How Facebook Can Deepen our Understanding of Behavior in Strategic Settings: Evidence from a Million Rock-Paper-Scissors Games*

Dimitris Batzilis[†] Sonia Jaffe[‡] Steven Levitt[§]

John A. List[¶] Jeffrey Picel[‖]

May 14, 2014

Abstract

Perhaps the most fundamental contribution to economics in the past century relates to understanding how agents behave in strategic settings. While our theoretical knowledge is deep, the empirical evidence remains behind. This study makes use of data from a Facebook application that had hundreds of thousands of people play a simultaneous move, zero-sum game – rock-paper-scissors – with varying information to provide empirical insights into whether play is consistent with extant theories. We report three major insights. First, we observe that many people employ strategies consistent with Nash equilibrium: that is, most people employ strategies consistent with Nash, at least some of the time. Second, players predictably respond to incentives in the game. For example, out of equilibrium, players strategically use information on previous play of their opponents, and they are more strategic when the payoffs for such actions increase. Third, experience matters: players with more experience use information on their opponents more

*We would like to thank Vince Crawford for his insights, Colin Camerer, William Diamond, Teck Ho, Scott Kominers, Jonathan Libgober, Guillaume Pouliot, and especially Larry Samuelson for helpful comments, the Roshambull development team for sharing the data, and Dhiren Patki and Eric Andersen for their research assistance. All errors are our own.

[†]Department of Economics, University of Chicago, batzilis@uchicago.edu

[‡](Corresponding author) Department of Economics, Harvard University, sjaffe@fas.harvard.edu; Graduate Student Office, Littauer Center, 1805 Cambridge Street, Cambridge MA 02138. Jaffe gratefully acknowledges the support of a National Science Foundation Graduate Research Fellowship and the hospitality of the Becker Friedman Institute for Research in Economics at the University of Chicago.

[§]Department of Economics, University of Chicago and NBER, slevitt@uchicago.edu

[¶]Department of Economics, University of Chicago and NBER, jlist@uchicago.edu

[‖]Department of Economics, Harvard University, jpicel@fas.harvard.edu,

efficiently than lesser experienced players, and are more likely to win as a result. We also explore the degree to which the deviations from Nash predictions are consistent with various non-equilibrium models. We find that both a level-k framework and a quantal response model have explanatory power: whereas one group of people employ strategies that are close to k_1, there is also a set of people who use strategies that resemble quantal response.

JEL classification: C72, D03
Keywords: play in strategic settings, large-scale data set, Nash equilibrium, non-equilibrium strategies

Over the last several decades it would be difficult to find an idea that altered the social science landscape more profoundly than game theory. Across economics and its sister sciences, elements of Nash equilibrium are included in nearly every analysis of behavior in strategic settings. For their part, economists have developed deep theoretical insights into how people should behave in a variety of important strategic environments – from optimal actions during wartime to more mundane tasks such as how to choose a parking spot at the mall. Understanding whether people actually behave in accord with theoretical predictions, however, has considerably lagged behind. Although there are important tests of game theory in lab experiments (see, e.g., Dufwenberg and Gneezy (2000); Dufwenberg et al. (2010); Lergetporer et al. (2014); Sutter et al. (2013)), credibly testing whether behavior conforms to theory in the field has been difficult (yet, see Chiappori et al. (2002); Güth et al. (2003); Goette et al. (2012) and the cites therein).

In this paper, we take a fresh approach to studying strategic behavior in the field, exploiting a unique dataset that allows us to observe play while the information shown to the player changes. In particular, we use data from over one million matches of rock-paper-scissors[1] played on a historically popular Facebook application. Before each match (made up of multiple throws), players are shown a wealth of data about their opponent's past history: the percent of past first throws in a match that were rock, paper, or scissors, the percent of all throws that were rock, paper, or scissors, and all the throws from the

[1] Two players each play rock, paper, or scissors. Rock beats scissors; scissors beats paper; paper beats rock. If they both play the same, it is a tie. The payoff matrix is in Figure 2.

opponents' most recent five games. These data thus allow us to investigate whether, and to what extent, players use information.

The informational variation makes the strategy space for the game potentially much larger than a basic rock-paper-scissors game. We show, however, that in Nash Equilibrium, players must expect their opponents to mix equally across rock-paper-scissors – same as in the one-shot game. Therefore, a player has no use for information on the opponent's history when the opponent is playing Nash.

To the extent that an opponent systematically deviates from Nash, however, knowledge of that opponent's past history can potentially be exploited.[2] Yet, despite the simplicity of the game and the transparency of the information, it is not clear how one should utilize the information provided. Players can use the information to determine whether an opponent's past play conforms to Nash, but they do not observe the past histories of the opponent's previous opponents; without seeing what information the opponent was reacting to, it is hard to guess what non-Nash strategy the opponent may be using. Additionally, players are not shown information about their own past play, so if a player wants to exploit an opponent's reaction, he has to keep track of his own history of play.

Because of the myriad of possible responses, we start with a reduced-form analysis of the first throw in each match to describe how players respond to the provided information. We find that players use information: for example, they are more likely to play rock when their opponent has played less paper (which beats rock) or more scissors (which rock beats) on previous first throws. Players have a weak negative correlation across their own first throws. Overall, we find that most players, at some point in their histories, employ strategies consistent with Nash equilibrium. Even so, we do find considerable evidence of disequilibrium play; for example, 53% of players are reacting to information about their opponents' history.

[2] If the opponent is not playing Nash, then Nash is no longer a best response. In symmetric zero-sum games like RPS, deviating from Nash is costless if the opponent is playing Nash (since all strategies have an expected payoff of zero), but there is a profitable deviation from Nash if a player thinks he knows what non-Nash strategy his opponent is using.

This finding motivated us to adopt a structural approach to evaluate the performance of two well-known disequilibrium models: level-k and quantal response. The level-k model posits that players are of different types according to the depth of their reasoning about the strategic behavior of their opponents (Stahl, 1993; Stahl and Wilson, 1994, 1995; Nagel, 1995). Players who are k_0 do not take into account their opponents' strategies or incentives. This can either mean that they play randomly (e.g. Costa-Gomes and Crawford (2006)) or that they play some focal or salient strategy (e.g. Crawford and Iriberri (2007); Arad and Rubinstein (2012)). Players who are k_1 respond optimally to a k_0 player, which in our context means responding to the focal strategy of the opponent's (possibly skewed) historical distribution of throws; k_2 players respond optimally to k_1, etc.[3]

Level-k theory acknowledges the difficulty of calculating equilibria and of forming beliefs, especially in one shot games. It has been applied to a variety of laboratory games (e.g. Costa-Gomes et al. (2001); Costa-Gomes and Crawford (2006); Hedden and Zhang (2002); Crawford and Iriberri (2007); Ho et al. (1998)), but this is one of the first applications of level-k theory to a naturally-occurring environment (e.g. Bosch-Domenech et al. (2002); Ostling et al. (2011); Gillen (2009); Goldfarb and Xiao (2011); Brown et al. (2012)). We also have substantially more data than most other level-k studies, both in number of observations and in the richness of the information structure.

We adapt level-k theory to our repeated game context. Empirically, we use maximum likelihood to estimate how often each player plays k_0, k_1, and k_2, assuming that they are restricted to those three strategies. We find that the majority of play is best described as k_0 (about 74%). On average, k_1 is used in 19% of throws. The average k_2 estimate is 7.7%, but for only 12% of players do we reject at the 95% level that they never play k_2. Most players use a mixture of strategies, mainly k_0 and k_1.[4] We also find that 20% of players

[3]Since the focal k_0 strategies can be skewed, our k_1 and k_2 strategies usually designate a unique throw, which would not be true if k_0 were constrained to be a uniform distribution.

[4]As we discuss in Section 4 there are several reasons that may explain why we find lower estimates for k_1 and k_2 play than in previous work. Many players may not remember their own history, which is necessary for playing k_2. Also, given that k_0 is what players would

deviate significantly from $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ when playing k_0.

An interesting result is that play is more likely to be consistent with k_1 when the expected return to k_1 is higher. This effect is larger when the opponent has a longer history – that is, when the skewness in history is less likely to be noise. The fact that players respond to the level of the perceived expected (k_1) payoff, not just whether it is the highest, is related to another model of non-equilibrium play based on Quantal Response Equilibrium (QRE) theory (McKelvey and Palfrey, 1995). QRE posits that players' probability of using a pure strategy is increasing in the relative perceived expected payoff of that strategy. Because we think players differ in the extent to which they respond to information and consider expected payoffs, we do not impose the equilibrium restriction that the perceived expected payoffs are correct. This version of quantal response can be thought of as a more continuous version of a k_1 strategy. Rather than always playing the strategy with the highest expected payoff as under k_1, the probability of playing a strategy increases with the expected payoff. As the random error in quantal response approaches zero (or the responsiveness of play to the expected payoff goes to infinity) this converges to the k_1 strategy.

On average, we find that increasing the expected payoff to a throw by one standard deviation increases the probability it is played by 5.2 percentage points (more than one standard deviation). The coefficient is positive and statistically significant for 60% of players. To interpret these results, one must consider that if players were using the k_1 strategy, then we would also find that expected payoffs have a positive effect on probability of play. Similarly, if players used quantal response, many of their throws would be consistent with k_1 and our maximum likelihood analysis would indicate some k_1 play.

Therefore, this evidence does not allow us to formally state which model is a better fit for the data. To preform that task, we compare the model likelihoods to test whether k_1 or quantal response better explains play. The quantal response model is significantly better than the maximum likelihood for

most likely play if they were not shown the information (i.e. when they play RPS outside the application), it may be more salient than in other contexts.

17.8 percent of players, yet the k_1 model is significantly better for 18.4 percent of players. We interpret this result as suggesting that there are some players whose strategies are close to k_1 and a distinct set of players who use strategies resembling quantal response. In sum, our data paint the picture that there is a fair amount of equilibrium play, and when we observe disequilibrium play, extant models have power to explain the data patterns.

The remainder of the paper is structured as follows. Section 1 describes the Facebook application in which the game is played and presents summary statistics of the data. Section 2 describes the theoretical model underlying the game, and the concept (and implications) of Nash equilibrium in this setting. Section 3 explores how players respond to the information about their opponents' histories. Section 4 explains how we adapt level-k theory to this context and provides parameter estimates. Section 5 adapts the quantal response model to our setting and Section 6 uses maximum likelihood to compare the level-k and quantal response models. Section 7 concludes.

1 An Introduction to Roshambull

Rock-Paper-Scissors, also known as Rochambeau and jan-ken-pon, is said to have originated in the Chinese Han dynasty, making its way to Europe in the 18th century. To this day, it continues to be played actively around the world. There is even a world Rock-Paper-Scissors championship sponsored by Yahoo.[5]

The source of our data is an early Facebook 'app' called Roshambull,[6] which allowed users to play rock-paper-scissors against other Facebook users. It was a very popular app for its era with 340,213 users ($\approx 1.7\%$ of Facebook

[5]Rock-paper-scissors is usually played for low stakes, but sometimes the result carries with it more serious ramifications. During the World Series of Poker, an annual \$500 per person rock-paper-scissors tournament is held, with the winner taking home \$25,000. Rock-paper-scissors was also once used to determine which auction house would have the right to sell a \$12 million Cezanne painting. Christie's went to the 11-year-old twin daughters of an employee, who suggested "scissors" because "Everybody expects you to choose 'rock'." Sotheby's said that they treated it as a game of chance and had no particular strategy for the game, but went with "paper" (Vogel, 2005).

[6]The name is a combination of a bastardized spelling of Rochambeau and the name of the firm sponsoring the app, Red Bull.

users) playing at least one match in the first three months of the game's existence. Users played best-two-out-of-three matches for prestige points known as 'creds.' They could share their records on their Facebook page and there was a leader board with the top players' records.

To make things more interesting for players, before each match the app showed them a "scouting sheet" with information on the opponent's history of play.[7] In particular, the app showed each player the opponent's distribution of throws on previous first throws of a match (and the number of matches) and on all previous throws (and the number of throws), as well as the opponent's lifetime win record and a play-by-play breakdown of the opponent's previous five matches. It also shows the opponent's win-loss records and the number of creds wagered. Figure 1 shows a sample screenshot from the game.

Our dataset contains 2,636,417 matches, all the matches played between May 23rd, 2007 (when the program first became available to users) and August 14th, 2007. For each throw, the dataset contains a player ID, match number, throw number, throw type, and the time and date at which the throw was made.[8] This allows us to create complete player histories at each point in time. Most players play relatively few matches in our three month window: the median number of matches is 5 and the mean number is 15.34.[9]

Some of our inference depends upon having a large number of observations per player; for those sections, our analysis is limited to the 7751 "experienced" players for whom we observe at least 100 clean matches. They play an average of 195.6 matches; the median is 151 and the standard deviation is 141.8.[10] Because these are the most experienced players, their strategies may not be representative; one might expect more sophisticated strategies in this group

[7]Bart Johnston, one of the developers said, "We've added this intriguing statistical aspect to the game... You're constantly trying to out-strategize your opponent" (Facebook, 2010).

[8]Unfortunately we only have a player id for each player; there is no demographic information or information about their out-of-game connections to other players.

[9]We exclude the small fraction of player-pairs for which one player won an implausibly high share of the matches (suggesting collusion). We of course include those matches when forming the histories.

[10]For some analyses we only use players who have 100 clean games with the relevant strategies indicate a unique throw, so we use between 5732 and 7751 players.

relative to the Roshambull population as a whole.

For all of the empirical analysis we focus on the first throw in each match. Modeling non-equilibrium behavior on subsequent throws is more complicated because in addition to their opponent's history, a player may also respond to the prior throws in the match. Table 1 summarizes the play and opponents' histories shown in the first throw of each match, for both the entire sample and the experienced players.

2 Model of the game

A standard game of rock-paper-scissors is a simple 3×3 zero-sum game. The payoffs are shown in Figure 2. The only Nash Equilibrium is for players to mix $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ across rock, paper, and scissors. Because each match is won by the first player to win two throws, and players play multiple matches, the strategies in Roshambull are potentially substantially more complicated: players could condition their play on various aspects of their own or their opponents' histories. A strategy would be a mapping from (1) the match history for the current match so far, (2) one's own history of all matches played, and (3) the space of information one might be shown about one's opponent's history, onto a distribution of throws.

In addition, Roshambull has a matching process operating in the background, in which players from a large pool are matched into pairs to play a match and then are returned to the pool to be matched again. In the Appendix, we formalize Roshambull in a repeated game framework.

Despite this potential for complexity, however, the equilibrium strategies are still simple.

Proposition 1. *In any Nash Equilibrium, for every throw of every match, each player correctly expects his opponent to mix $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ over rock, paper, and scissors.*[11]

Proof. See the Appendix. □

[11] Players could use aspects of their history that are not observable to the opponent as a private randomization devices, but conditional on all information available to the opponent, they must be mixing $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$.

The proof uses the fact that it is a symmetric, zero sum game to show that players continuation values at the end of every match must be zero. Therefore players are only concerned with winning the match, and not with the effect of their play on their resulting history. We then show that for each throw in the match, if player A correctly believes that player B is not randomizing $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$, then player A has a profitable deviation.

Nash Equilibrium implies that players randomize $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ both unconditionally and conditional on any information available to the opponent. Out of equilibrium, players may condition their throws on their or their opponents' histories in a myriad of ways. The resulting play may or may not result in an unconditional distribution of play that differs substantially from $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$. In Section 3, we present evidence that 83% of players have throw distributions that do not differ from from $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$. Yet, when throw distributions are exploitable, players respond to their opponents' histories.[12]

3 Players respond to information

Before examining the data for specific strategies players may be using, we present reduced-form evidence that players respond to the information available to them. To keep the presentation clear and simple, for each analysis we focus on rock, but the results are similar for paper and scissors.

We start by examining the dispersion across players in how often they play rock. Figure 3 shows the distribution across experienced players of the fraction of their last 100 throws that are rock. It also shows this distribution for simulated players who play rock, paper and scissors with probability one-third on each throw. The distribution from the actual data is substantially more disperse than the simulations, suggesting that the fraction of rock played deviates from one-third more than one would expect from pure randomness. Doing a chi-squared test on all of players' throws we reject uniform random play for 17% of experienced players.

Given this dispersion in the frequency with which players play rock, we test

[12] We also find serial correlation both across throws within a match and across matches, which is inconsistent with Nash Equilibrium.

whether players respond to the information they have about their opponents tendency to play rock – the opponents' historical rock percentage. Table 2 bins opponents by their historical percent rock and shows the fraction of paper, rock, and scissors play. Note that the percent paper is increasing across the bins and percent scissors is decreasing. Paper goes from less than a third chance to more than a third chance (and scissors goes from more to less) right at the cutoff where rock goes from less often than random to more often than random. The percent rock a player throws does not vary nearly as much across the bins.

For a more quantitative analysis of how this and other information presented to players affects their play, Table 3 presents regression results. The dependent variable is binary, indicating whether a player throws rock. The coefficients all come from one regression. The first column is the effect for all players, the second column is the additional effect of the covariates for players in the restricted sample; the third column is the additional effect for those players after their first 99 games. For example, a standard deviation increase in the opponents historical fraction of scissors (.18) increases the probability than an inexperienced player plays rock by 4.5 percentage points $(100 \cdot .18 \cdot .2531)$ and for an experienced player who already played at least 100 games, the increase is 9 percentage points $(100 \cdot .18 \cdot (.2531 + .1409 + .1379))$. As expected, the effects of the opponents percent of first throws that were paper is positive and gets stronger with experience, and the effect for scissors is negative and gets stronger with experience. The effect of the opponent's distribution of all throws and lagged throws is less clear.[13]

The consistent and strong reactions to the opponent's distribution of first throws motivates our use of that variable in the structural models.

The fact that players respond to their opponents' histories makes their play somewhat predictable and potentially exploitable. To measure this exploitability, we run the regression from Table 3 on half the restricted sample and use

[13] If we run the regression with just the distribution of all throws or just the lags, the signs are as expected, but that seems to be mostly picking up the effect via the opponent's distribution of first throws.

the coefficients to predict the probability of playing rock on each throw for the other half of the restricted sample. We do the same for paper and scissors. We then calculate how often a player who was optimally responding to this predicted play would win, draw, and lose a throw. We compare these rates to the players in our sample, keeping in mind that predicted play is based largely on the opponent's history, so responding to it optimally would require that the opponent know his own history. Table 4 presents the results. If players bet \$100 on each throw, the average experienced player would win \$1.49 on the average throw. This is better than the \$.66 the average player wins,[14] but someone responding optimally to their predictability would win \$16.86 on average.

Given these incentives to exploit players' predictability, we want to check whether their opponents do. They do not appear to. Given the predicted probabilities of play for experienced players, we calculate the expected payoff to an opponent of playing rock. Table 5 bins throws by the expected payoff to playing rock and shows the distribution of opponent throws. The probability of playing rock bounces around – opponents are not more likely to play rock when the actual expected payoff is high; the predictability of players' throws is not effectively exploited.

Since players are mostly responding to their opponent's history, exploiting those response requires that a player remember her own history of play (since the game does not show one's own history). So it is perhaps not surprising that players' predictability is not exploited and therefore unsurprising that they react in a predictable manner. If we do the analysis at the player level, 53% of players significantly respond to their opponents' historical distributions of past throws.

Having described in broad terms how players react to the information presented, we turn to existing structural models to test whether play is consistent with these hypothesized non-equilibrium strategies.

[14]The average overall must be zero, but our cleaning of the data left us with .66% more wins than losses.

4 Level-k behavior

While level-k theory was developed to analyze single shot games, it is a useful framework for exploring how players incorporate information about their opponent.[15] The k_0 strategy is to ignore the information about one's opponent and play a (possibly random) strategy independent of the opponent's history. While much of the existing literature assumes that k_0 is uniform random, some studies assume that k_0 players use a salient or focal strategy. In this spirit, we allow players to randomize non-uniformly (imperfectly) when playing k_0 and assume that the k_1 strategy best responds to a focal strategy for the opponent — k_1 players best respond to the opponent's past distribution of first throws.[16] It seems natural that a k_1 player who assumes his opponent is non-strategic would use this description of past play as a predictor of future play. When playing k_2, players assume that their opponents are playing k_1 and respond accordingly.

Given players' assumptions about their opponents' play, their strategies then depend on the value function they are maximizing. We assume that players are myopic and ignore the effect of their throw on their continuation value.[17] This approach is consistent with the literature that analyzes some games as "iterated play of a one-shot game" instead of as an infinitely repeated game (Monderer and Shapley, 1996). More generally, we think it is a reasonable simplifying assumption. While not impossible, it is hard to imagine how one would manipulate one's history to affect future payoffs with an effect large enough to outweigh the effect on this period's payoff.[18]

[15]Though players play multiple games, they might struggle to form accurate beliefs about opponents' strategies since players are playing against many different opponents each of whom may be using a complicated mixed strategy.

[16]The reduced form results indicate that players react much more strongly to the distribution of first throws than to the other information provided.

[17]In the proof of Proposition 1 we show that in Nash Equilibrium, histories do not affect continuation values, so in equilibrium it is a result, not an assumption, that players are myopic. However, out of Nash Equilibrium, it is possible that what players throw now can affect their probability of winning later rounds.

[18]One statistic that we thought might affect continuation values is the skew of a player's historical distribution. As a player's history departs further from random play, the more opportunity for opponent response and player manipulation of opponent response. We ran

Formal definitions of the different level-k strategies in our context are as follows:

Definition. When a player uses the k_0 *strategy* in a match, his choice of throw is unaffected by his history or his opponent's history.

We should note that using k_0 is not necessarily unsophisticated. It could be playing the Nash equilibrium strategy. However there are two reasons to think that k_0 might not represent sophisticated play. First, for some players the frequency distribution of their k_0 play differs significantly from $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$, suggesting that if they are trying to play Nash, they are performing poorly. Second, more subtly, it is not sophisticated to play the Nash equilibrium if your opponents are failing to play Nash. With most populations who play the beauty contest game, people who play Nash do not win (Nagel, 1995). In RPS, if there is a possibility that one's opponent is playing something other than Nash, there is a strategy that has a positive expected return, whereas Nash always has a zero expected return. (If it turns out the opponent is playing Nash, then every strategy has a zero expected return and so their is little cost to trying something else.) Given that some players differ from $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ when playing k_0 and most don't always play k_0, Nash is frequently not a best response.

Definition. When a player uses the k_1 *strategy* in a match, he plays the throw that has the highest expected payoff on this throw if his opponent randomizes according to the opponent's historical distribution of first throws.[19]

We have not specified how a player using k_0 chooses a throw, but provided the process is not changing over time, his past throw history is a good predictor

multinomial logits for each player on the effect of own history skewness on the probability of winning, losing or drawing. The coefficients were significant for less than 5% of players. The mean coefficient implied that if a player's skewness is a standard deviation higher (relative to the population) her probability of winning is .37 percentage points higher. This provides some support to our assumption that continuation values are not a primary concern.

[19]Sometimes opponents' distributions are such that there are multiple throws that are tied for the highest expected payoff. For our baseline specification we ignore these throws. As a robustness check we define alternative k_1-strategies where one throw is randomly chosen to be the k_1 throw when payoffs are tied or where both throws are considered consistent with k_1 when payoffs are tied. The results to do not change substantially.

of play in the current match. To calculate the k_1 strategy for each throw, we calculate the expected payoff to each of rock, paper, and scissors against a player who randomizes according to the distribution of the opponent's history. The k_1 strategy is the throw that has the highest expected payoff.

Definition. When a player uses the k_2 *strategy* in a match, he plays the throw that is the best response if his opponent randomizes between the throws that maximize expected payoff against the player's own historical distribution.

The k_2 strategy is to play "the best response to the best response" to one's *own* history. In this particular game k_2 is in some sense harder than k_1 because the software shows only one's opponent's history, but players could keep track of their own history.

Having defined the level-k strategies in our context, we now turn to the data for evidence of level-k play.

4.1 Reduced-form evidence for level-k play

One proxy for k_1 and k_2 play is players choosing throws that are consistent with these strategies. Whenever a player plays k_1 (or k_2) her throw is consistent with that strategy. However, the converse is not true. Players playing the NE strategy of $\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ would be consistent with k_1 a third of the time (on average).

For each player we calculate the fraction of throws that are k_1-consistent; these fractions are upper bounds on the amount of k_1 play. The highest percentage of k_1-consistent behavior for an individual in our restricted sample is 84.6 percent, indicating that no player uses k_1 consistently. Figure 4a shows the distribution of the fraction of k_1-consistency across players. It suggests that at least some players use k_1 at least some of the time: the distribution is to the right of the vertical $\frac{1}{3}$-line and there is a right tail. To complement the graphical evidence, we formally test whether the observed frequency of k_1-consistent play is significantly greater than expected under random play. Using this test, we can reject the null of no k_1 play at a 95 percent confidence interval for 71.7 percent of players in the sample.

Given that players seem to play k_1 some of the time, players could benefit from playing k_2. Figure 4b shows the distribution of the fraction of actual throws that are k_2-consistent. Perhaps unsurprisingly given that players are not shown the necessary information, we do not find much evidence of k_2 play. The observed frequency of k_2 play is significantly greater than expected with random play for only 7.5 percent of players, barely more than the 5% we would find if no one played k_2.

If we assume that players use either k_0, k_1, or k_2 we can use the percentage of throws that are consistent with neither k_1 nor k_2 to obtain a lower bound on how often each player is playing k_0. We calculate $\underline{k}_0 = 1 - \bar{k}_1 - \bar{k}_2$. We do not expect this bound to be tight because, in expectation, a randomly chosen k_0 play will be consistent with either the k_1 or k_2 strategy relatively often. The mean lower bound across players is 37 percent. The minimum is 9.3 percent and the maximum is 74 percent.

4.1.1 Multinomial Logit

Before turning to the structural model, we can use a multinomial logit model to explore whether a throw being k_1-consistent increases the probability that a player choses that throw. For each player, we estimate a multinomial logit where the utilities are

$$U_i = \alpha_i + \beta \cdot 1\{k_1 = i\} + \varepsilon_i,$$

where $i = r, p, s$ and $1\{k_1 = i\}$ is an indicator for when the k_1-consistent action is to throw i. Figure 5 shows the distribution of β's across players. The mean is .532.

The marginal effect varies slightly with the baseline probabilities, $\frac{\partial Pr\{i\}}{\partial x_i} = \beta \cdot Pr\{i\}(1 - Pr\{i\})$, but is approximately $\frac{1}{3}\left(1 - \frac{1}{3}\right) = \frac{2}{9}$ times the coefficient. Hence, on average, a throw being k_1-consistent means it is 12 percentage points more likely to be played. Given that the standard deviation across players in the percent of rock, paper, or scissors throws is about 5 percentage points, this is a large average effect. The individual level effect is positive and significant for 61 percent of players.

4.2 Maximum likelihood estimation of a structural model of level-k thinking

The results presented in the previous sections provide some evidence as to what strategies are being employed by the players in our sample, but they do not allow us to identify with precision the frequency with which strategies are employed. To obtain point estimates of each player's proportion of play by level-k, along with standard errors, we need additional assumptions.

Assumption 1. *All players use only the k_0, k_1, or k_2 strategies in choosing their actions.*

Assumption 1 restricts the strategy space, ruling out any approach other than level-k, and also restricting players not to use levels higher than k_2. We limit our modeling to levels k_2 and below, both for mathematical simplicity and because there is little reason to believe that higher levels of play are commonplace, both based on the low rates of k_2 play in our data, and rarity of k_3 and higher play in past experiments.[20]

Assumption 2. *Whether players chose to play k_0, k_1, or k_2 on a given throw is independent of which throw (rock, paper, or scissors) each of the strategies would have them play.*

Assumption 2 implies, for example, that the likelihood that a player chooses to play k_2 will not depend on whether it turns out that the best k_2 action is rock or is paper. This independence is critical to the conclusions that follow. Note that Assumption 2 does not require that a player commit to having the same probabilities of using k_0, k_1, and k_2 strategies across different throws.

Given these assumptions we can calculate the likelihood of observing a given throw in terms of 4 parameters, given in Table 6. Given these parameters, the probability of observing a given throw i is

$$\hat{k}_1 \cdot 1\{k_1 = i\} + \hat{k}_2 \cdot 1\{k_2 = i\} + (\hat{r}_0 \cdot 1\{i = r\} + \hat{p}_0 \cdot 1\{i = p\} + \hat{s}_0 \cdot 1\{i = s\}),$$

[20] As an aside, in the case of rock-paper-scissors the level $k + 6$ strategy is identical to the level k strategy for $k \geq 1$, so it is impossible to identify levels higher than 6. This also implies that the k_1 play we observe could in fact be k_7 play, but we view this as highly unlikely.

where $1\{\cdot\}$ is an indicator function, equal to one when the statement in braces is true and zero otherwise. This reflects the fact that the throw will be i if the player plays k_1 and the k_1 strategy says to play i ($\hat{k}_1 \cdot 1\{k_1 = i\}$) or the player plays k_2 and the k_2 strategy says to play i ($\hat{k}_2 \cdot 1\{k_2 = i\}$) or the player plays k_0 and chooses i ($\hat{r}_0 \cdot 1\{i = r\} + \hat{p}_0 \cdot 1\{i = p\} + \hat{s}_0 \cdot 1\{i = s\}$).

For each player, the overall log-likelihood depends on twelve statistics from the data. For each throw type ($i = R, P, S$), let n^i_{12} be the number of throws that are of type i and consistent with k_1 and k_2, n^i_1 the number of i throws consistent with just k_1, n^i_2 the number of i throws consistent with just k_2, and n^i_0 the number of i throws consistent with neither k_1 nor k_2. Given these statistics, the log-likelihood function is

$$\mathcal{L}(\hat{k}_1, \hat{k}_2, \hat{r}_0, \hat{p}_0, \hat{s}_0) = \sum_{i=r,p,s} \left(n^i_{12} \ln(\hat{k}_1 + \hat{k}_2 + \hat{\imath}_0) + n^i_1 \ln(\hat{k}_1 + \hat{\imath}_0) + n^i_2 \ln(\hat{k}_2 + \hat{\imath}_0) + n^i_0 \ln(\hat{\imath}_0) \right).$$

For each player we use maximum likelihood to estimate $\left(\hat{k}_1, \hat{k}_2, \hat{r}_0, \hat{p}_0, \hat{s}_0\right)$. Given these estimates, the standard errors are calculated analytically from the inverse of the Hessian of the likelihood function.

Table 7 summarizes the estimates of k_0, k_1, and k_2: the average player uses k_0 for 73.8 percent of throws, k_1 for 18.5 percent of throws and k_2 for 7.7 percent of throws. Weighting by the precision of the estimates or by the number of games does not change these results substantially. As the minimums and maximums suggest, these averages are not the result of some people always playing k_1 while others always play k_2 or k_0. Most players mix, using a combination of mainly k_0 and k_1.[21]

Table 8 reports the share of players for whom we can reject with 95 percent confidence their never playing a particular level-k strategy. Almost all players (93 percent) appear to use k_0 at some point. About 63 percent of players use k_1 at some stage, but we can reject exclusive use of k_1 for all but two out of 6,399 players. Finally, only about 12 percent of players are we confident use

[21] Other work has found evidence of players mixing levels of sophistication of across different games, e.g. Georganas et al. (2010).

k_2.

For each player, we can also examine the estimated fraction of rock, paper, and scissors when they play k_0. The distribution differs significantly from random uniform for 1257 or 20% of players, similar to the fraction of players whose raw throw distributions differ significantly from uniform (17%).

4.3 Cognitive Hierarchy

The idea that players might use a distribution over the level-k strategies naturally connects to the Cognitive Hierarchy model of Camerer et al. (2004). They also model players as having different levels of reasoning, but the higher types are more sophisticated than in level-k. The strategies for ch_0 and ch_1 are the same as k_0 and k_1; ch_2 assumes that other players are playing either ch_0 or ch_1, in proportion to their actual use in the population, and best responds to that mixture. To test if this more sophisticated version of two levels of reasoning fits the data better, we do another maximum likelihood estimation. The definitions of ch_0 and ch_1 are the same as k_0 and k_1.

Definition. When a player uses the ch_2 *strategy* in a match, he plays the throw that is the best response to the opponent playing according to the opponent's historical distribution 79.94% of the time and randomizing between the throws that maximize expected payoff against the player's own historical distribution the other 20.05% of the time.

The percents come from observed frequencies in the level-k estimation. When players play either k_0 or k_1, they play k_0 $\frac{73.79}{73.79+18.52} = 79.94\%$ of the time.[22]

Analogous to Assumptions 1 and 2 above, we assume that players use only ch_0, ch_1 and ch_2, and that which strategy they chose is independent of what throw the strategy dictates.

Table 9 summarizes the estimates: the average player uses ch_0 for 74.9 percent of throws, ch_1 for 16.31 percent of throws, and ch_2 for 8.85 percent of throws. Weighting by the precision of the estimates or by the number of games a player plays does not change these substantially. These results are

[22]To fully calculate the equilibrium, we could repeat the analysis using the frequencies of ch_0 and ch_1 found below and continue until the frequencies converged, but the numbers are very similar, so we do not think this computationally intense exercise is necessary.

similar to what we found for level-k strategies, suggesting that the low rates of using two iterations of reasoning were not a result of restricting that strategy to ignoring k_0 play.

4.4 Naive level-k strategies

Even if a player expects his opponent to play as she did in the past, he may not calculate the expected return to each strategy. Instead he may employ the simpler strategy of playing the throw that beats the opponent's most common historical throw. Put another way, he may only consider maximizing his probability of winning instead of weighing it against the probability of losing as is done in an expected payoff calculation. We consider this play *naive* and define alternative versions of k_1 and k_2 accordingly.

Definition. When a player uses the *naive* k_1 *strategy* in a match, he plays the throw that will beat the throw that his opponent has played most frequently in the past.

Definition. When a player uses the k_2 *strategy* in a match, he plays the throw that is the best response if his opponent plays the throw that beats the throw that the player has played most frequently in the past. If two throws are tied for most frequent in the player's historical distribution, he assumes his opponent randomizes between the two throws that beat one of the throws tied for most frequent.

Table 10 summarizes the estimates for naive play. The average player uses k_0 for 72.3 percent of throws, *naive* k_1 *strategy* for 21.1 percent of throws and *naive* k_2 *strategy* for 6.7 percent of throws. As before, weighting by the precision of the estimates or by the number of games a player plays does not change these results substantially. Most players use a mixed strategy, mixing primarily over k_0 and *naive* k_1 *strategy.*

Compared with our standard level-k model, the results for our naive level-k model show slightly more use of the *naive* k_1 *strategy.* The average player used it for 21.1 percent of throws compared to 18.6 percent for standard k_1 in the former model. In the naive level-k model, players used k_0 and the *naive* k_2 *strategy* slightly less often compared to their standard level-k counterparts.

The opposite naive strategy would be for players to minimize their probability of losing, playing the throw that is least likely to be beat. Running the same model for that strategy we find almost no evidence of k_1 or k_2 play, suggesting that players are more focused on the probability of winning.

4.5 Comparison to the literature

All three of these related models suggest that players of Roshambull use considerably fewer levels of iteration in their reasoning process compared to estimates from other games and other experiments. Bosch-Domenech et al. (2002) found that less than a fourth of the players who used the k-strategies discussed in this paper were k_0 players. Whereas, depending on the model, we found between 72.25% and 74.85% of plays involved zero iterations of reasoning. Camerer et al. (2004) suggest that players iterate 1.5 steps on average in many games. In comparison, in our level-k model we find that our average player uses $1 * .185 + 2 * .077 = .339$ levels of iterated rationality. Stahl (1993) reported that an insignificant fraction of players were k_0, 24 percent were k_1 players, 49 percent were k_2 players, and the remaining 27 percent were "Nash types." In contrast, we found that the majority of plays were k_0 (ch_0) and that k_1 (ch_1) outnumbered k_2 (ch_2), though in this game k_0 is closer to Nash than either k_1 or k_2. The dearth of k_2 play is especially striking in our context given the high returns to playing k_2.

One explanation for the differences between our results and the past literature is that most of the players do not deviate substantially from equilibrium play, making the expected payoffs to k_1 relatively small. Also, the set-up of rock-paper-scissors does not suggest a level-k thinking mindset as strongly as the p-beauty contest games or other games specifically designed to measure level-k behavior. Our more flexible definition of k_0 play may also explain its higher estimate. The low level of k_2 play is likely a result of the fact that the Facebook application did not show players their histories so players had to keep track of that on their own in order to effectively play k_2.

Another explanation is that we restrict the strategy space to exclude both Nash Equilibrium and different ways in which the players can react rationally to their opponent's information. It seems players respond more to the first throw history than other information, but there may be many other strategies which are rationalizable in ways which we do not model. Bosch-Domenech et al. (2002), for example, considered equilibrium, fixed point, degenerate and non-degenerate variants of iterated best response, iterated dominance, and even experimenter strategies. Not all of these translate into the RPS set-up, but any strategies that our model left out might look like k_0 play when the strategy space is restricted.

4.6 When are players' throws consistent with k_1?

Though we find relatively low levels of k_1 play, we do find some and the result that many of the players seem to be mixing strategies raises the question of when they chose to play k_0, k_1, and k_2. Our structural model assumes that which strategy players chose is independent of the throw dictated by each of the strategies. It does not require that which strategy they chose be independent of the expected payoffs, but the MLE model cannot give us insight into how expected payoffs may affect play. This is partially because the MLE model does not allow us to categorize individual throws as being a given strategy.

Therefore, to try to get at when players use k_1, we return to using k_1-consistency as a proxy for possible k_1 play. We test two hypotheses. First, the higher the expected payoff to playing k_1, the more likely a player is to play k_1. For example, if an opponent is k_0, the expected returns to playing k_1, relative to playing randomly, are much higher if the opponent's history (or expected distribution) is 40 percent rock, 40 percent paper, 20 percent scissors than if it is 34 percent rock, 34 percent paper, 32 percent scissors.

The second hypothesis is that a player will react more to a higher k_1 payoff when his opponent has played more games. A 40 percent rock, 40 percent paper, 20 percent scissors history is more informative if it is based 100 past throws than if it is based on only 10 throws.[23]

We also analyze whether these effects vary by player experience; we interact all the covariates with whether a player is in the restricted sample (they eventually play $\geq$ 100 matches) and whether they have played 100 matches

[23] Similar predictions could be made about k_2 play; however, since we find that k_2 is used so little, we do not model k_2 play in this section

before the current match.

Tables 11 present empirical results from testing these hypotheses. The k_1 payoff is the expected payoff to playing k_1 assuming the opponent randomizes according to his history. Its standard deviation is .25, so for inexperienced players a one standard deviation increase in payoff to the k_1 strategy, increases the probability the throw is k_1-consistent by 1.8 percentage points when opponents have a short history, 10 percentage points when opponents have a medium history and 15 percentage points when opponents have played over 94 games. Given that 45% of all throws are k_1-consistent, these latter two effects are substantial. Experienced players react slightly less to the k_1-payoff when opponents have short histories, but their reactions to opponents with medium or long histories are somewhat larger.

If we run a logit analysis at the player level of the effect of k_1 payoff and opponent history length on k_1 consistency, the mean marginal effect across players is very close to the OLS coefficients. We only report the overall results because the individual analyses lack power — very few of the individual-level coefficients are statistically different from zero.

While we expect the correlation between opponent's history-length and playing k_1 to be negative – since longer histories are less likely to show substantial deviation from random – we do not have a good explanation for why the direct effect of opponent's history length is negative, even when controlling for the k_1 payoff. Perhaps the players are more wary of trying to exploit a more experienced player.

5 Quantal Response

The above evidence that k_1-consistent play is more likely when the expected payoff is higher, naturally leads us to a model of play that is more continuous. In some sense level-k strategies are all or nothing. If a throw has the highest expected payoff against the opponent's historical distribution, then the k_1 strategy says to play it, even if expected payoff is very small. A related, but different strategy is for players to choose each throw with a probability that is increasing in its expected payoff against the opponent's historical distribution of play. This is related to the idea behind Quantal Response Equilibrium (McKelvey and Palfrey, 1995), but without requiring that players be in equilibrium. In this context, players doing one iteration of reasoning would have probabilities of play

$$P_i \propto \exp(\alpha_i + \beta E[i \mid \text{opponent's distribution}]).$$

Their probability of playing a given throw is increasing in expected return to that throw, assuming the opponent plays according to his historical distribution. This smooths the threshold response of the k_1 strategy into a more continuous response.[24]

Figure 6 shows the distribution of coefficients across individuals. The mean coefficient is .01.[25] The expected return is the percent chance of winning minus the percent chance of losing, so it ranges from -100 to 100. The standard deviation is 23.2, so, on average, a standard deviation increase in the expected return to an action, increases the percent chance it is played by

$$23.2 \cdot .01 \cdot 100 \cdot \frac{2}{9} = 5.2 \text{ percentage points.}$$

(We multiply by 100 to convert to percentages and by 2/9 to evaluate the margin at the means.) The standard deviation across players in the percent of the time they play a throw is 5%, so this is significant, but not a huge effect.

The expected return has a significant effect for 60.7% of players. The mean of the effect size conditional on being significant is .025. Converting to margins, this corresponds to a standard deviation increase in expected return resulting in a 12.9 percentage point increase in the probability of playing a given throw, which is quite large.

[24] A second level of reasoning would expect opponents to play according to the distribution induced by one's own history and would play with probabilities proportional to the expected payoff against that distribution. However, given the low levels of k_2 play we find and the econometric difficulties of including own history in the logit, we only analyze the first iteration of reasoning.

[25] In the reduced form results (Table 5) we showed that players did not respond to the expected payoff calculated from *predicted* opponent play, whereas this shows that players do respond to expected payoffs calculated from *historical* opponent play.

6 Likelihood Comparison

Which is a better model of player behavior, the discrete "if it's the k_1 throw, play it" or the more continuous "if its k_1 payoff is higher, play it with higher probability"? Since the strategies are similar, if players were using one there would still be some evidence for the other, so we use a likelihood test to see which model better fits players' behavior. For each player we can calculate the likelihood of observing the set of throws he plays given the level-k maximum likelihood model and given the quantal response model. To facilitate the comparison, we estimate a version of the maximum likelihood model with no k_2 so that each model has three independent parameters.

If we assume that one of two models, level-k (LK) and QE, generated the data and have a flat prior over which it was, then the probability that it was QE is

$$P(\text{ QE } \mid \text{ data }) = \frac{P(\text{ data } \mid \text{ QE })}{P(\text{ data } \mid \text{ LK }) + P(\text{ data } \mid \text{ QE })}.$$

Figure 7 plots the distribution of the probability of the quantal response model across players. The probability is less than .5 for 3,711 players, so for 55.6% of players the quantal response is a better model. More interestingly, there are substantial numbers of players both to the left of .05 and to the right of .95. For 1,228 players (18.4%) the MLE model is a statistically better fit and for 1,186 players (17.8%) the quantal response is a statistically better fit. This suggests some players' strategies focus more on whether the throw has the highest expected return (k_1) and others' strategies respond more to the level of the expected return (quantal response).

7 Conclusion

The 20th century witnessed several break through discoveries in economics. Arguably the most important revolved around understanding behavior in strategic settings, which originated with John von Neumann's (1928) minimax theorem. In zero-sum games with unique mixed-strategy equilibria, minimax logic dictates that strategies should be randomized to prevent exploitation by one's opponent. The work of Nash enhanced our understanding of optimal play in games, and several theorists since have made seminal discoveries.

We take this research in a different direction by analyzing an enormous set of naturally generated data on rock-paper-scissors with information about opponents' past play. In doing so, we are able to explore the models — both equilibrium and non-equilibrium — that best describe the data. While we find that most people employ strategies consistent with Nash, at least some of the time, there is considerable deviation from equilibrium play. Adapting level-k thinking to our repeated game context, we use maximum likelihood to estimate the frequency with which each player uses k_0, k_1 and k_2. We find that about three-quarters of all throws are best described as k_0. A little less than one-fifth of play is k_1, with k_2 level play accounting for less than one-tenth of play. Interestingly, we find that most players are mixing over at least two levels of reasoning. Since players mix across levels, we explore when they are most likely to play k_1. We find that consistency with k_1 is increased when the expected return to k_1 is higher.

We also explore the quantal response model. Our adapted version of quantal response has players paying attention to the expected return to each strategy. We find that a one standard deviation increase in expected return increases the probability of a throw by 5.2 percentage points. In addition, for about a fifth of players the quantal response model fits significantly better than the level-k model, but for another one-fifth the level-k model fits significantly better. It seems that some players focus on the levels of the expected returns, while others focus on which throw has the highest expected return.

Beyond theory testing, we draw several methodological lessons. First, while our setting is very different from the single-shot games that level-k theory was originally developed, the evidence that players mix across strategies raises questions for experiments that attempt to categorize players as a k-type, based on only a few instances of play. Second, with large data sets subtle differences in theoretical predictions can be tested with meaningful power. As the internet continues to provide unique opportunities for such large-scale data, we hope that our study can serve as a starting point for future explorations of behavior in both strategic and non-strategic settings.

References

Arad, A. and Rubinstein, A. (2012). The 11-20 money request game: A level-k reasoning study. *The American Economic Review*, 102(7):3561–3573.

Bosch-Domenech, A., Montalvo, J. G., Nagel, R., and Satorra, A. (2002). One, two, (three), infinity,...: Newspaper and lab beauty-contest experiments. *The American Economic Review*, 92(5):1687–1701.

Brown, A. L., Camerer, C. F., and Lovallo, D. (2012). To review or not to review? Limited strategic thinking at the movie box office. *American Economic Journal: Microeconomics*, 4(2):1–26.

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.

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.

Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235.

Costa-Gomes, M. A. and Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: An experimental study. *The American Economic Review*, 96(5):1737–1768.

Crawford, V. P. and Iriberri, N. (2007). Fatal attraction: Salience, naivete, and sophistication in experimental "hide-and-seek" games. *The American Economic Review*, 97(5):1731–1750.

Dufwenberg, M. and Gneezy, U. (2000). Measuring beliefs in an experimental lost wallet game. *Games and Economic Behavior*, 30(2):163–182.

Dufwenberg, M., Sundaram, R., and Butler, D. J. (2010). Epiphany in the game of 21. *Journal of Economic Behavior & Organization*, 75(2):132–143.

Facebook (2010). Developer partner case study: Red bull application.

Georganas, S., Healy, P. J., and Weber, R. A. (2010). On the persistence of strategic sophistication. Ohio State Working Paper.

Gillen, B. (2009). Identification and estimation of level-k auctions. Working Paper.

Goette, L., Huffman, D., Meier, S., and Sutter, M. (2012). Competition between organizational groups: Its impact on altruistic and antisocial motivations. *Management science*, 58(5):948–960.

Goldfarb, A. and Xiao, M. (2011). Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. *The American Economic Review*, 101(7):3130–3161.

Güth, W., Schmidt, C., and Sutter, M. (2003). Fairness in the mail and opportunism in the internet: A newspaper experiment on ultimatum bargaining. *German Economic Review*, 4(2):243–265.

Hedden, T. and Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, 85(1):1–36.

Ho, T.-H., Camerer, C., and Weigelt, K. (1998). Iterated dominance and iterated best response in experimental "p-beauty contests". *The American Economic Review*, 88(4):947–969.

Lergetporer, P., Angerer, S., Glätzle-Rützler, D., and Sutter, M. (2014). Third-party punishment increases cooperation in children through (misaligned) expectations and conditional cooperation. *Proceedings of the National Academy of Sciences*, page 201320451.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38.

Monderer, D. and Shapley, L. S. (1996). Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1):258–265.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326.

Ostling, R., Wang, J., Chou, E., and Camerer, C. (2011). Testing game theory in the field: Swedish LUPI lottery games. *American Economic Journal: Microeconomics*, 3(3):1–33.

Stahl, D. O. (1993). Evolution of smart$_n$ players. *Games and Economic Behavior*, 5(4):604–617.

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.

Stahl, D. O. and Wilson, P. W. (1995). On players models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.

Sutter, M., Kocher, M. G., Glatzle-Rutzler, D., and Trautmann, S. T. (2013). Impatience and uncertainty: Experimental decisions predict adolescents' field behavior. *The American Economic Review*, 103(1):510–531.

Vogel, C. (2005). Rock, paper, payoff: Child's play wins auction house an art sale. *The New York Times*. April 29.

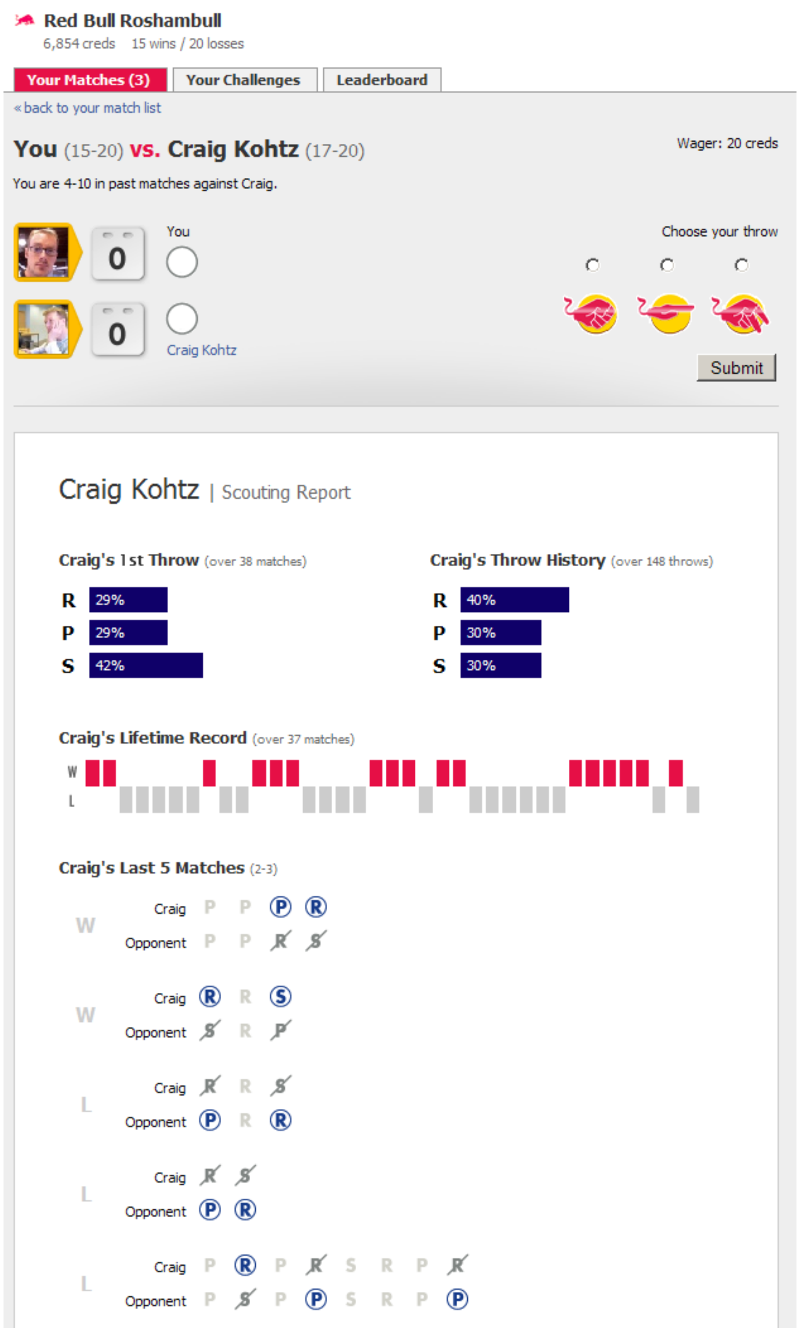

Figure 1: Screenshot of the Roshambull App.

Note: This is a sample screenshot from the start of a game of Roshambull.

		Player 2:		
		Rock	Paper	Scissors
	Rock	(0,0)	(-1,1)	(1,-1)
Player 1:	Paper	(1,-1)	(0,0)	(-1,1)
	Scissors	(-1,1)	(1,-1)	(0,0)

Figure 2: Payoffs for a single throw of rock-paper-scissors.

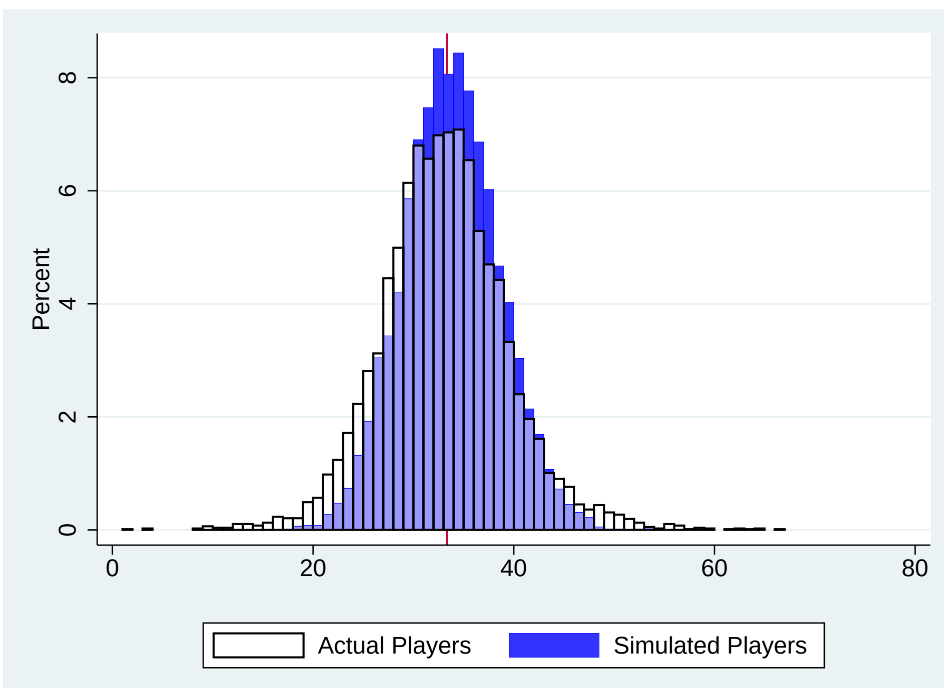

Figure 3: Observed and Simulated Percent of Throws that are Rock

Note: For each of the 7751 players with at least 100 matches we calculate the percent of his or her last 100 throws that were rock (white distribution). We overlay this on the distribution of the percent of throws that are rock for 7751 simulated players who each play 100 throws and throw rock, paper and scissors each with a probability one-third (blue distribution).

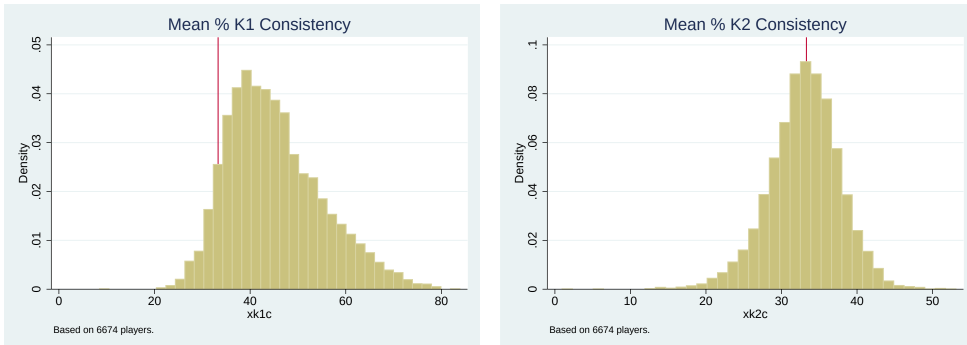

(a) Fraction of k_1-consistent throws. (b) Fraction of k_2-consistent throws.

Figure 4: Level-k consistency

Note: These graphs show the distribution across players of the fraction of throws that are k_1- and k_2-consistent. They include the 6674 players who have 100 games with well-defined k_1 and k_2 strategies. The vertical line indicates $\frac{1}{3}$, which we would expect to be the center of the distribution if throws were random.

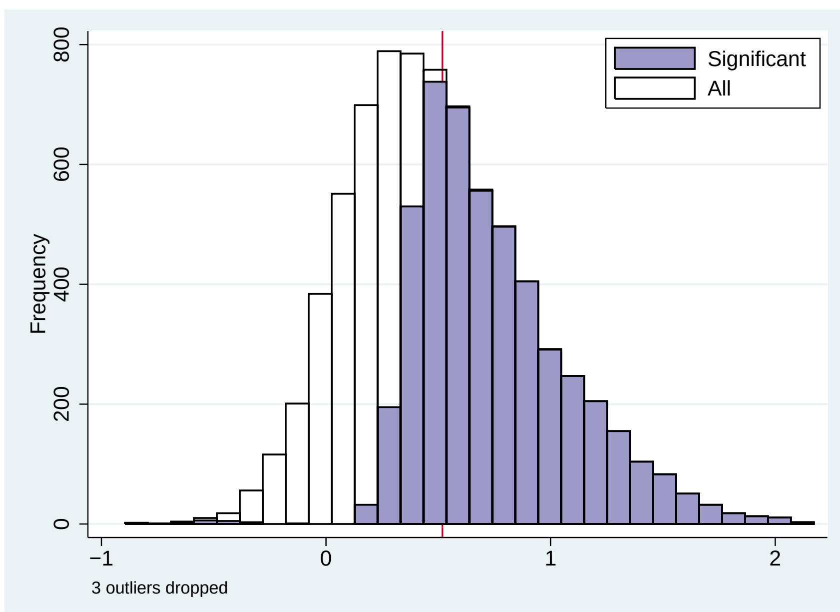

Figure 5: Distribution across players of the coefficient in the multinomial logit.

Note: The graph shows distribution across 6670 players of the coefficient from the multinomial logit $U_i = \alpha_i + \beta \cdot 1\{k_1 = i\} + \varepsilon_i$, where $i = r, p, s$ and $1\{k_1 = i\}$ is an indicator for when i is the k_1-consistent thing to do.

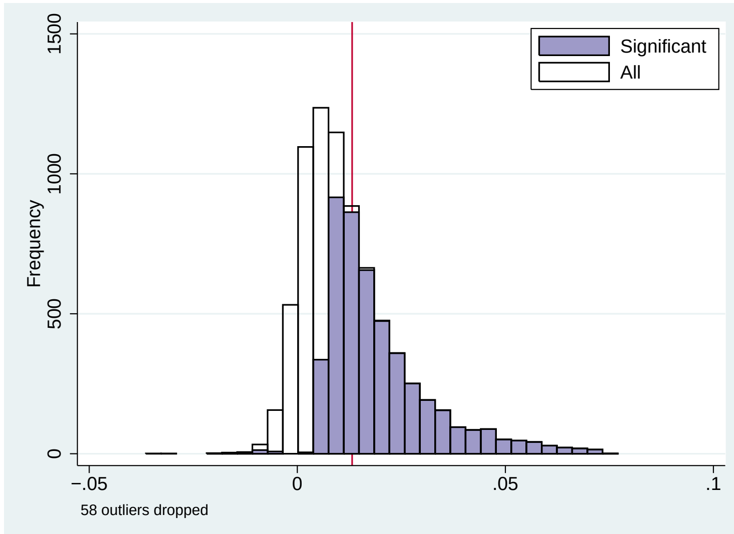

Figure 6: Distribution across players of the coefficient in the quantal response model.

Note: The distribution across 6670 players.

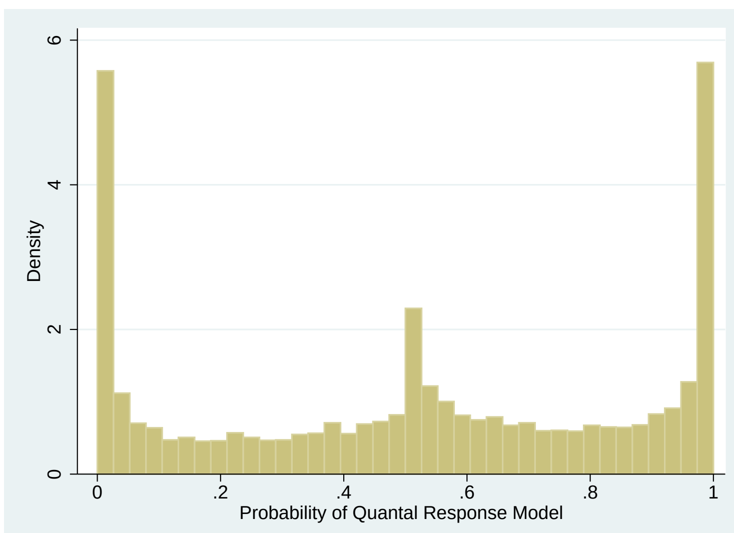

Figure 7: Distribution across players of the probability that the quantal response model and not the maximum likelihood model generated the data.

Note: The distribution across 6670 players, assuming a flat prior.

Table 1: Summary Statistics of First Throws

Variable	Full Sample		Restricted Sample	
	Mean	(SD)	Mean	(SD)
Throw Rock (%)	33.39	(47.16)	32.57	(46.86)
Throw Paper (%)	34.78	(47.63)	34.57	(47.56)
Throw Scissors (%)	31.83	(46.58)	32.86	(46.97)
Opp's Historical %Rock	34.15	(18.40)	33.59	(13.71)
Opp's Historical %Paper	35.12	(18.12)	34.76	(13.45)
Opp's Historical %Scissors	30.73	(17.03)	31.65	(12.83)
Opp's Historical Skew	9.68	(17.29)	5.39	(12.43)
Opp's Historical %Rock (all)	35.36	(11.49)	34.81	(8.58)
Opp's Historical %Paper (all)	34.03	(11.25)	34.10	(8.35)
Opp's Historical %Scissors (all)	30.61	(10.62)	31.09	(7.98)
Opp's History Length	59.06	(125.12)	99.13	(162.14)
Total observations	4,596,464		1,471,159	

The Restricted Sample uses data only from players who play at least 100 matches. The first 3 variables are dummies for when a throw is rock (R), paper (P), or scissors (S). The next 3 are the percentages opponent's past first throws that were R, P or S. The "all throws" are the corresponding percentages for all of the opponents' past throws. Skew measures the extent to which the opponent's history of first throws deviates from random. Opp's Historical Length is the number of previous matches the opponent played.

Table 2: Response to Percent Rock

Opp's Historical %Rock	Throws (%)			
	Paper	Rock	Scissors	N
0% - 25%	24.33	34.12	41.55	216542
25% - 30%	25.75	34	40.25	259141
30% - $33\frac{1}{3}$ %	29.84	33.98	36.17	259427
$33\frac{1}{3}$ % - 37%	35.41	33.48	31.1	294979
37% - 42%	42.74	31.38	25.87	185369
42% - 100%	51.11	27.61	21.28	219066

Note: Using only players with at least 100 matches, we bin matches by the opponent's historical percent of rock throws prior to the match. For each range of opponent's historical percent rock, we show the distribution of rock, paper, and scissors throws the players use.

Table 3: Probability of Playing Rock

Dependent Variable: Dummy for Throwing Rock			
	Overall Effect	Additional Effect on Experienced Players	Additional Effect when ≥100 Games
	(1)	(2)	(3)
Opp's Frac Paper (first)	-0.0415***	-0.0695***	-0.0955***
	(0.0022)	(0.0056)	(0.0087)
Opp's Frac Scissors (first)	0.2531***	0.1409***	0.1379***
	(0.0023)	(0.0061)	(0.0094)
Opp's Frac Paper (all)	0.0018	0.0224*	0.0041
	(0.0033)	(0.0088)	(0.0138)
Opp's Frac Scissors (all)	0.0401***	0.0245**	-0.0209
	(0.0035)	(0.0094)	(0.0146)
Opp's Paper Lag	0.0042***	-0.0012	-0.0040*
	(0.0007)	(0.0016)	(0.0019)
Opp's Scissors Lag	0.0145***	0.0048**	-0.0014
	(0.0008)	(0.0016)	(0.0020)
Own Paper Lag	0.0001	0.0073***	0.0052**
	(0.0007)	(0.00015)	(0.0019)
Own Scissors Lag	0.0039***	0.0029	-0.0046*
	(0.0007)	(0.00015)	(0.0019)
Constant	0.3385***	-0.0106***	-0.0015
	(0.0007)	(0.0009)	(0.0018)
R^2	0.0142	0.0239	0.0258
N	4210005		

Note: All columns show OLS coefficients for the effect on whether a throw is rock; all the coefficients come from one regression. The first column is the effect for all players, the second column is the additional effect of the covariates for players in the restricted sample; the third column is the additional effect for those players after their first 99 games. Opp's Fraction Paper refers to the fraction of the opponent's previous first throws that were paper. Opp's Fraction Paper (all) refers to the fraction of all the opponent's previous throws that were paper. Opp's Paper Lag is a dummy for whether the opponent's most recent first throw in a match was paper. Own Paper Lag is a dummy for whether the player's own most recent first throw in a match was paper (similarly for scissors). The regression also control for the opponent's number of previous matches.

Table 4: Win Percentages

Data	Wins (%)	Draws (%)	Losses (%)	W - L (%)
Full Sample	34.14	32.38	33.48	0.66
Experienced Sample	34.65	32.18	33.16	1.49
Best Response to Predicted	42.14	29.7	25.28	16.86

Note: Experienced refers to players who play at least 100 games. "Best Response" is how a player playing against the experienced sample would do if she always played the best response to how the player is predicted to play by Table 3. We used half the sample to calculate the coefficients, which were used to predict the play for the other half. W-L is the difference between the first and third column; it equals the average winnings per throw if players bet $100 on a throw.

Table 5: Opponents' Response to Expected Payoff of Rock

Opponent's Expected	Opponent's Throw (%)			N
Payoff of Rock	Paper	Rock	Scissors	
$[-1, -0.666]$	29.5	42.57	27.93	3183
$[-0.666, -0.333]$	29.57	41.34	29.09	14723
$[-0.333, 0]$	32.65	33.81	33.54	354952
$[0, 0.333]$	34.53	32.46	33	329709
$[0.333, 0.666]$	35.42	33.85	30.73	11441

Note: Using the players' predicted play from Table 3 (and similar for paper and scissors), we calculate the expected payoff to their opponents of playing rock. This table shows the distribution of opponents' play for different ranges of that expected payoff.

Table 6: Parameters of the Structural Model

Variable	**Definition**
$\hat{r}_0$	fraction of the time a player plays k_0 *and* chooses rock
$\hat{p}_0$	fraction of the time a player plays k_0 *and* chooses paper
$\hat{s}_0$	fraction of the time a player plays k_0 *and* chooses scissors
$\hat{k}_1$	fraction of the time a player plays k_1
$(\hat{k}_2)$	$1 - \hat{k}_1 - \hat{r}_0 - \hat{p}_0 - \hat{s}_0$ (not an independent parameter)

Note: $\hat{r}_0$ is not equal to the fraction of k_0 throws that are rock; that conditional probability is given by $\frac{\hat{r}_0}{\hat{r}_0+\hat{p}_0+\hat{s}_0}$.

Table 7: Summary of k_0, k_1, and k_2 estimates.

Variable	**Mean**	**SD**	**Median**	**Min**	**Max**	**N**
k_0	73.79	15.7	75.2	18.65	100	6635
k_1	18.52	14.48	16.12	0	76.66	6635
k_2	7.69	7.76	5.87	0	40.65	6635

Note: Based on the 6635 players who have 100 clean matches where the k_1 and k_2 strategies are well-defined.

Table 8: Percent of players we reject always or never playing a strategy.

Variable	**N**	**95% CI does not include 0**	**95% CI does not include 1**	**95% CI does not include 0 or 1**
k_0	6399	93.06	58.26	57.45
k_1	6399	62.78	99.97	62.78
k_2	6399	11.55	95.55	11.55

Note: All percentages refer to the 6399 players who have 100 matches where the k_1 and k_2 strategies are well-defined and for whom we can calculate standard errors on the estimates.

Table 9: Summary of ch_0, ch_1, and ch_2 estimates.

Variable	**Mean**	**SD**	**Median**	**Min**	**Max**	**N**
ch_0	74.85	16.03	76.84	16.99	100	6853
ch_1	16.31	14.08	13.78	0	76.66	6853
ch_2	8.85	6.89	8.12	0	50.61	6853

Note: Based on the 6853 players who have 100 clean matches where the ch_1 and ch_2 strategies are well-defined.

Table 10: Summary of Naive k_0, k_1, and k_2 estimates.

Variable	**Mean**	**SD**	**Median**	**Min**	**Max**	**N**
k_0	72.25	17.73	74.72	6.54	100	5732
k_1	21.06	16.71	18.04	0	93.46	5732
k_2	6.69	7.71	4.39	0	45.75	5732

Note: Based on the 5732 players who have 100 clean matches where the naive k_1 and k_2 strategies are well-defined.

Table 11: Effect of Expected k_1 Payoff on k_1-Consistency (OLS)

	(1)	(2)	(3)
K1 Payoff	0.076***	0.107***	0.071***
	(0.0012)	(0.0014)	(0.0017)
High Opp Exp			-0.076***
			(0.0017)
Medium Opp Exp			-0.056***
			(0.0015)
K1 Payoff X High Opp Exp			0.541***
			(0.011)
K1 Payoff X Medium Opp Exp			0.324***
			(0.0051)
Experienced	0.014***	.	.
	(0.0016)	.	.
Exp X K1 Payoff	0.061***	0.043***	-0.005
	(0.0039)	(0.0039)	(0.0051)
Exp X High Opp Exp			-0.007**
			(0.0038)
Exp X Medium Opp Exp			-0.012***
			(0.0037)
Exp X K1 Payoff X High Opp Exp			0.072***
			(0.022)
Exp X K1 Payoff X Medium Opp Exp			0.075***
			(0.012)
Own Games>100	-0.002	0.002	0.031***
	(0.0022)	(0.0017)	(0.0044)
Own Games>100 X K1 Payoff	0.081***	0.066***	-0.026***
	(0.0068)	(0.0059)	(0.0080)
Own Games>100 X High Opp Exp			-0.017***
			(0.0049)
Own Games>100 X Medium Opp Exp			-0.010*
			(0.0052)
Own Games>100 X K1 Payoff X High Opp Exp			0.005
			(0.024)
Own Games>100 X K1 Payoff X Medium Opp Exp			0.065***
			(0.017)
Observations	3921798	3921798	3921798
Player Fixed Effects	No	Yes	Yes
Adjusted R^2	0.003	0.004	0.008

* $p < .10$, ** $p < .05$, *** $p < .01$

S.E.'s are clustered by player.

Note: 'k_1 Payoff' is the expected payoff to playing k_1 if the opponent randomizes according to his history (ranges from -1 to 1). 'High opp exp' is a dummy for opponents who have 95 or more past games; 'Medium opp exp' is a dummy for opponents with between 31 and 94 past games. 'Experienced' is a dummy for players who eventually play at ≥ 100 games. 'Own Games >100' indicates the player has already played at least 100 games. The 'X' indicates the interaction between the dummies and other covariates. SEs are clustered by player.

Appendix

We formalize the process of playing rock-papers-scissors over Facebook as a sequence of best-of-three matches described by the game Γ nested inside a larger game $\hat{\Gamma}$, which includes the matching process that pairs the players. We do not specify the matching process, as it turns out that it does not matter and the following holds for any matching process. Players may exit the game (and exit may not be random) after any subgame, but not in the middle of one. All players have symmetric payoffs and discount factor δ across subgames.

Each nested game Γ is a "best-of-three" match of rock-paper-scissors played in rounds, which we will call "throws." For each throw, both players simultaneously choose actions from $A = \{r, p, s\}$ and the outcome for each player is a win, loss, or tie; r beats s, s beats p, and p beats r. A player wins Γ by winning two throws. The winner of Γ receives a pay-off of 1 and the loser gets -1. Note that Γ is zero-sum. Therefore, at any stage of $\hat{\Gamma}$ the sum across players of all future discounted payoffs is zero.

Each match consists of at least two throws. Because of the possibility of ties, there is no limit on the length of a match. Let

$$\mathcal{K}^l = \underbrace{A^2 \times A^2 \ldots \times A^2}_{l \text{ times}}$$

be the set of all possible sequences of l throws by two players. Let $K^l \subset \mathcal{K}^l$ be the set of possible *complete matches* of length l: sequences of throw pairs such that no player had 2 wins after $l-1$ throws, but a player had two wins after the l^{th} throw. Let $K = \cup K^l$ be the set of possible complete matches of any length. Let $\hat{K}^l \subset \mathcal{K}^l$ be the set of possible *incomplete matches* of length l: sequences of throw pairs such that no player has 2 wins.

A player's overall history after playing t matches is the sequence of match histories for all matches he has played,

$$h_i^t \in H^t = \underbrace{K \times K \ldots \times K}_{\text{t times}}.$$

Players may not observe their opponents' exact histories. Instead a player observes some public summary information of his opponent's history. Let $f : H^t \to S^t \; \forall t$ be the function that maps histories into summary information. Denote by s^t an element of S^t.

A strategy for a player having played t complete matches, facing an opponent with m complete matches after l throws of the current match is a mapping from the player's history, the information the player has about his opponent's history, and partial history of the current match to a distribution of actions $\sigma_{t,m,l} : H^t \times S^r \times \hat{K}^l \to \Delta A$.

It is helpful to define a function $\#win_i : \hat{K}^l \cup K^l \to \{0,1,2\} \; \forall \; l$, which denotes the number of wins for player i after a match history. Similarly $\#win_j$ is the number of wins for player j and $\#win = max\{\#win_i, \#win_j\}$ is the number of wins of the player with the most wins.

In Nash Equilibrium, at any throw of any match, the distribution chosen must maximize the expected payoff. The payoff consists of the flow payoff plus the continuation value from future matches if the match ends on this throw, or the expected payoff from continuing the match with the updated match history if the match does not end on this throw. So,

$$\sigma_{t,m,l}(\hat{k}^l_{ij}, h_i^{t_i}, s_j^{t_j}) \in \arg\max_{\sigma_i} \mathbb{E}_{\sigma_i,\sigma_j}\left[u(\hat{k}^l_{ij}, a_i, a_j)\right],$$

where

$$u(\hat{k}^l_{ij}, a_i, a_j) = \begin{cases} 1 + \delta \cdot \eta((h_i^{t_i}, (\hat{k}^l_{ij}, (a_i, a_j)))) & \#win_i((\hat{k}^l_{ij}, (a_i, a_j))) = 2, \\ -1 + \delta \cdot \eta((h_i^{t_i}, (\hat{k}^l_{ij}, (a_i, a_j)))) & \#win_j((\hat{k}^l_{ij}, (a_i, a_j))) = 2, \\ v_i((\hat{k}^l_{ij}, (a_i, a_j)), h_i^{t_i}, s_j^{t_j}) & \#win((\hat{k}^l_{ij}, (a_i, a_j))) < 2. \end{cases}$$

In the first two cases there is the immediate payoff from the match ending plus the inter-match continuation value, $\eta(h_i^{t_i+1})$ (the value of going back into the pool and playing future matches with the updated player history). In the last case there is just the intra-match continuation value, $v(k^{l+1}, h_i^{t_i}, s_j^{t_j})$ (the value of continuing this match with the updated match history), which includes the later inter-match continuation value from when the match eventually ends. Note that we have implicitly set the intra-match discount factor to 1.[26]

Because each match is symmetric and zero sum, the inter-match continuation values are unimportant.

Lemma 1. *Under Nash equilibrium play, $\eta(h^t) = 0 \; \forall h^t \in H^t \; \forall t$.*

Proof. Suppose that $\eta(h_i^t) \neq 0$. If $\eta(h_i^t) < 0$, then a player with that history could deviate to always playing $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. The player would win half their matches and lose half their matches and have continuation value 0. If $\eta(h_i^t) > 0$, then since $\hat{\Gamma}$ is zero sum at every stage, there must exist h_j^r such that $\eta(h_j^r) < 0$, which by the same logic cannot happen in equilibrium. $\square$

[26]Because what matters for the result is the symmetry across strategies at all stages, having an intra-match discount factor does not change the result, but substantially complicates the proof.

Since inter-match continuation values are all zero, the intra-match continuation value is just the probability of eventually winning the match minus the probability of eventually losing the match. This means that continuation values are zero-sum: $v_i(\hat{k}_{ij}, h_i, s_j) = -v_j(\hat{k}_{ij}, h_i, s_j)$.

The symmetry of the match also implies that, regardless of history, whenever players are tied for the number of wins in the match, the continuation value going forward is 0 for both players.

Lemma 2. *In Nash equilibrium if $\#win_i(\hat{k}^l_{ij}) = \#win_j(\hat{k}^l_{ij})$, then $v_i(\hat{k}^l_{ij}, h_i, s_j) = v_j(\hat{k}^l_{ij}, h_j, s_i) = 0$.*

Proof. Assume $\#win_i(\hat{k}^l_{ij}) = \#win_j(\hat{k}^l_{ij})$. Suppose $v_i(\hat{k}_{ij}, h_i, s_j) \neq 0$. If $v_i(\hat{k}_{ij}, h_i, s_j) < 0$, player i has a profitable deviation to play $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for the remainder of the match and they will win with probability one-half, giving them a continuation value of zero. Similarly for j. We therefore have $v_i(\hat{k}_{ij}, h_i, s_j) \geq 0$ and $v_j(\hat{k}^l_{ij}, h_j, s_i) \geq 0$ for all match histories.

If $v_i(\hat{k}^l_{ij}, h_i, s_j) > 0$ then since the match is zero-sum, there must exist an $\hat{h}_i, \hat{h}_j$ such that $v_i(\hat{k}^l_{ij}, \hat{h}_i, f(\hat{h}_j)) < 0$ or $v_j(\hat{k}^l_{ij}, \hat{h}_j, f(\hat{h}_i)) < 0$, which contradicts the above result. Similarly for $v_j > 0$. □

Lemma 3. *In Nash equilibrium, if $\#win_i(\hat{k}^l_{ij}) = 1$ and $\#win_j(\hat{k}^l_{ij}) = 0$ then $v_i(\hat{k}^l_{ij}, h_i, s_j) = \frac{1}{2}$ and $v_j(\hat{k}^l_{ij}, h_j, s_i) = -\frac{1}{2}$.*

Proof. Assume $\#win_i(\hat{k}^l_{ij}) = 1$ and $\#win_j(\hat{k}^l_{ij}) = 0$. Suppose $v_i(\hat{k}^l_{ij}, h_i, s_j) < \frac{1}{2}$. This implies that player i has a profitable deviation to play $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for the remainder of the match. If they do so they will win the match with probability

$$\underbrace{\frac{1}{3}}_{\text{win this throw}} \cdot 1 + \underbrace{\frac{1}{3}}_{\text{lose this throw}} \cdot \frac{1}{2} + \underbrace{\frac{1}{3}}_{\text{draw this throw}} \left(\frac{1}{3} + \frac{1}{3} \cdot \frac{1}{2} + \frac{1}{3}(\cdots) \right) = \frac{1}{3} \cdot \frac{3}{2} \left(1 + \frac{1}{3} \cdots \right) = \frac{3}{4}.$$

and lose with probability $\frac{1}{4}$, giving them an intra-match continuation value of $\frac{1}{2}$. Therefore $v_i(\hat{k}^l_{ij}, h_i, s_j) \geq \frac{1}{2}$. Similar logic guarantees that $v_j(\hat{k}^l_{ij}, h_j, s_i) \geq -\frac{1}{2}$.

If $v_i(\hat{k}^l_{ij}, h_i, s_j) > \frac{1}{2}$ then since the match is zero-sum, there must exist an $\hat{h}_i, \hat{h}_j$ such that $v_i(\hat{k}^l_{ij}, \hat{h}_i, f(\hat{h}_j)) < \frac{1}{2}$ or $v_j(\hat{k}^l_{ij}, \hat{h}_j, f(\hat{h}_i)) < -\frac{1}{2}$, which contradicts the above result. Similarly for $v_j(\hat{k}^l_{ij}, h_j, s_i) > -\frac{1}{2}$. □

Since the number of wins each player has after any history that has not ended the match is either 0 or 1, Lemmas 2 and 3 encompass all the possible

$\hat{k}^l_{i,j}$. We therefore see that the continuation value depends only on the number of wins each player has

$$v_i(\hat{k}^l_{i,j}, h_i, s_j) = \tilde{v}_i(\#win_i(\hat{k}^l_{i,j}), \#win_j(\hat{k}^l_{i,j})).$$

So, for every $\hat{k}^l_{i,j}$ we can calculate $u_i((\hat{k}^l_{i,j}, (a_i, a_j)))$ for each own throw and opponent throw. Figure 8 gives these payoffs for each possible stage in a match. For each of these stages, it can only be a Nash Equilibrium of each player expects their opponent to play $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

A\B	R	P	S
R	0,0	$-\frac{1}{2},\frac{1}{2}$	$\frac{1}{2},-\frac{1}{2}$
P	$\frac{1}{2},-\frac{1}{2}$	0,0	$-\frac{1}{2},\frac{1}{2}$
S	$-\frac{1}{2},\frac{1}{2}$	$\frac{1}{2},-\frac{1}{2}$	0,0

(a) Start of Match

A\B	R	P	S
R	$\frac{1}{2},-\frac{1}{2}$	0,0	1,-1
P	1,-1	$\frac{1}{2},-\frac{1}{2}$	0,0
S	0,0	1,-1	$\frac{1}{2},-\frac{1}{2}$

(b) A has 1 win, B has zero

A\B	R	P	S
R	$-\frac{1}{2},\frac{1}{2}$	-1,1	0,0
P	0,0	$-\frac{1}{2},\frac{1}{2}$	-1,1
S	-1,1	0,0	$-\frac{1}{2},\frac{1}{2}$

(c) B has 1 win, A has zero

A\B	R	P	S
R	0,0	-1,1	1,-1
P	1,-1	0,0	-1,1
S	-1,1	1,-1	0,0

(d) A has 1 win, B has 1 win

Figure 8: Total payoffs (flow + continuation) for each stage of the match.