Introducing Nash equilibria via an online casual game which people actually play

David Aldous and Weijian Han

October 31, 2014

Abstract

This is an extended write-up of a lecture introducing the concept of *Nash equilibrium* in the context of an auction-type game which 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 (amongst 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

The first author teaches an undergraduate course "Probability in the Real World" with a non-traditional format. There are twenty 80-minute lectures on very different topics; each lecture is (ideally) "anchored" by some interesting new real-world data; and I talk about some theory relevant to understanding this specific data. From the traditional viewpoint this format has the obvious defect that one can hardly go in depth into any topic in a single lecture. But my goal is not to teach mathematical technique or statistical methodology, but to inspire students to think about the relationship between textbook material and real data. Instead of homework and exams, students do a course project of their choice which is (ideally) in the same spirit of finding and studying some new data relevant to some topic in probability or statistical theory.

In a typical lecture I give a brief overview of the topic, do a little mathematics aimed in the direction of the particular data-set, and then jump (if necessary) to stating that more advanced theory gives some specific prediction for what we would expect to see in data like this; and finally compare the prediction to the particular new real-world data. Of course some of the jumped-over theory can be outlined in the lecture write-up. See [1] for this style of lecture on the topic of prediction markets and martingales, and see [2] for a typical student project (finding data to check Benford's Law). This article describes a lecture in this style on the topic of game theory. Sections 2 - 6 are an expanded version of what I actually do in class.

Game theory is an appealing mathematical topic, and there are perhaps a hundred books giving introductory accounts in different styles. Styles using minimal mathematics range from "popular science" [6] to airport bookstore Business section bestseller [5]. A wide-ranging account with a modicum of mathematics is provided in [12], while careful rigorous expositions at a lower division mathematical level can be found in [13, 14] and the recent e-book [11]. A representative of the numerous textbooks aimed at students of Economics is [7], and an erudite overview from that viewpoint is given in [9]. But such books either contain no real data, or occasionally quote data obtained by others. Where can we obtain some interesting new real data, by ourselves?

One common "experimental" answer is to have volunteers (typically one's students) play games, such as Prisoners' Dilemma or the Ultimatum Game, that explicitly fit the mathematical game theory framework. But people's behavior in such staged tournaments is not necessarily representative of behavior in the other aspects of life to which it is often claimed that game theory can be applied – aspects as varied as economic behavior, political or military conflict, or evolutionary biology. So ideally I would like data from such "real-world" settings. Now observing any one-time game-like setting without objectively recognizable quantitative payoffs, one can surely devise payoffs so that the observed outcome is consistent with game-theoretic predictions; and game-theoretic explanations in such settings are hardly more than the "just-so stories" famously satirized by Stephen Jay Gould. Economics settings, where payoffs are simply money, provide a more promising direction to seek repeatable data. Analyzing auction data for commodity items (e.g. iPhones) on eBay is a popular project for my students, with somewhat of a game-theoretic flavor. But it is hard to think of precise quantitative game-theoretic predictions with which to compare such observed data.

To get repeatable data in a context which is (partially) amenable to theoretical analysis, I chose a game (in the everyday sense of "game") which has an incidental game-theoretical aspect; the players likely have no technical familiarity with game theory, but simply play in the intuitive way that ordinary people play recreational table games. The game is `pogo.com`'s *Dice City Roller (DCR)*. In class, I start by spending a couple of minutes demonstrating the game by actually participating in it, in real time. In this article the written description of the DCR game is deferred to section 6; readers may wish to read it now, or go online and play it themselves, before reading further. For our mathematical analysis, the following abstracted model of the game is sufficient, *with italicized comments on actual play.*

1.1 Model of the DCR game

- There are M items of somewhat different known values, say $b_1 \geq b_2 \geq \dots b_M$ (*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*).
- 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 4. 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: bid late on an item that few or no others have bid on, seeking to obtain it cheaply, or bid early on a valuable item to discourage others from bidding on it. 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 described in section 7 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 – for instance the "prediction market" lecture mentioned above motivated the research paper [3].

We obtained data from 300 instances of the DCR game, and various statistical aspects of the data are shown in section 5. As mentioned above we do not have a precise formula for the NE strategy in the real game, but nevertheless we can formulate plausible approximations to the NE strategy. And the bottom line, discussed in section 5, is rather ambiguous. On one hand the "ordinary people" playing this game are not bidding in a way that is close to the NE strategy, but their deviations are not "foolish" in any specific way.

2 Starting the game theory lecture

In the lecture I give a quick bullet point overview of game theory.

1. Setting: players each separately choose from a menu of actions, and get a payoff depending (in a known way) on all players' actions.

2. Rock-paper-scissors illustrates why one should use randomized strategy, and why we assume a player's goal is to maximize their expected payoff. There is a complete theory of such two-person zero-sum games.

3. 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.

4. More advanced theory is often devoted to settings where Nash equilibria are not optimal in some sense, as with Prisoners' Dilemma, and to understanding why human behavior is not always selfish. For a glimpse of contemporary research see [4].

This lecture will focus on point 3. The DCR game is one which, to a game theorist, fits exactly into the setting of point 3. The "learning-adjust" theory predicts that players who play repeatedly and play selfishly – being unable or unwilling to collaborate with other players – will tend to adjust their strategies to approximate the Nash equilibrium (which we now abbreviate to NE) strategy. Further discussion of NE can be found in many places, for instance the textbook treatment in [8] or the high-level discussion in the "why study NE" section of [9], whose thesis is that **if** there is some natural way to play a game **then** that way must be a NE, but not conversely. Instead of general discussion, what I will do in this lecture is

- calculate the NE strategy in somewhat simplified versions of the real game;
- compare this with the data on what players actually do.

I do not seek to introduce and explain much standard game-theory terminology – for instance, the concept of NE refers, strictly speaking, to a *strategy profile*, that is a strategy for each player, but in our "symmetric over players" context we look only for NE strategy profiles in which each player uses the same strategy.

3 The 2-player 2-item game

3.1 A simple game played earlier in class

In the first class of the course, students do several exercises to generate data that will 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 – e.g. "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 – so 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, only the higher bidder gets the item.

Write down how much you would bid, and on which item. After class I will match your bid against a random other student's bid.

This was designed as the simplest possible variant of the real game – 2 players, 2 items, no time window. I remind students that we already have this small data-set – the 35 bids by students – so let's calculate the NE and compare that with the data.

3.2 Analysis of the 2-player 2-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 their expected gain (rather than their 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 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. A reader familiar with game theory might wish to think what can be said about the NE strategy without doing any calculations.

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

Proposition 1. *The unique Nash equilibrium strategy is* (F_1, F_b) *given in (14,15) below.*

Proof. We will give a somewhat pedantic "proof from first principles" using mathematical symbols. Some key ideas will be restated in words in section 3.3, and understanding these ideas will enable us to carry through the general case analysis (section 4) surprisingly easily, by omitting fussy details.

Your opponent's strategy is a function[2] (f_1, f_b) and your strategy is a function (g_1, g_b). Your expected gain equals

$$\int_0^1 (1-x)g_1(x)[F_1(x) + F_b(b)]\ dx + \int_0^b (b-y)g_b(y)[F_b(y) + F_1(1)]\ dy. \quad (4)$$

To explain the first term: 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 $F_1(x)$). The second term arises similarly from the case of bidding on the second item.

We now point out an obvious fact. Consider a function $h(x) \geq 0$ with $h^* := \sup_x h(x) < \infty$. Consider the functional $L(g) := \int h(x)g(x)dx$ as being defined on the space $\mathcal{F}$ of probability density functions g, and recall support(g) is the closure of $\{x : g(x) > 0\}$. Then

the functional $L(\cdot)$ *attains its maximum at* g_0 *iff*

$$h(x) = h^* \text{ for all } x \in \text{support}(g_0). \quad (5)$$

So given your opponent's strategy (f_1, f_b), your expected gain is maximized by choosing (g_1, g_b) satisfying, for some constant c (the h^* in (5))

$$\begin{aligned} (1-x)[F_1(x) + F_b(b)] &= c \text{ on support}(g_1) && (6)\\ &\leq c \text{ off support}(g_1) && (7)\\ (b-y)[F_b(y) + F_1(1)] &= c \text{ on support}(g_b) && (8)\\ &\leq c \text{ off support}(g_b). && (9)\end{aligned}$$

Now the definition of (f_1, f_b) being a NE is precisely the assertion that (6 - 9) hold for $(g_1, g_b) = (f_1, f_b)$. So assume that, for the remainder of the proof. We will show that (6 - 9), together with "boundary conditions" (3) and $F_1(0) = F_b(0) = 0$, determine f_1, f_b, c uniquely. Write x^* for the supremum of support(f_1). So $F_1(x^*) = F_1(1)$ and from (6)

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

[2] We envisage (f_1, f_b) as a single function defined on a union of two disjoint intervals.

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

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

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

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

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

and so for general x, (10) becomes

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

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 $y^* = x^* + b - 1 = b^2/(1 + b)$.

Now we can write everything explicitly:

$$F_1(x) = \tfrac{b}{1+b}(\tfrac{1}{1-x} - 1) \text{ on } 0 \leq x \leq \tfrac{1}{1+b} \quad (14)$$

$$F_b(y) = \tfrac{1}{1+b}(\tfrac{b}{b-y} - 1) \text{ on } 0 \leq y \leq \tfrac{b^2}{1+b}. \quad (15)$$

The corresponding densities are

$$f_1(x) = \tfrac{b}{1+b}(1 - x)^{-2} \text{ on } 0 \leq x \leq \tfrac{1}{1+b} \quad (16)$$

$$f_b(y) = \tfrac{b}{1+b}(b - y)^{-2} \text{ on } 0 \leq y \leq \tfrac{b^2}{1+b}. \quad (17)$$

This essentially completes the proof. Careful readers will observe that we actually proved that, if (6 - 9), together with (3), have a solution, it must be (16,17); such readers can check for themselves that this really is a solution. □

3.3 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 $b/(1+b)$.

Note that (i) says to bet proportional to the value of item; alas this simple rule does not extend to $N > 2$ players (see (21) 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 aspect can be seen without calculation: if your opponent's strategy had maximum bids x^*, y^* with $1 - x^* < b - y^*$, say, then taking your strategy as

> the modification of the opponent's strategy in which (for small $\varepsilon > 0$) bids on the first item in $[x^* - \varepsilon, x^*]$ are replaced by bids on the second item in $[y^*, y^* + \varepsilon]$

will increase your expected gain. The same "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 4, as explained below.

An important general principle about NE (not restricted to this particular game) was hidden in the argument around (5).

> If opponents play the NE strategy then any non-random 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 "constant expected gain" property 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.

This property allows us to extend the "equal gap principle" of this particular game to say, for general numbers of players and items,

> in the NE strategy, the gap between your maximum bid on an item and the item's value is the same for all items, and equals the expected gain at NE.

3.4 Comparison with class data

As mentioned in section 3.1 we obtained data for this game with $b = 1/2$ by asking students to make one bid. The top two frames in Figure 1 compare the NE distribution functions F_1 and F_b at (14,15) 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 is 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 2 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.

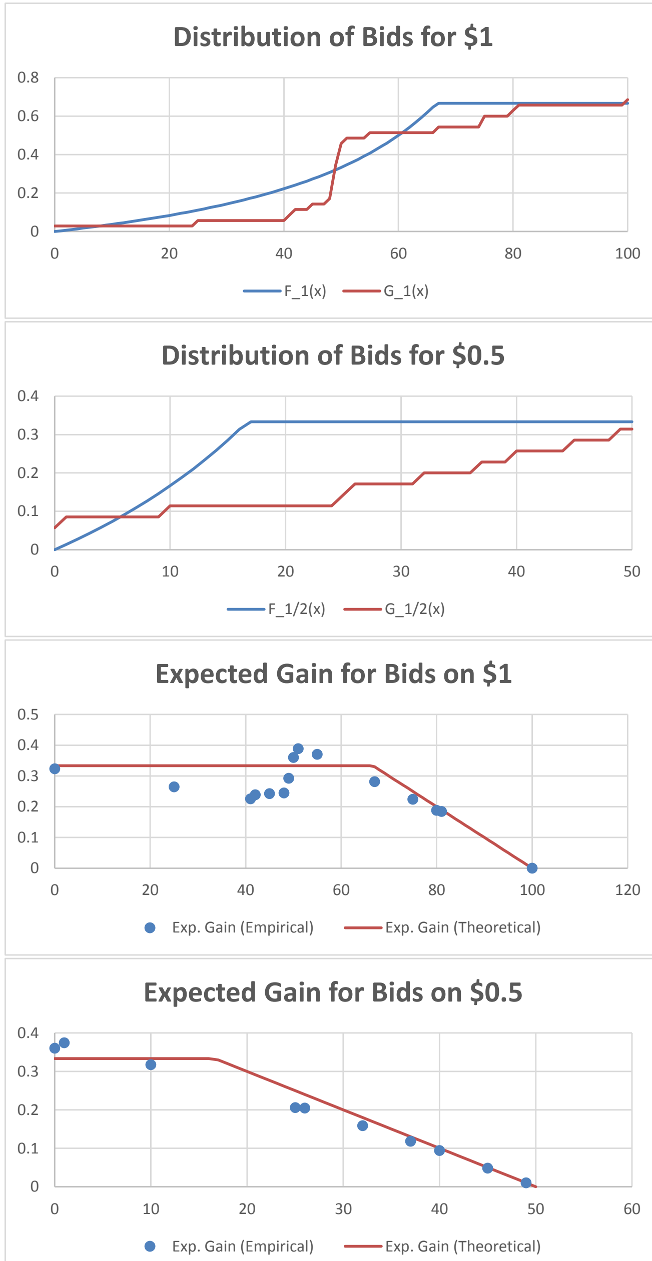

Figure 1. Class data compared with the NE.

4 N players and M items

We now consider the general case of $N \geq 2$ players and $M \geq 2$ items of values $b_1 \geq b_2 \geq \ldots \geq b_M > 0$. Armed with the general "constant expected gain" property of NE and the special (to our model) "equal gap principle" described in section 3.3, 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

$$\text{expected gain to a player at NE } = c := \left(\frac{M-1}{\sum_i b_i^{-1/(N-1)}} \right)^{N-1} \tag{18}$$

and the NE strategy is defined by the density functions at (22) below.

To derive the formula, define as at (1,2)

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

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_i on the i'th item, the "constant expected gain" property says

$$(b_i - x)\ (1 - (F_i(x_i^*) - F_i(x)))^{N-1} = c, \quad 0 \leq x \leq x_i^* := b_i - c. \tag{19}$$

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 (19) with $x = 0$ we have

$$1 - F_i(x_i^*) = (c/b_i)^{1/(N-1)} \tag{20}$$

and so

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

which rearranges to (18). So the probability that (at NE) you bid on item i is, by (20),

$$F_i(x_i^*) = 1 - (c/b_i)^{1/(N-1)} = 1 - \frac{b_i^{-1/(N-1)}}{\sum_j b_j^{-1/(N-1)}}(M-1). \tag{21}$$

Now (19) gives an explicit formula for $F_i(x)$, and differentiating gives the density

$$\begin{aligned} f_i(x) &= \tfrac{1}{N-1} c^{1/(N-1)} (b_i - x)^{-N/(N-1)}, \quad 0 \le x \le b_i - c \\ &= \frac{M-1}{N-1} \frac{1}{\sum_j b_j^{-1/(N-1)}} (b_i - x)^{-N/(N-1)}, \quad 0 \le x \le b_i - c. \quad (22) \end{aligned}$$

The distribution function can be written as

$$F_i(x) = c^{1/(N-1)} \left(\left(\tfrac{1}{b_i - x}\right)^{1/(N-1)} - \left(\tfrac{1}{b_i}\right)^{1/(N-1)} \right), 0 \le x \le b_i - c \qquad (23)$$

where again c is the expected gain to a player at NE, at (18).

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

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

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

The mistake is that we implicitly assumed that the NE strategy included a bid on each item (*include* means "assigns non-zero probability to"). We can fix the mistake as follows. Recall we order item values as $b_1 \ge b_2 \ge \ldots \ge b_M > 0$. 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 stategy which does not include a bid on any of $b_{m+1}, \ldots, b_M$. Otherwise continue to $m+1$. However, we only need to do this procedure if the original formula (18) for expected gain is manifestly wrong, in giving a value greater than the smallest value b_M.

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.

4.1 Minimum bid rule

An extra feature of the DCR game is that there is a minimum allowed bid on each item, say minimum bid $\theta_i < b_i$ on the i'th item. Fortunately the analysis above extends to this case with only minor changes: (18, 23) are replaced by

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

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

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

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

5 Comparing data from the DCR game with NE theory

Comparing data from the DCR game 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).

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

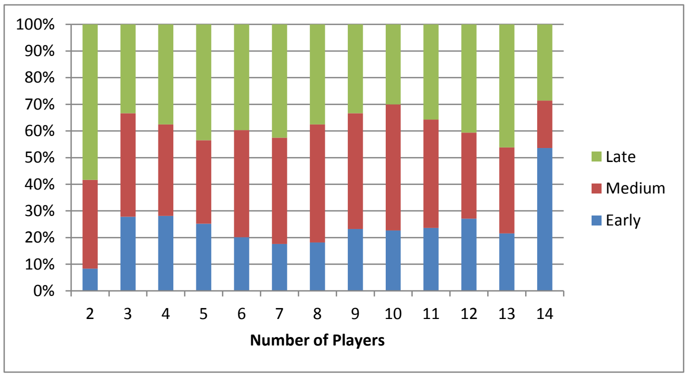

Figure 2. Distribution of early/medium/late bids, for varying numbers of players.

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 at (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 a "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 a "NE theory" distribution function for value of winning bid, and this is plotted as a distribution function G in Figure 3.

Deferring some further details and approximations to section 6, Figure 3 shows the comparison between data and NE theory. The labels "150 match" etc are our names for some of the items (explained in section 6), and this data is for $N = 8$ players.

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. A typical player self-description is "age 63, retired nurse: interests church, crafts, grandkids"; on this basis 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. A third possibility is described in section 6.1.

Figure 3. Comparison of winning bid distribution from data and from NE theory.

6 Dice City Roller

The game motivating this article is called *Dice City Roller* (DCR), and found on pogo.com, a free online casual gaming website offering over 150 different games. At a typical time there may be about 10 different active "rooms" each containing typically having 5 - 15 competing players – other rooms with 1 or 2 players do not concern us. The underlying game is illustrated by the screenshot in Figure 4 (details below not relevant to our mathematics, until further notice). An instance of the game consists of 12 repetitions of the following "turn". The player is shown five rolled dice, allocates them onto "cards" to fill out specified combinations (over several turns); when a card is completed the player earns points and a new card is offered.

Figure 4. Screenshot of basic play of the underlying game in DCR.

For instance, in Figure 4 the rolled dice show $1, 6, 2, 4, 1$. The player could place a 1 on the Full House, place a 2 and 6 on the Straight and place a 1 on the 3 Of A Kind, to complete 3 cards, placing the remaining 4 on one of the other cards. The player has 15 seconds to decide upon and execute these placements.

This underlying game is more subtle than it may appear, because there is a bonus for completing several cards on the same turn, so a simple greedy scheme for filling cards is not optimal. However this activity is not game-theoretic, because there is no interaction between players – one simply seeks to maximize one's score, that is one's total number of points at the end of the game.

What is relevant to us is the "auction" version which adds the following step, 3 times during the game. Players are allowed to use some of their points to bid for one of 5 prizes, a prize being the chance to earn extra points. The bidding proceeds as described in section 1.1:

> During a 20 second time window, players see how many bids have already been placed on each item, but do not see the bid amounts.

The screenshot in Figure 5 shows a situation 5 seconds before the window closes. Three players have placed bids, on different cards – these numbers of bids are shown in the disc at the cards' bottom left corner. Three other players had not yet placed bids. In the 5 seconds remaining after the screenshot, it happened that two players bid on the top right card (with 0 earlier bids) and one player bid on the bottom left card (with 1 earlier bid).

Figure 5. Screenshot of auction in progress.

In each auction there are 5 cards, from a set of 12 different cards. The prizes, if you win a bid, are a random number of points. In our mathematical analysis we took the prizes to be the (non-random) expected value for each card, so we are implicitly assuming

(a) players seek to maximize their expected number of points won

(b) players know the expectations, for instance implicitly learned by experience.

As noted in section 6.1 below, one can actually calculate the expected values. Issue (a) is more subtle. One way in which DCR resembles real life rather than a staged tournament is that the players' objectives are rather ill-defined. A player accumulates "pogo points" over the many different games offered by `pogo.com`. This is distinct from your score (total number of points) in one DCR game. The number of pogo points you earn from one DCR game depends partly linearly on your score, and partly on whether your score exceeds a certain threshold; actually winning a game (scoring more than the other players) earns you kudos but not pogo points. So a sophisticated player would switch between risk-averse and risk-seeking actions depending on their progress toward the threshold or toward winning; the relative importance of these two goals depending on some mental "exchange rate" between pogo points and kudos.

6.1 More details

1. As seen in Figure 5, each card shows the minimum bid allowed, the maximum possible prize and the "type" which is mostly *match* or *scratch*: our names like "150 match" refer to type and maximum prize. In "150 match" there are 6 covered numbers, and the winning bidder uncovers each until finding two equal numbers, and that number becomes the prize. One can learn that the 6 covered numbers are 50, 100 and 150, with two copies of each. So the prize is equally likely to be 50, 100, 150, with expectation 100. In a "scratch" card there are also 6 covered numbers; except that one is a bomb; the player uncovers numbers until reaching the bomb, and the prize is the sum of the values uncovered. For such a "scratch" card the maximum prize is the sum of the 5 numbers and the expected value is **exactly** half of this maximum. Learning all these numerical values requires careful observation, and we suspect typical players do not explicitly know these expected values. In particular, for "match" cards the expected value is always **more than** half of the maximum prize shown on the card. A player unaware of this distinction is liable to underbid on the "match" cards or overbid on the "scratch" card, which appears to be happening in the Figure 3 data.

2. The NE distribution for bid amounts on card i depends not only on i and N = number of players, but also on the other cards present in the auction. To produce Figure 3 we did the NE calculation separately for each auction in our data with a given combination (N, i) and averaged the distributions (which in fact vary little over different such auctions). Because of the large number of combinations of (N, i) we see small-sample fluctuations in our data for any particular combination. In Figure 3 we in fact averaged over the cases $N = 7, 8, 9$ to smooth the data.

3. Two of the 12 cards are "extra die next round"; we assigned an equivalent point value to those prizes for the purpose of calculating the NE bid distribution for other cards (amongst plausible values, the exact choice of value has negligible effect).

4. The site shows the total number of pogo points that players have accumulated, from which we can infer that many players have spent many thousands of hours playing different games on `pogo.com`. From this and the players' self-descriptions, we tell students to envisage a typical player as "your grandmother, who plays a mean game of gin rummy".

7 Introducing a time element

In seeking to model the "20 second window" aspect of the 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 3-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. *The 2-player 2-item 2-stage game, with item values* 1 *and* $b \leq 1$, *has the following Nash equilibrium strategy, which 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*

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

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

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 1-stage NE strategy (14,15).*

One can check that the gain c in (27) is larger than the corresponding gain $b/(1+b)$ in the 1-stage game. Perhaps unexpectedly this NE is not unique, as observed by Dan Lanoue, because there is the following rather anti-social NE strategy which essentially reduces the 2-stage game to the 1-stage game.

> 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 1-stage NE strategy.

The point is that this strategy punishes any opponent who bids in the first round, so forces the reduction to the 1-stage game.

8 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 NE in settings like ours often involves little more than basic calculus. Berkeley, like most large universities, has an undergraduate course in game theory, which is an optional course parallel to my (also optional) course described in the introduction. An incidental advantage of our "continuous" model is that it is novel even to students who have taken the game theory course.

2. As implied in section 2, an expert might remark that merely calculating a NE is not getting to grips with the essence of modern game theory. Agreed; but my general theme is comparing theory with data, and this is the best I can do in an 80 minute class.

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 [10] for the NE analysis.

Acknowledgements. We thank David Kreps for expert comments on the game theory background, Yufan Hu for a detailed analysis of prizes in the DCR game, and Dan Lanoue for comments on the time window setting.

References

[1] David Aldous. Using prediction market data to illustrate undergraduate probability. *Amer. Math. Monthly*, 120(7):583–593, 2013.

[2] David Aldous and Tung Phan. When can one test an explanation? Compare and contrast Benford's law and the fuzzy CLT. *Amer. Statist.*, 64(3):221–227, 2010.

[3] David Aldous and Mykhaylo Shkolnikov. Fluctuations of martingales and winning probabilities of game contestants. *Electron. J. Probab.*, 18:no. 47, 1–17, 2013.

[4] Valerio Capraro. A model of human cooperation in social dilemmas. *PLoS ONE*, 8(8):e72427, 08 2013.

[5] Avinash K. Dixit and Barry J. Nalebuff. *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life.* Norton, 1993.

[6] Len Fisher. *Rock, Paper, Scissors: Game Theory in Everyday Life.* Basic Books, 2008.

[7] Robert Gibbons. *Game Theory for Applied Economists.* Princeton University Press, 1992.

[8] David M. Kreps. *A Course in Microeconomic Theory.* Princeton University Press, 1990.

[9] David M. Kreps. *Game theory and economic modelling.* Oxford University Press, 1990.

[10] Robert Östling, Joseph Tao-yi Wang, Eileen Y. Chou, and Colin F. Camerer. Testing game theory in the field: Swedish LUPI lottery games. *American Economic Journal: Microeconomics*, 3(3):1–33, 2011.

[11] Erich Prisner. *Game Theory Through Examples.* Mathematical Association of America, Washington, DC, 2014. e-book; Classroom Resource Materials.

[12] Edward C. Rosenthal. *The Complete Idiot's Guide to Game Theory.* ALPHA, 2011.

[13] Saul Stahl. *A gentle introduction to game theory*, volume 13 of *Mathematical World.* American Mathematical Society, Providence, RI, 1999.

[14] Philip D. Straffin. *Game theory and strategy*, volume 36 of *New Mathematical Library*. Mathematical Association of America, Washington, DC, 1993.