Introducing Nash Equilibria via an Online Casual Game That People Actually Play

David Aldous and Weijian Han

Abstract. This is an extended write-up of a lecture introducing the concept of *Nash equilibrium* in the context of an auction-type game that one can observe being played by "ordinary people" in real time. In a simplified model, we give an explicit formula for the Nash equilibrium. The actual game is more complicated and more interesting; players place a bid on one item (among several) during a time window; they can see the numbers but not the values of previous bids on each item. A complete theoretical analysis of the Nash equilibrium now seems a challenging research problem. We give an informal analysis and compare with data from the actual game.

1. INTRODUCTION. *Game theory* is an appealing mathematical topic, and there are perhaps 100 books giving introductory accounts in different styles. Styles using minimal mathematics range from "popular science" [**2**] to bookstore business section bestseller [**1**]. A wide-ranging account with a modicum of mathematics is provided in [**8**], while careful rigorous expositions at a lower division mathematical level can be found in [**9, 10**] and the recent e-book [**7**]. A representative of the numerous textbooks aimed at students of economics is [**3**], and an erudite overview from that viewpoint is given in [**4**].

However, game theory is invariably taught in the *Suppose Alice and Bob* ... paradigm of hypothetical games; textbooks either contain no real data or occasionally quote data obtained by others. The first author teaches an undergraduate course for statistics majors that seeks to explore the breadth of contexts where probability arises, via lectures which are (ideally) "anchored" by some interesting, new, real-world data. This article is based on a lecture attempting to do this for game theory, with data gathered from an online game that one can watch being played in real time. The game is `pogo.com`'s *Dice City Roller (DCR) [**6**]*, and from players' self-descriptions ("age 63, retired nurse: interests church, crafts, grandkids" is typical) we suppose the typical player is not a student of game theory. But do they play roughly in accordance with the predictions of game theory?

The actual game has other elements, in addition to the auction aspect we will study, but we will not describe these extra elements. Many students, and the authors, find it an intriguing game to play, so readers may wish to go online and play it themselves. For our mathematical analysis, the following abstracted model of the game is sufficient, *with italicized comments on actual play*.

Model of the DCR game.

- There are M items of somewhat different known values, say $b_1 \geq b_2 \geq \cdots \geq b_M > 0$ (*always* $M = 5$*, but the values vary between instances of the game*).
- There are N players (N *varies but* 5–12 *is typical*).
- A player can place a sealed bid for (only) one item during a window of time (*20 seconds*).

http://dx.doi.org/10.4169/amer.math.monthly.124.6.506
MSC: Primary 91A06

- During the time window, players see how many bids have already been placed on each item but do not see the bid amounts.

Of course, when time expires, each item is awarded to the highest bidder on that item. We assume players are seeking to maximize their expected gain. So a player has to decide three things: when to bid, which item to bid on, and how much to bid.

It turns out that, without the time element (that is, if players make sealed bids without any information about other players' bids), the game above is completely analyzable, as regards Nash equilibria (see Section 3). This is the mathematical content of the paper, and the results are broadly in line with intuition.

The time element makes the game more interesting because various strategies suggest themselves (see Section 5). Alas, theoretical analysis seems intractable, at least at an undergraduate level. One can analyze the simplest case (two players, two items, two discrete time periods), and the result is given as Proposition 2, but the answer is clearly special to that case and does not illuminate the general case. Continuing theoretical analysis of this "time element" setting is therefore a research project. Indeed, an incidental benefit of looking at real data is that it often suggests research-level theoretical problems.

Game theory and Nash equilibria. Here is a bullet point overview of game theory.

- Setting: players each separately choose from a menu of actions and get a payoff depending (in a known way) on all players' actions.
- Rock–paper–scissors illustrates why one should use randomized strategy and why we assume a player's goal is to maximize his or her expected payoff. There is a complete theory of such two-person zero-sum games.
- For other games, a fundamental concept is *Nash equilibrium* strategy: one such that, if all other players play that strategy, then you cannot do better by choosing some other strategy. This concept can be motivated mathematically by the idea that, if players adjust their strategies in a selfish way to maximize their own payoff and if the strategies converge to some limit strategy from which a player cannot improve by further adjustment, then by definition the limit strategy is a Nash equilibrium (which we now abbreviate to NE). This is the "learning-adjust" theory of why we might expect players to actually use NE strategies.

Overview. In this article, we will

- calculate the NE strategy in somewhat simplified versions of the real game (Sections 2 and 3);
- compare this with the data on what players actually do (Sections 1 and 4).

2. THE TWO-PLAYER TWO-ITEM GAME.

A simple game played earlier in class. In the first class of the course, students did several exercises to generate data that would be useful later, and this was one exercise in the Fall 2014 course.

> Imagine you and another player in the following setting. There are two items, a \$1 bill and a 50 cent coin. You can write a bid on one item, for instance, "I bid 37 cents for the \$1" or "I bid 12 cents for the 50 cent coin." If you and the other player bid on different items, then both get their item, and so you would make a gain of 63 cents in the first case or 38 cents in the second case. If you both bid on the same item, then only the higher bidder gets the item.
>
> Write down how much you would bid and on which item. After class, we will match your bid against a random other student's bid.

This was designed as the simplest possible variant of the DCR game. Let's calculate the NE and compare that with the class data.

Analysis of the two-player two-item game. To generalize very slightly the game above, there are two items, of values 1 and $b \in (0, 1]$, and there are two players. Each player places a sealed bid for one of the items, and when the bids are unsealed, the winners are determined. We assume each player is seeking to maximize his or her expected gain (rather than his or her gain relative to the other player's gain, which would make it a zero-sum game). We assume that bids on the first item are real numbers in $[0, 1]$, and bids on the second item are real numbers in $[0, b]$. A player's strategy is a pair of functions (F_1, F_b):

$$F_1(x) = \mathbb{P}(\text{ bid an amount} \leq x \text{ on the first item}), \quad 0 \leq x \leq 1 \tag{1}$$

$$F_b(y) = \mathbb{P}(\text{ bid an amount} \leq y \text{ on the second item}), \quad 0 \leq y \leq b \tag{2}$$

where

$$F_1(1) + F_b(b) = 1. \tag{3}$$

In the arguments below, we assume for simplicity that these distribution functions[1] have densities $f_1(x) = F_1'(x)$, $f_b(y) = F_b'(y)$, and we work with these densities where convenient. The reader familiar with measure theory will see that the general case requires only notational changes.

Intuition for playing this game seems simple: Bid more often on the more valuable item and, typically, bid low for the less valuable item or bid somewhat higher for the more valuable item. We will see this is borne out in the NE.

Proposition 1. *The unique Nash equilibrium strategy is*

$$F_1(x) = \tfrac{b}{1+b}(\tfrac{1}{1-x} - 1) \text{ on } 0 \leq x \leq \tfrac{1}{1+b} \tag{4}$$

$$F_b(y) = \tfrac{1}{1+b}(\tfrac{b}{b-y} - 1) \text{ on } 0 \leq y \leq \tfrac{b^2}{1+b}. \tag{5}$$

The functions F_1 and F_b for $b = 1/2$ are shown in Figure 1 below.

Proof. Any strategy has a *support*, the set of actions that might be used in that strategy; in our context, the support of strategy (F_1, F_b) is $\{x : f_1(x) > 0\} \cup \{y : f_b(y) > 0\}$. Our proof will exploit the following fundamental general *principle of indifference* for NE.

> If opponents play the NE strategy, then any nonrandom choice of action you make in the support of the NE strategy will give you the same expected gain (which equals the expected gain if you play the random NE strategy), and any other choice will give you smaller (or equal) expected gain.

This is true because the NE expected gain is an average gain over the different choices in its support; if these gains were not constant, then one would be larger than the NE gain, contradicting the definition of NE.

So now we suppose there is an NE strategy (F_1, F_b) in our game, and seek to calculate what it is. If you bid x on the first item, then you gain $1 - x$ if your opponent either bids on the second item (chance $F_b(b)$) or bids less than x on the first item (chance

[1] We slightly abuse terminology in calling these *distribution functions* because their individual masses are less than 1.

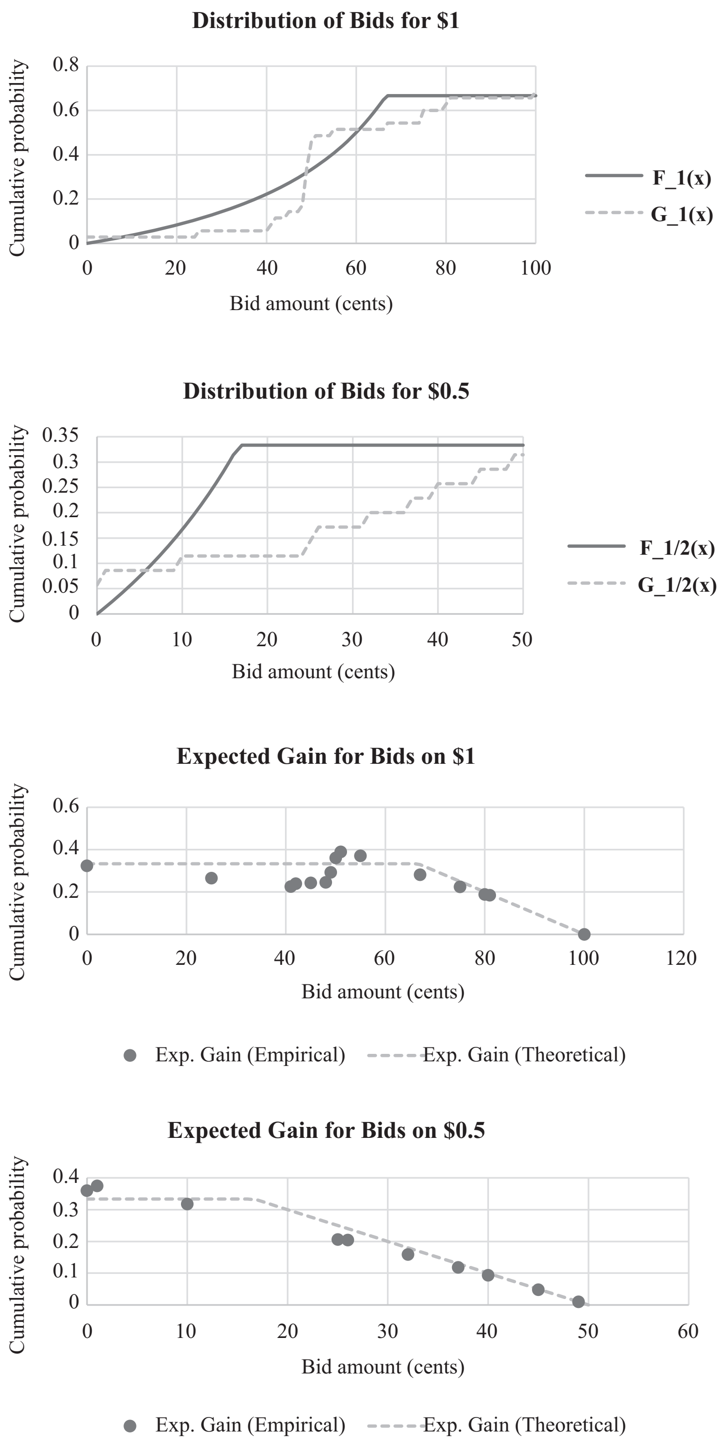

Figure 1. Class data compared with the NE.

$F_1(x)$). So your expected gain, if your opponent plays NE, equals $(1 - x)[F_1(x) + F_b(b)]$. Similarly, if you bid y on the second item, then your expected gain, if your opponent plays NE, equals $(b - y)[F_b(y) + F_1(1)]$. So the principle of indifference implies that, if you also play NE,

$$(1-x)[F_1(x)+F_b(b)] = c \text{ on support}(f_1) \quad (6)$$

$$\leq c \text{ off support}(f_1) \quad (7)$$

$$(b-y)[F_b(y)+F_1(1)] = c \text{ on support}(f_b) \quad (8)$$

$$\leq c \text{ off support}(f_b) \quad (9)$$

where c is the expected gain under NE.

What we need to show is that (6)–(9), together with "boundary conditions" (3) and $F_1(0) = F_b(0) = 0$, determine f_1, f_b, c uniquely. We give the argument in full detail below, though the essential calculations are quite simple, as we shall see with the general case later. Write x^* for the supremum of support (f_1). So $F_1(x^*) = F_1(1)$ and from (6) or directly from the principle of indifference,

$$1 - x^* = c.$$

We cannot have $c = 0$ (F_1 would put mass 1 at 1), so take $c > 0$. Using (6) and (7), we have

$$F_1(x) = \tfrac{1-x^*}{1-x} - F_b(b) \text{ on support}(f_1) \quad (10)$$

$$\leq \tfrac{1-x^*}{1-x} - F_b(b) \text{ off support}(f_1). \quad (11)$$

From (10), support(f_1) must be some interval $[x_*, x^*]$ because, if the support contained a gap (a, b), then $F_1(a) = F_1(b)$ is inconsistent with the strict monotonicity of the function $F_1(\cdot)$ in (10). Next, if $x_* > 0$, then (11) and the fact $F_1(x^*) = 0$ would force $F_1(x) < 0$ on $0 < x < x_*$, which is impossible. So we have shown support$(f_1) = [0, x^*]$. Then (10) for $x = 0$ shows that

$$F_b(b) = 1 - x^*, \quad (12)$$

and so for general x, (10) becomes

$$F_1(x) = (1-x^*)[\tfrac{1}{1-x} - 1], \;\; 0 \leq x \leq x^*,$$

and, in particular,

$$F_1(x^*) = x^*.$$

Now we can repeat for equation (8) the analysis done for (6); the same argument shows that support(f_b) must be an interval $[0, y^*]$. Then (8), together with facts $F_1(x^*) = x^*$ and $c = 1 - x^*$, gives

$$F_b(y) = \tfrac{1-x^*}{b-y} - x^*, \; 0 \leq y \leq y^*. \quad (13)$$

Because $F_b(0) = 0$, we can now identify the value of x^* as the solution of $\frac{1-x^*}{b} - x^* = 0$, that is,

$$x^* = \tfrac{1}{1+b}.$$

And because

$$F_b(y^*) = F_b(b) = 1 - F_1(1) = 1 - x^*,$$

applying (13) at y^* identifies y^* as the solution of $\frac{1-x^*}{b-y^*} = 1$, that is, $y^* = x^* + b - 1 = b^2/(1+b)$. This establishes the explicit formulas for F_1 and F_b stated in (4) and (5), and the corresponding densities are

$$f_1(x) = \tfrac{b}{1+b}(1-x)^{-2} \text{ on } 0 \le x \le \tfrac{1}{1+b} \tag{14}$$

$$f_b(y) = \tfrac{b}{1+b}(b-y)^{-2} \text{ on } 0 \le y \le \tfrac{b^2}{1+b}. \tag{15}$$

This essentially completes the proof. Careful readers will observe that we actually proved that, **if** (6)–(9), together with constraint (3), have a solution, then it must be (14), (15); such readers can check for themselves that this really is a solution. ■

Discussion of Proposition 1. The NE strategy is consistent with the qualitative intuitive strategy noted above the statement of Proposition 1. Here are some quantitative properties of the strategy that can be read off from the formulas above.

(i) Bid on the more valuable prize with probability $1/(1+b)$ and on the less valuable with probability $b/(1+b)$.

(ii) Conditional on bidding on the more valuable prize, your median bid is $1/(1+2b)$ and your maximum bid is $1/(1+b)$; conditional on bidding on the less valuable prize, your median bid is $b^2/(2+b)$ and your maximum bid is $b^2/(1+b)$.

(iii) Your expected gain is $b/(1+b)$.

(iv) The $b \downarrow 0$ limit is the NE strategy for the single item auction, which is to bid the full value.

Note that (i) says to bet on each item with probability proportional to the value of item; alas, this simple rule does not extend to $N > 2$ players (see (19) for the correct extension). In fact, the only aspect that generalizes nicely is that the gap between your maximum bid and the item's value is the same for both items; $1 - 1/(1+b) = b - b^2/(1+b) = b/(1+b)$. This follows from the principle of indifference because these are the expected gains from your maximum bid on each item. This "equal gap principle" works by the same argument for general numbers of players and items and provides a key simplification for the argument in Section 3.

Comparison with class data. We obtained data for this game with $b = 1/2$ by asking 35 students each to make one bid. The top two frames in Figure 1 compare the NE distribution functions F_1 and F_b defined at (4), (5) with the corresponding empirical distribution functions G_1 and G_b from the data. The bottom two frames in Figure 1 compare the NE expected gain from bidding different amounts with the corresponding empirical mean gain from the amounts bid by students. That is, a bid of 49 cents on the \$1 item had, when matched against a random other bid, mean gain of 29 cents, and this is represented by a point at (49, 29).

Here the data are not close to the NE. Students had some apparent intuition to bid around 50 cents on the \$1, and those who bid on the 50 cent items tended to overbid. But recall from Section 1 that the NE concept is motivated by the idea that, if players play repeatedly and adjust their strategies in a selfish way, then strategies should typically converge to some NE. So it is not reasonable to expect NE behavior the first time a game is played.

But in contrast, the actual DCR game is played repeatedly, and so it is more meaningful to ask whether players' strategies do in fact approximate the NE.

3. N PLAYERS AND M ITEMS.

We now consider the general case of $N \geq 2$ players and $M \geq 2$ items of values $b_1 \geq b_2 \geq \cdots \geq b_M > 0$. Using again the general principle of indifference for the NE and the special (to our model) "equal gap principle" described in Section 1, the actual calculation of the NE is surprisingly simple if we omit some details. The bottom line (with a caveat noted below reflecting omitted details!) is the formula

$$\text{expected gain to a player at NE } = c := \left(\frac{M-1}{\sum_i b_i^{-1/(N-1)}} \right)^{N-1}, \tag{16}$$

and the NE strategy is defined by the density functions at (20) below.

To derive the formula, define as in (1) and (2)

$$F_i(x) = \mathbb{P}(\text{bid on item } i, \text{bid amount} \leq x).$$

By the "equal gap principle," we take the NE strategy $(F_i(\cdot), 1 \leq i \leq M)$ to be such that each F_i is supported on $[0, b_i - c]$, where c is the expected gain to a player at NE. Writing out the expression for the expected gain when you bid x on item i, the principle of indifference says

$$(b_i - x)\,(1 - (F_i(x_i^*) - F_i(x)))^{N-1} = c, \quad 0 \leq x \leq x_i^* := b_i - c. \tag{17}$$

This is the generalization of (6), (8). Because a strategy is a probability distribution, we have $\sum_i F_i(x_i^*) = 1$, and so

$$\sum_i (1 - F_i(x_i^*)) = M - 1.$$

Now using (17) with $x = 0$, we have

$$1 - F_i(x_i^*) = (c/b_i)^{1/(N-1)}, \tag{18}$$

and so

$$\sum_i (c/b_i)^{1/(N-1)} = M - 1,$$

which rearranges to (16). So the probability that (at NE) you bid on item i is, by (18),

$$F_i(x_i^*) = 1 - (c/b_i)^{1/(N-1)} = 1 - \frac{b_i^{-1/(N-1)}}{\sum_j b_j^{-1/(N-1)}}(M-1). \tag{19}$$

Now (17) gives an explicit formula for $F_i(x)$, and differentiating gives the density

$$\begin{aligned} f_i(x) &= \frac{1}{N-1} c^{1/(N-1)} (b_i - x)^{-N/(N-1)}, \quad 0 \leq x \leq b_i - c, \\ &= \frac{M-1}{N-1} \frac{1}{\sum_j b_j^{-1/(N-1)}} (b_i - x)^{-N/(N-1)}, \quad 0 \leq x \leq b_i - c. \end{aligned} \tag{20}$$

The distribution function can be written as

$$F_i(x) = c^{1/(N-1)} \left(\left(\tfrac{1}{b_i - x}\right)^{1/(N-1)} - \left(\tfrac{1}{b_i}\right)^{1/(N-1)} \right), 0 \le x \le b_i - c, \tag{21}$$

where again c is the expected gain to a player at NE, at (16).

Reality check and caveat. As a reality check, consider the case of $N = 2$ players and $M = 3$ items of values $(b_1, b_2, b_3) = (1, 1, b)$ where $0 < b \le 1$. From the formulas above, we find

$$x_1^* = x_2^* = \tfrac{1}{1+2b}, \; x_3^* = \tfrac{b(2b-1)}{1+2b}; \tag{22}$$

$$F_1(x_1^*) = F_2(x_2^*) = \tfrac{1}{1+2b}, \; F_3(x_3^*) = \tfrac{2b-1}{1+2b}. \tag{23}$$

But for $b < 1/2$, this says $x_3^* < 0$ and $F_3(x_3^*) < 0$, which cannot be correct.

The mistake is that we implicitly assumed that the support of the NE strategy included a bid on every item. In the case above, with two items of values $(1, 1)$ the NE expected gain equals $1/2$, so availability of other items of value less than $1/2$ would not be helpful; players would play the two-item NE. This argument works in the general case, as follows. Recall we order item values as $b_1 \ge b_2 \ge \cdots \ge b_M > 0$. Inductively for $m = 2, 3, \ldots, M - 1$, calculate the NE and the expected gain assuming we have only the first m items available. If the expected gain is greater than b_{m+1}, then stop and use this NE strategy, which does not include a bid on any of $b_{m+1}, \ldots, b_M$. Otherwise, continue to $m + 1$.

However, one can show that it is only necessary to do this procedure if the original formula (16) for expected gain is manifestly wrong, in giving a value greater than the smallest value b_M, which would correspond to $x_M^* < 0$.

Discussion. If the items have equal value b, then we can find the expected gain more easily. The NE strategy will be symmetric over items, so the chance that no opponent bids on item 1 equals $((M-1)/M)^{N-1}$. So, assuming that the NE strategy includes bidding an amount close to 0, bidding such an amount earns you expected gain of $b((M-1)/M)^{N-1}$, and by the "constant expected gain" property, this is the expected gain at NE. Note that for large M and N the expected gain is around $b \exp(-M/N)$. The fact this depends on the ratio M/N — the average number of bids per item — is very intuitive, but the fact it decreases exponentially rather than polynomially fast is perhaps not so intuitive.

Minimum bid rule. An extra feature of the actual DCR game is that there is a minimum allowed bid on each item, say minimum bid $\theta_i < b_i$ on item i. Fortunately, the analysis above extends to this case with only minor changes: (16), (21) are replaced by

$$\text{expected gain to a player at NE } = c := \left(\frac{M-1}{\sum_i (b_i - \theta_i)^{-1/(N-1)}} \right)^{N-1} \tag{24}$$

$$F_i(x) = c^{\frac{1}{N-1}} \left(\left(\tfrac{1}{b_i - x}\right)^{\frac{1}{N-1}} - \left(\tfrac{1}{b_i - \theta_i}\right)^{\frac{1}{N-1}} \right), \quad \theta_i \le x \le b_i - c. \tag{25}$$

And the chance $F_i(b_i - c)$ of bidding on item i becomes

$$p_i := 1 - \left(\tfrac{c}{b_i - \theta_i}\right)^{\frac{1}{N-1}}. \tag{26}$$

4. COMPARING DATA FROM THE DCR GAME WITH NE THEORY. We obtained, via screenshots, data from 300 instances of the DCR game. Comparing these data with NE theory requires a certain fudge and involves a small complication. As mentioned before, the "time window" aspect makes the game more interesting because various strategies suggest themselves: Bid late on an item that few or no others have bid on, or bid early on a valuable item to discourage others from bidding on it. Figure 2 shows some data on when players place their bid. Instead of recording exact time of bids, we recorded bid times, via screenshots, as

early (20–14 seconds before deadline)
medium (14–5 seconds before deadline)
or *late* (5–0 seconds before deadline).

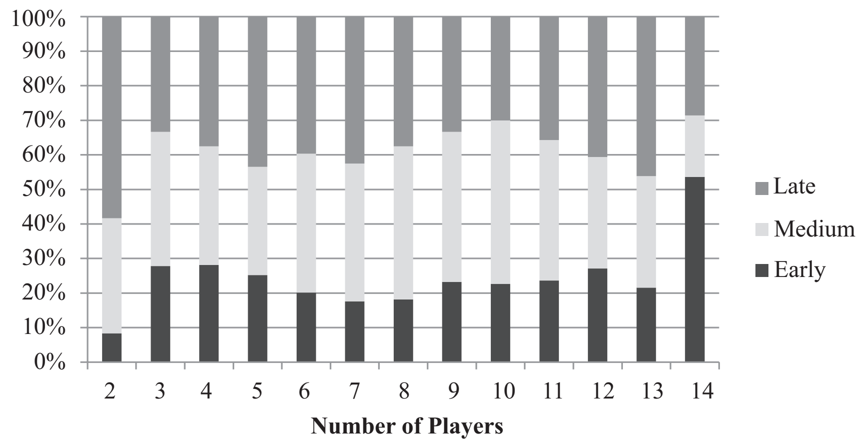

Figure 2. Distribution of early/medium/late bids, for varying numbers of players.

There is no clear pattern of bid times versus number of players, though bid times are widely spread over the window.

Using the analysis we have done, which ignores the "time window" aspect that players can see how many bids have been made on which items, the number of bids on item i at NE would have Binomial(N, p_i) distribution for p_i in (26). But the strategic considerations involve different players seeking to bid on different items, and therefore, we expect the distribution of the number of bids on a given item in the DCR game would be more concentrated around its mean than the corresponding binomial. And indeed this can be clearly seen in the data.

Another complication is that the observable data in the DCR game is the number of bids, and the value of the winning bid, on each item, but we cannot see the values of losing bids. So, for a given pair (N, i) of (number of players, item), the data we have available is the empirical distribution of values of winning bids over auctions where there was at least one bid. This is plotted as a distribution function G^* in Figure 3. We want to compare that to an "NE theory" distribution, and we obtain this by assuming that the amounts of bids follow the NE distribution (25), but (to allow for strategic effects) we use the true empirical distribution for the *number* of bids. Then we can numerically calculate an "NE theory" distribution function for value of winning bid, and this is plotted as a distribution function G in Figure 3.

Figure 3 shows the comparison between data and NE theory. The labels "150 match," etc. are our names for some of the items, and these data are for $N = 8$ players. The amounts of rewards and bids are in "points" internal to the game, not convertible to real-world money.

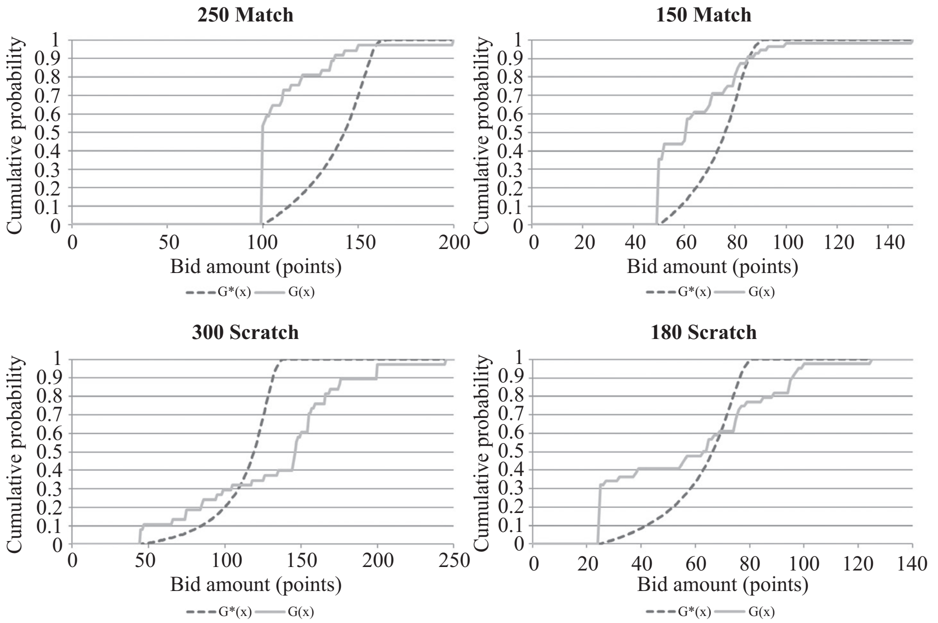

Figure 3. Comparison of winning bid distribution from DCR data and from NE theory. The smooth curve G is the theoretical distribution function,and the staircase function G^* is the empirical distribution function.

One's first reaction to the Figure 3 data is that the players' bids are not very close to what NE theory would predict. One could imagine many reasons for this discrepancy. From a typical player self-description, we suppose the typical player is not a student of game theory so might not consider the idea of conscious randomization. The fact that the winning bid is, in roughly a third of these cases, the minimum allowed bid is clearly a consequence of time-window strategy (making a last-second bid on an item no one else has bid on) not taken into account in our theory, so the data might be closer to the true NE than to our approximate NE.

5. INTRODUCING A TIME ELEMENT.

In seeking to model the "20-second window" aspect of the actual DCR game, a natural start is to discretize time into s stages. So the model is:

> At the start of each stage, you are told the numbers of bids on the different items in previous stages, and if you have not already placed a bid, then you can place a bid in that stage.

Note that in collecting data from the DCR game, we were anticipating comparing it to the NE for the three-stage model.

Developing NE theory in this setting turns out to be a challenging research project, so let us just state the result in the simplest possible case.

Proposition 2. *Consider the two-player two-item two-stage game, with item values* 1 *and* $b \leq 1$. *This game has (at least) the following two Nash equilibrium strategies.*

Strategy A *has mean gain per player*

$$c := \frac{b(b^2 + b + 1)}{(b + 1)(b^2 + 1)}. \tag{27}$$

The probabilities of (bid on item 1 in first stage, bid on item 2 in first stage, wait) are $(1-c, 1-\frac{c}{b}, c+\frac{c}{b}-1)$, *and the bid amounts have the distributions*

$$F_1(x) = c(\tfrac{1}{1-x}-1),\ 0 \le x \le x^* = 1-c; \quad F_1(x^*) = F_1(1) = 1-c \quad (28)$$

$$F_b(y) = c(\tfrac{1}{b-y}-\tfrac{1}{b}),\ 0 \le y \le y^* = b-c; \quad F_b(y^*) = F_b(b) = 1-\tfrac{c}{b}. \quad (29)$$

If you wait and the opponent bids in the first stage, then you bid 0 *on the other item in the second stage. If you both wait, then in the second stage, you bid according to the one-stage NE strategy (4), (5).*

Strategy B *has mean gain per player =* $b/(1+b)$. *Never bid in the first stage. If opponent bids in the first stage, then bid the whole value (*1 *or* b*) on the same item in the second stage. If opponent does not bid in the first stage, then in the second stage, bid according to the one-stage NE strategy.*

One can quickly check that the gain c in (27) is larger than the corresponding gain $b/(1+b)$ in the one-stage game, and in that sense, strategy A is better than strategy B. Strategy B is quite counterintuitive, but clearly **is** an NE because an opponent cannot gain from bidding in the first stage.

To outline the argument for Proposition 2, it is clear that any NE strategy has the following properties:

(a) If neither player bids in the first stage, then bid according to the one-stage NE strategy in the second stage.

(b) If opponent bids on item 2 in the first stage and you do not bid, then bid 0 on item 1 in the second stage.

(c) If opponent bids on item 1 in the first stage and you do not bid and if you choose to bid on item 2 in the second stage, then bid 0.

Thus, to specify a strategy reduces to specifying the functions F_1 and F_b in (1), (2) representing distributions of bid amounts in the first stage, now constrained by $F_1(1) + F_b(b) \le 1$, and also specifying the analogous function G_1 for the bid amount on item 1 in the second stage, given that opponent bids on item 1 and you do not bid in the first stage.

Now consider a special strategy satisfying (a), (b), and also (d).

(d) If opponent bids on item 1 in the first stage and you do not bid, then always bid 0 on item 2 in the second stage.

Specifying such a strategy requires only specifying F_1 and F_b, and the analysis is similar to the analysis in the one-stage case, Proposition 1, with an extra "gain" term arising from the second stage. Solving the analogous equations given by the principle of indifference for a strategy satisfying (d) to be an NE with respect to opponents constrained by (d) leads to the formulas stated as strategy A. To complete the proof that strategy A really is an NE, we need to check that an opponent cannot benefit by deviating from (d), that is, by bidding instead on item 1 in that situation, and this can be checked by calculation.

We strongly suspect that strategies A and B are the **only** NE. To prove this seems to require a detailed analysis using the three functions F_1, F_b, G_1, and we do not see a simple way to carry through the analysis. The existence of strategy B raises the possibility of other counterintuitive NEs, which means one must be careful about details of proof. Indeed, on a first analysis, we overlooked strategy B. For this reason, we were fussy about details in writing the proof of Proposition 1.

6. FINAL REMARKS.

1. Our setting differs from the usual setting of introductory game theory in that we use continuous, rather than discrete, actions. But our arguments show that calculating an NE in settings like ours often involves little more than basic calculus.

An incidental advantage of our "continuous" model is that it is novel even to students who have taken a game theory course.

2. In the actual DCR game, the prizes, if you win a bid, are a random number of points. In our mathematical analysis, we took the prizes to be the (nonrandom) expected value, so we are implicitly assuming

(a) players seek to maximize their expected number of points won;
(b) players know (learned from past experience) the expectations of the random number of points gained by winning bids on the various items.

In fact, after careful observation, one can actually calculate the expected values, which we used in calculating the theoretical distribution G in Figure 3.

3. Another interesting context for NE involves the "least unique positive integer" game, whose brief implementation as a real Swedish Lottery game attracted 50,000 players before it was realized that a consortium could "cheat" by buying sufficiently many numbers — see [**5**] for the NE analysis.

ACKNOWLEDGMENTS. We thank David Kreps for expert comments on the game theory background, Yufan Hu for a detailed analysis of prizes in the DCR game, Dan Lanoue for noticing strategy B in the time window setting, and two anonymous reviewers for helpful suggestions.

REFERENCES

1. A. K. Dixit, B. J. Nalebuff, *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*. Norton, New York, 1993.
2. L. Fisher, *Rock, Paper, Scissors: Game Theory in Everyday Life*. Basic Books, New York, 2008.
3. R. Gibbons, *Game Theory for Applied Economists*. Princeton Univ. Press, Princeton, NJ, 1992.
4. D. M. Kreps, *Game Theory and Economic Modelling*. Oxford Univ. Press, Oxford, 1990.
5. R. Östling, J. T. Wang, E. Y. Chou, C. F. Camerer, *Testing game theory in the field: Swedish LUPI lottery games*, *American Economic Journal: Microeconomics* **3** (2011) 1–33.
6. Pogo.com, "How to play Dice City Roller", 2016. http://www.pogo.com/game/help.do?game=ytz.
7. E. Prisner, *Game Theory through Examples*. Mathematical Association of America, Washington, DC, 2014.
8. E. C. Rosenthal, *The Complete Idiot's Guide to Game Theory*. Alpha Books, Indianapolis, IN, 2011.
9. S. Stahl, *A Gentle Introduction to Game Theory*. American Mathematical Society, Providence, RI, 1999.
10. P. D. Straffin, *Game Theory and Strategy*. Mathematical Association of America, Washington, DC, 1993.

DAVID ALDOUS received his Ph.D. in mathematics from Cambridge University, and has been on the faculty of the U.C. Berkeley Statistics department since 1979. A central research theme has been the study of large finite random structures, obtaining asymptotic behavior as the size tends to infinity via consideration of some suitable infinite random structure. He is currently engaged in articulating critically what mathematical probability says about the real world.
Department of Statistics, University of California, Berkeley, CA 94720-3860
aldous@stat.berkeley.edu

WEIJIAN HAN received his B.S. degree in statistics and economics from U.C. Berkeley in 2016 and has started pursuing a career as an actuary. His academic interests include game theory, econometrics, and machine learning, and he is particularly interested in using simulations to study statistical theories.
Mercer Consulting, 4 Embarcadero Center, #400, San Francisco, CA 94111
james.han1008@gmail.com