any point in time), buyers can make bids (at any point in time). For example, if a buyer bids £70, that means that she is willing to pay £70. Each new bid must be higher than the current bid, so bids must ‘improve’ over time.
- At any stage, you may accept a bid that has been made by a buyer. To do this, raise your hand and wait until the auctioneer points at you. Once this has occurred, state your identity and that you want to accept the current bid. For example, if you are seller 2 and want to accept a bid of £70, say ‘seller 2 accepts 70’.
- If multiple sellers want to make an ask or accept a bid, priority will be given to the seller who has raised their hand first.
- Once you have made a trade, you become inactive and cannot make any further trades in that round — so it is as if you possess just one unit of the fictitious commodity. Likewise, each buyer is only able to make at most one trade within a round.
- While each ask needs to be lower than the previous ask, everything is reset following a trade. In other words, once a trade has been made, you are free to submit any ask that you choose – even if that ask is higher than previously submitted asks.

Example

- Seller 3 wants to ask for £60 so raises her hand. Once she is pointed at by the auctioneer, she says ‘seller 3 asks 60’.
- Seller 1 wants to undercut her, so raises his hand. Once he is pointed at by the auctioneer, he says ‘seller 1 asks 57’.
- Buyer 4 wants to bid £5, so raises her hand. Once she is pointed at by the auctioneer, she says ‘buyer 4 bids 5’.
- Seller 2 wants to accept the bid, so raises his hand. Once he is pointed at by the auctioneer, he says ‘seller 2 accepts 5’. A trade has occurred (at a price of £5).
- Since seller 2 has just traded with buyer 4, they both become inactive and remain silent for the rest of the round. Meanwhile, other participants continue to bargain over prices. Since a trade has just occurred, subsequent asks are no longer constrained to be below £57 (the previous lowest ask).
-

Rounds

- Trade will continue in this fashion until there is a pause in market activity for roughly 20 seconds. At that point, the auctioneer will ask buyers if they would like to make any new bids, or would like to accept the current market ask. The auctioneer will then ask sellers if they would like to make any new asks, or would like to accept the current market bid. If all traders remain silent, then the auctioneer will close the market and conclude the round.
- Within each round, you are only allowed to sell (at most) one unit of the commodity. However, each round starts afresh: so even if you have made a trade in a particular round, you are free to make trades in subsequent rounds.
- After several rounds of trading have concluded, you will be given a new cost for the fictitious commodity. So even if your current cost is rather high, you might be luckier later on!
- At the end of the experiment, one round will be randomly selected to ‘count’ for calculating your earnings. Since any of the rounds could turn out to be the one that counts, you should do your best to maximise your net earnings within each round.

Summary

- There is a group of sellers, and a group of buyers. Sellers may make asks or accept bids at any point in time. Similarly, buyers may make bids or accept asks at any point in time.
- If you want to make an ask or to accept a bid, you need to first raise your hand.
- Asks must be whole numbers; and new asks must be lower than previous asks (until a trade is made).
- If you manage to make a trade, you earn the price paid by the buyer minus your cost (assuming that this trade is selected to ‘count’ for your earnings).
- You can sell up to one unit within every round.