Conflict Begets Cooperation in Socialized KingPawn

version 13/04/07

J. Stephen Judd
Computer and Information Science
University of Pennsylvania

Abstract

We report experiments on a two-person game in which human subjects continuously choose to play either a dominant role (called a *King*) or a submissive one (called a *Pawn*). A King receives a higher payoff rate, but only if the partner is a Pawn. Straightforward equilibrium analysis predicts there will be no conflict (time when both players are Kings), but the behavioral data is utterly at odds with that, as is intuition.

We decompose the game into parts and offer an analytic solution for an important subgame. This yields a much more satisfying prediction of conflict behavior, and indeed the recorded games are extremely well correlated with the prediction. Both the theory and the data show structure in conflict that unequivocally encourages cooperative sharing behavior.

Introduction

It has often been declared that people are not "rational" when they play some interesting games. Conventional game-theoretic equilibrium analysis predicts one outcome, but people resolutely behave in some other fashion, so analysts throw up their hands and expel them from the ranks of homo economicus. In the KingPawn game we study here, a naive game theoretic analysis indeed predicts one type of behavior and our behavioral data utterly conflicts with it. To resolve this, we introduce a novel theoretical analysis of the game that makes more believable predictions and also matches the behavioral data much better. It also leaves the notion of homo economicus quite intact.

The experiments involved 172 human subjects from far-flung countries, who were repeatedly assigned to play each other in randomly chosen pairs. During a game, each subject could choose to be a *King* or a *Pawn*. A King is paid at a higher rate (twice as much as a Pawn), but only if the partner has not also chosen to be a King; a King in conflict receives no payments. The naive game-theoretic viewpoint suggests that no conflicts should occur, but the behavioral results firmly suggest otherwise. Since only one of the two players can be a King for either to be paid, such a configuration is inherently conflicting, giving rise to considerable tensions that sometimes provoke expensive punishments.

Other reports discuss the emotional (Fehr and Gächter 2000) or neural (Sanfey et al. 2011) or genetic (Cesarini et al. 2008) basis for preferring fairness during decision-making, but we prefer to believe that the sharing phenomena we have found are explained in simpler strategic terms. We find that people understand the dynamic of this game and a sufficient number of them practice punishing behaviors that result in a global expected payoff that is monotonically increasing with the degree of cooperativity and fairness that a player exhibits. This results in a meta-game with a stable cooperative equilibrium that pays near maximum social welfare and pays both players nearly the same. Fairness develops because the average group behavior makes it statistically rational in the meta-game, even though it is not rational in the prima facie game. This remarkable phenomenon is achieved without any communication between the punishers. Because of this lack of communication, it *appears* that such a system could only arise after a long enculturation period in which players play or observe similar games, observing both actions and payoffs often enough to be able to predict statistics of the social cohort. But this begs the question. It does not win anything to explain away a remarkably cooperative coordination in a game of conflict by assuming a long and remarkably cooperative development of strategies to play it. The long and far-sighted effort it would take to study the problem, learn the solution, and broadly apply the lessons are too unlikely to assume they would "just happen". Once again, we might go looking for some emotional or neural or genetic substrate that could sustain this undertaking.

However, we also offer a theoretical analysis that provides a straightforward game-theoretic explanation for the structure of the meta-game. It decomposes KingPawn and identifies a subsumed game called Will-Testing. Unlike the naive analysis, the equilibrium of Will-Testing *does* involve conflict, and the quantity it predicts is extremely well correlated with the behavioral data. Any amount of understanding of Will-Testing – be it intuitive, instinctual, intellectual or experiential – would certainly help players deploy appropriate responses to KingPawn. Since Will-Testing is very general it may underlie many other games, and hence there may be wide-spread experience with it among humans. This may be why human subjects are so quick to play KingPawn in a sophisticated manner.

The experiments described here are part of a broader and ongoing program of behavioral experiments in strategic and economic interaction on social networks conducted at Penn (Kearns, Suri, and Montfort 2006; Judd and Kearns 2008; Kearns et al. 2009), and are an effort to apply the methods of behavioral game theory (Camerer 2003) to the study of social networks.

Related Literature

The subject of fairness in human interactions has a very long history. Sociologists and social psychologists view it as central to many social phenomena, and have well-developed theories of fair exchange and reciprocity (exchange/equity theory) (Brown 1986), although a true appreciation of equity of exchanges needs to consider long term "accounting" and the subjective evaluation of fairness.

There is considerable economic literature on fairness. Rabin (1993) offers a theory that incorporates fairness into traditional game theoretic models of pairwise interactions. The economic experiments of Fehr and Gächter (2000) show that people frequently punish non-altruistic behavior and derive pleasure from doing so. Our results also document punishment, but we interpret them in strictly game-theoretic terms, and feel the "pleasure" element is only an interpretive side story.

Bowles and Gintis (Bowles and Gintis 2003) discuss explanations of cooperation ranging from kin selection and genetics through to specifically-human cognitive and linguistic capabilities. Their view of the human ability to internalize norms may be the closest to what we witness in this paper.

KingPawn is related to the well-known Ultimatum Game.

One-Shot Version

The one-shot version of this game is given by the following game payoff matrix:

	Pawn	King
Pawn	(1,1)	(1,2)
King	(2,1)	(0,0)

When one player plays a mixed strategy with equal probability for each role, the other player will get an expected payoff of 1 regardless of strategy employed. The equilibrium is for both players to use this mixed strategy, and consequently both players earn the pawn rate. An echo of this result will occur in the analysis of continuous-time KingPawn.

Theoretical Analysis

The Continuous KingPawn Game

The continuous-time game has the extra element of time to deal with, and a new artifice will be used to deal with it.

There are two pure equilibrium states, and each one has exactly one King. A stasis appears unavoidable, because the only player that dislikes the state has no option except to enter an even worse state (one involving 2 kings, which pays nothing). Furthermore, there is no clock that assists in the synchronizing of state changes; this makes the element of time hard to reason about.

The crux of equilibrium reasoning, though, is simply that the players will be in *an* equilibrium. It does not much care *which* equilibrium, and it is also unconcerned about the route to get into an equilibrium, or even transitions between them. If there were something extraneous that caused players to jump from one equilibrium to another, the players would still be expected to stay in the (new) equilibrium. Therefore if there were some mechanism working through time that caused such transitions to occur, the only game theoretic prediction about the outcomes of KingPawn is that no time would be spent in a non-equilibrium state. This means that the max social welfare is always obtained, the final ratio of earnings could be anything from .5 to 1, and the two players will fall somewhere on the dashed lines of Figure 1. Call this the naive analysis of KingPawn.

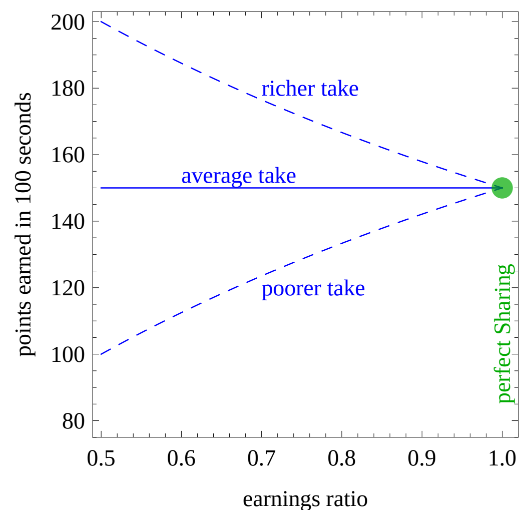

Figure 1: Naive payoffs for rational players in continuous-time KingPawn. The two players must fall on the dashed lines, but their average is always on the solid line.

We will enrich this analysis by breaking the game into smaller components.

Decomposing KingPawn

The play of a KingPawn game can be decomposed in various ways. The way we choose to do it is into the following three separate decisions: 1) deciding the ratio of earnings, 2) deciding which player gets the larger share, and 3) deciding how much conflict there is. Mathematically, there is a fourth element which specifies how much time is spent in the Pawn-Pawn state, but since this is neither rational in theory or common in practice, we deem it to be zero. By doing so, the answers to the three decisions fully determine an outcome of KingPawn, and vice versa.

We will avoid Decision 1 and carry it along as a parameter during analysis of the latter two. The question about who will be the richer one seems to be the core contest in KingPawn, and some conflict is often employed to resolve it. It is to that interesting negotiation that we now turn.

The Will-Testing Game

The Will-Testing game gives the larger payoff to whichever player is more tolerant of conflict. It is played by two people in continuous time through a fixed time interval, T. Either player may capitulate at any time, thereby ending the play. The payoff for the victor is some payscale V times the amount of time remaining in the interval; the capitulator gets half as much as the victor[1].

The tension in the game is that both players want someone to capitulate in a hurry, but both want the other one to do so, and hence they both stall. There is no equilibrium to this game. The maximum social welfare is $(V + \frac{1}{2}V)T$ which occurs only when one player capitulates immediately.

The tension changes as we alter the fraction that the capitulator gets. If his fraction is very small (say 1/100) then there is a huge incentive to stall. If his fraction is close to 1 then it does not matter very much who goes first, and it's relatively easy to make the sacrifice to capitulate. Since the fraction thus makes a big difference, we will parameterize the game and call the capitulator's fraction ρ where $0 \leq \rho \leq 1$.

In spite of its description as being played in continuous time, we will formalize it as a one-shot game where each player i selects a real number s_i in the fixed interval $[0, T]$. Then the payoffs (before multiplication by V) for players 1 and 2 respectively are:

$$\begin{cases} (T - s_2) & \text{and} \quad \rho \times (T - s_2) \quad \text{if } s_1 > s_2 \\ \rho \times (T - s_1) & \text{and} \quad \rho \times (T - s_2) \quad \text{if } s_1 = s_2 \\ \rho \times (T - s_1) & \text{and} \quad (T - s_1) \quad \text{if } s_1 < s_2 \end{cases}$$

There is no pure equilibrium strategy for this game. There is a mixed strategy equilibrium, which we reveal below, but instead of thinking of this as a 2-player game (which is a poor model of our experimental setup), we prefer to recast it as a game played in a large society of anonymous players.

The *social version* of the game has a large set of players who each choose an action once; their payoffs are the average of the payoffs using that choice against all other players. This captures the idea of a group of social animals that play the game repeatedly with partners picked uniformly at random from the same group. The randomization of strategy that would occur in a repeated 2-person version of KingPawn is replaced in the social version by random selection of 2 players who each have pure strategies, and the expectation over choices is replaced by the average over all possible pairings. Thus the two views are equivalent.

There is no pure equilibrium for this game when there are a finite number of players. In the limit of infinite players there is a density function for the equilibria, but it does not specify which player uses which strategy so all permutations are equivalent. Finding this distribution involves an equation that makes the usual assumption that all players get equal payoffs, so the players are quite indifferent to which permutation occurs.

One strong interesting feature of the equilibria is that all players must choose *different* actions. Birds and mammals and most other animals are generally perceived as being individuals who are all slightly different, populating a spectrum of behavior space. They are also widely seen as having social structures like dominance hierarchies wherein pairwise relations between individuals can be characterized. We view the social version of Will-Testing to be an example of a game that produces a similar ordered structure where the existence of differences in individual behaviors is a straightforward requirement of the game's solution. In the case of this particular game, the social dynamic will be interpretable as a total ordering (like a pecking order) simply because the action space is real and one dimensional.

[1] Will-Testing is very related to War of Attrition, but there is a major difference in that for both players the value of winning the contest goes to zero at the end of the game. In War of Attrition, the value of the prize is a constant, which gives a perpetual incentive to bid higher even after the price is more than the value of the prize. The equilibria are thus quite different.

Equilibria of Social Will-Testing

We seek a density over the action space, $\text{den}(t) \geq 0$ and $\int_0^T \text{den}(t)dt = 1$. To find the average over all possible pairwise interactions in the social version of the game, we must integrate over that space. The average pay of action s is thus:

$$\text{pay}_\rho(s) = V \int_0^s \text{den}(t)(T-t)\,dt + V \int_s^T \rho\,\text{den}(t)(T-s)\,dt \tag{1}$$

Now make the usual assertion about the nature of the equilibrium that no action is more profitable than any other (otherwise some player would choose that action rather than the one she did), That is, set pay to be constant:

$$\begin{aligned} \frac{\text{pay}_\rho(s)}{V} = c \quad &= \quad \frac{\text{pay}_\rho(0)}{V} = 0 + \int_0^T \rho\,\text{den}(t)(T-0)\,dt \\ &= \quad \rho T \int_0^T \text{den}(t)\,dt = \rho T \end{aligned}$$

The pay per time spent is thus just ρ for every player, the same as what the capitulator gets if he acts immediately.

To make the Will-Testing game have the same Max Social Welfare as KingPawn (3 points/second), the pay scale V must be adjusted for each earnings ratio: $V(\rho) = 3/(1+\rho)$, as plotted by the dash-dot line in Figure 2. The average pay of both players in any game is $(R + \rho R)/2 = \rho T V(\rho)$, where R is the richer player's take. The richer player earns $R = 2\rho\ 100 \times 3/(1+\rho)^2$, which is shown as the dashed line in the same plot. The poorer take is ρR, also shown as a dashed line.

Experimental Design

Game Design

A GUI for the game was implemented as a Java applet and deployed as an internet service. Via a mouse click, each player could instantly change to be a King or a Pawn at any time. As a King, the player enjoyed a high pay rate (2 points per second), but payments only accrued when his partner was not a King. The Pawn earned a steady income, albeit only half of a King's (1 point per second). Social welfare was maximized when one player was King and one was Pawn; together they earned 3 points/second.

Each player was given a view of both his own and his partner's state. Payments accrued continuously and were

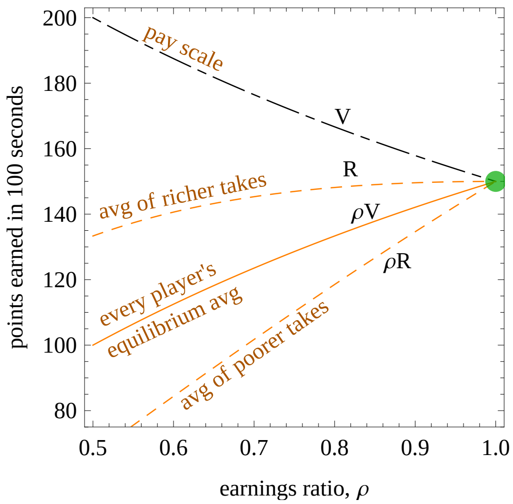

Figure 2: Equilibrium payoffs for Will-Testing, as a function of the earnings ratio. Every player's average take falls on the solid line. Ignoring individuals and averaging over games, the richer and poorer takes fall on the dotted lines.

displayed for each player in real time, along with a clock that measured out the 100 second duration of the game. A screenshot of the GUI is shown in Figure 3.

This continuous-time game is easily recognized as an iterated version of the one-shot game. Our equipment had a clock with 1-millisecond resolution, and hence it was actually a finitely repeated game but since this granularity is far smaller than humans can detect, it is better described as continuous time.

Human Subject Methodology

All experiments were held in a single session lasting 4.5 hours with workers from Amazon Mechanical Turk as participants. The players came from multiple far-flung time zones. They had been recruited and trained on earlier days and then told of a collective game-playing time and website where they could all congregate to be matched with other players. We paid US $0.003 per point accumulated.

The players were repeatedly assigned to play each other in randomly chosen pairs, according to the following procedure. The player with the fewest number of games so far was selected to play so long as s/he also met these constraints: 1) was not busy playing someone else, 2) had played fewer than 15 games so far, 3) had declared themself ready for another game, and 4) had internet connections that were currently fast enough. Such a player was matched randomly with any other player who also fit the above constraints and who had not played the first one selected. The server then initiated a game process for them. Both players had to respond to an alert, locate a popped-up browser window, bring it to the foreground, and privately select a starting role for himself before the game was allowed to begin. After it started, all actions were visible to both players and were asynchronous.

Players were given no indication of who they were playing, and since they were randomly assigned to play with

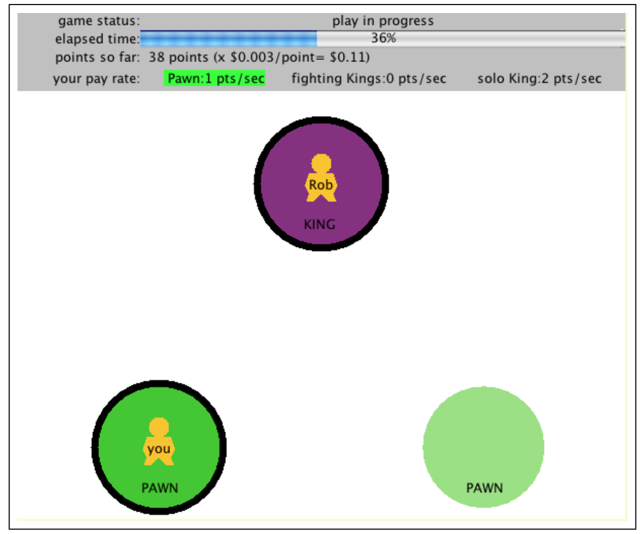

Figure 3: A screenshot of a player's GUI showing the user in Pawn mode and his partner in King mode.

anyone else in the widely and thinly spread online cohort, there was negligible chance that any outside communication occurred between subjects.

The game session was advertised as lasting for at least an hour and was forecast to have scores of players attending. People understood that there would be no re-encounters with previous partners, so there was neither historical nor future interactions to consider.

Experimental Results

Attendance

Of the 201 people who arrived at the site during the session, 172 played at least one game. Figure 4 gives a histogram of how many games they played. They were limited to 15

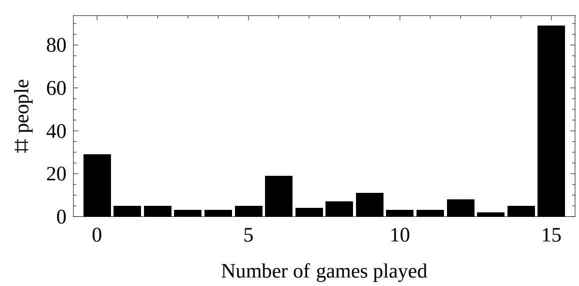

Figure 4: Histogram of number of games played

games apiece, but many people played fewer either because they arrived late or left early. Some people were prevented from playing due to technical problems like their internet connection was too slow.

A total of 974 games got played. Most players stayed at the site for more than an hour, and on average they played about 5.7 games during the session, each with a different partner.

Social Welfare

Over all games, the average earning rate was 1.21 points/second. This is above the pawn rate, but below the max social welfare rate of 1.5 points/player/second. Hence the players as a group were 80.6% efficient at extracting money from the experimenter's pocket.

Distribution of Takes

Figure 5 plots the number of points earned in each encounter (being twice as many as there are games). Of these, 21% were less than the pawn rate, and 8% were greater than the perfect sharing rate, leaving about 71% in between.

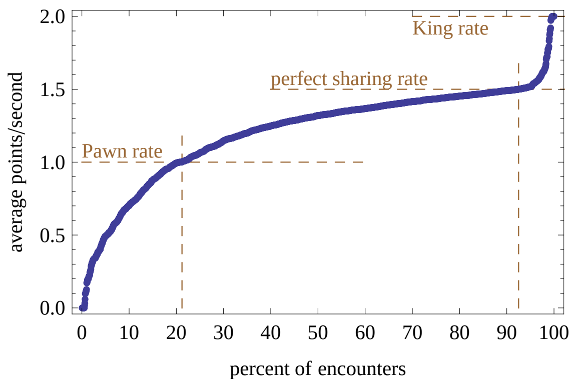

Figure 5: Cumulative distribution of individual takes (shown with the customary axes interchanged.)

Ratios of Takes

Figure 6 plots the outcomes of all games where the two axes show the earnings of the two players involved.

Figure 7 is a smoothed histogram of the number of games with a given earnings ratio[2]. It is distinctly bimodal, with a lump near 0.5 caused by the existence of people we call Demanders (those who insist on being King all or most of the time), and a bigger lump near 1.0 populated by Sharers (those who are willing to trade the King role back and forth with their partner).

Predictive Power

Simple game-theoretic analysis predicts no conflicts will persist: it is "irrational" to play King when your partner is a King and thereby earn 0 if you could play Pawn instead and earn 1 point per second. Also, it is irrational to play Pawn if your partner is a Pawn, since a lone King's take is higher than a Pawn's. Hence every game will earn the maximum social welfare and the average take between the two players in each game is exactly 150 points, irrespective of earnings ratio. The solid blue line in Figure 8 represents this prediction.

Nevertheless, the observed games show abundant conflict; Figure 8 reveals how much is sustained and where it occurs. Every black dot represents the mean of the takes of

[2]There were 5 games where both players earned 0, and since their earnings ratio is not defined they are not included in the plot.

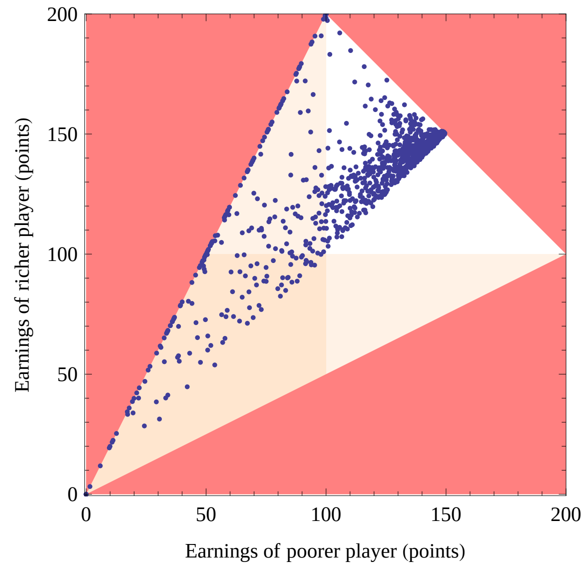

Figure 6: Earnings in each game. The outcome of each game is plotted with the richer player arbitrarily shown on the vertical axis. The pink zones are not reachable; zones below the pawn rate are shaded.

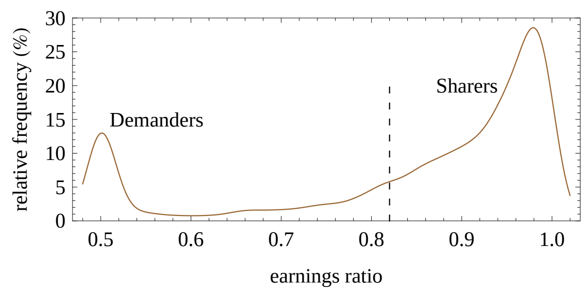

Figure 7: Smoothed histogram showing bimodal distribution of earnings ratios.

all players whose games had a given earnings ratio[3]. The solid orange line are the equilibria of social Will-Testing, also shown in Figure 2. It is clear this solid orange line is a much better approximator of the data points than the solid blue line. The blue line has zero predictive value for the data, but the Pearson correlation coefficient between the orange predictor and the data is 0.974. The probability of that correlation being a chance occurrence is $< 10^{-6}$.

Cost of Exploration

A distributed system can converge to some target distribution if every player can detect anomalies in the current distribution and move their own choice of s to escape overly dense zones in favor of overly sparse ones. But of course to

[3]To avoid arbitrary decisions about the number and size of bins and the locations of their boundaries, the data points each represent a set of 70 games (7.2% of the total). Each set has been selected to be as narrow in its spread of earnings ratios as possible, and each set overlapped its neighbor but was different in 4 data points. This technique helps suppress noise, especially in the very sparse region of data with earnings ratios between 0.55 and 0.75. (See Figure 7.)

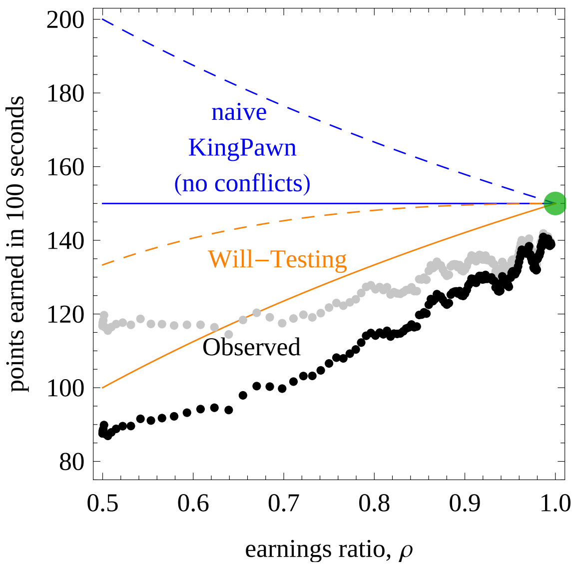

Figure 8: Take plotted against earnings ratio, comparing KingPawn, Will-Testing, and human behavior. The solid lines are averages over all players and should be compared with the black dots of observed human behavior. The dashed lines are averages of just the richer players and should be compared with the gray dots of observed human behavior.

do so they need to be able to sample the distribution. Furthermore, for this to work the distribution cannot be changing faster than it can be sampled. Hence, assume that every player starts with some preferred action, repeatedly samples the current distribution and slowly moves his value toward a sparse zone. The process must be capable of moving the values in either direction.

One quirk of Will-Testing is that the revealed data at the end of each game is not symmetric; the player with the bigger s value gets to see the other player's (smaller) s value, but the player with the lesser value learns only that the other player had some bigger value. The capitulator thus gets less information, and in particular he gets no information about the zone just higher than his current value of s. If it is sparse, then it will be beneficial to move there.

These observations leads us to hypothesize that players will engage in conflict beyond what the equilibrium predicts in order to discover that extra valuable information. As long as each player can gather information about a small zone just higher than his current preferred action, the group as a whole can correctly populate the space. The cost of performing this exploration will cause the average social welfare of Will-Testing to be less than the theory here suggests. A new cohort will not be in equilibrium right away, and until it is there will be excess conflict.

Indeed, the black dots of Figure 8 do have a systemic offset. The observed values are everywhere lower than what the Will-Testing equilibrium predicts. It would be useful to have a model to quantify this cost of exploration.

Punishment Profile

There is an easy way for any player to guarantee that he is no poorer than his partner, so it is tempting for someone to examine the payoffs exclusively from the point of view of the richer player. Given how attractive this mindset might become, it is worthwhile to look not just at the *average* take but to examine the special experience of the *richer* player, and see how that might influence his behavior. By separating the Demander's results from his victim's, we can draw an even clearer conclusion. Every gray dot in Figure 8 shows the mean of the takes of just the richer player in all games which had a given earnings ratio. Note that even the average *richer* take was everywhere less than the perfect Sharer's take of 150. This is an important strategic signal that will surely influence play.

In fact, human Sharers were not really making 150 points; their effective limit was about 140 points, caused by slight overlaps during role switches. Even so, the statistics are unequivocal: for every gray dot with a ratio less than 0.82 (i.e. for the entire space to the left of the dashes in Figure 7), the mean of that set of games passed a statistical test showing it to be less than 140, with $P < 0.01$ in every case.

Conclude therefore that the expected income of Demanders was thus convincingly less than that of Sharers. And conclude from this conclusion that cooperation is thus mathematically encouraged, and that all players should develop a bias toward outcomes of higher earnings ratios.

For every ratio of earnings, the dashed blue curve of Figure 8 is the amount that the richer player would take in naive KingPawn analysis; It is easy to grasp this and be seduced by its possibility into preferring outcomes with low earnings ratios. Most remarkably, this curve and the actual takes have distinctly opposing trends. The punishment (i.e. conflict) incurred for trying to get more than the Sharing level is proportional to and larger than the amount of temptation there is to try. The punishment profile is clearly encouraging the richer player to allow the poorer player to earn more.

Distribution of Conflict

The dots in Figure 8 are merely the averages of conflict. This section examines how conflict times are distributed, in order to see whether players are really using actions drawn according to the equilibria of Will-Testing.

A solution to equation (1) for $\mathrm{den}(\cdot)$ is available in (Gurvits and Judd 2012). It provides a closed form description of the equilibrium probability density function:

$$\mathrm{den}(\rho, s) = \frac{\rho(1-s)^{-\frac{2\rho-1}{\rho-1}}}{1-\rho}$$

Its cumulative distribution is

$$\mathrm{cdf}(\rho, s) = \int_0^s \mathrm{den}(\rho, x)\, dx = 1 - (1-s)^{\frac{\rho}{1-\rho}}$$

which is plotted in Figure 9 for various values of ρ.

The capitulator's action is the *smaller of two values*, and if those two values are independently drawn from den(), the cumulative distribution of the measured conflict will be

$$\mathrm{cdf}_2(\rho, s) = 1 - (1 - \mathrm{cdf}(\rho, s))^2 = 1 - (1-s)^{\frac{2\rho}{1-\rho}}$$

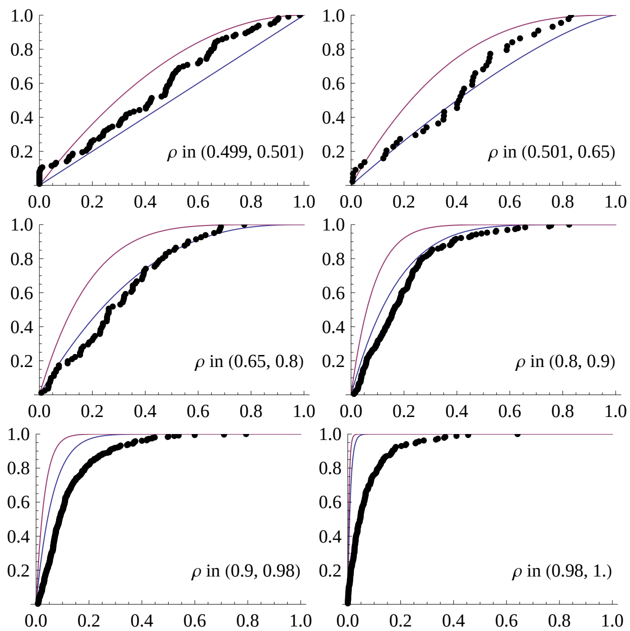

Figure 9: Cumulative distributions of conflict in various zones of earnings ratio, ρ. The horizontal axes represent s, the Will-Testing action choice of the capitulator, which is identical to the amount of conflict endured; the vertical axes are the fraction of games that see that amount of conflict or less. The dots represent data from all games collected with earnings ratios in the noted zone. The blue lines represent the equilibrium cdf in the center of the zone. The purple lines represent cdf_2.

which is also plotted in Figure 9.

The individual data points in the same figures each display the amount of conflict in one game. They reasonably approximate the predictions, but are generally further to the right (meaning that there is more conflict than predicted) which is to be expected due to the cost of exploration. In addition to that discrepancy, it is clear here that there is progressively wider divergence with higher earning ratios.

This progression is not fully understood, but it may be partially caused by overhead incurred during role switches. A friendly switchover is typically produced by the Pawn player deciding to change to King, followed by a reaction time delay of up to 2 seconds, followed by the other player switching to Pawn. The brief time of double Kings counts as "conflict" in our calculations, but would be viewed by the humans as merely a signal that it was time to change roles. The Sharing games tend to have more role switches than others. Demanders do not respond to such signals, so friendly switchovers are progressively rarer in lower earnings ratios.

Discussion

Source of Discipline

That the average Demander's take should be less than the Sharer's take is strategically desirable from the point of view of Sharers, who clearly have an interest in encouraging sharing behavior amongst everyone else. As a group they have jointly punished the Demanders and produced an expected payoff curve that tends to increase with earnings ratio.

But there is a deep puzzle as to how this group effect gets organized. Note that the Demander's take is indeed an *average*. It arises from the action of the set of their victims who each separately act to punish the Demander. But this set of victims *cannot observe each other*. They are unlikely even to have played each other, and none of them could ever directly measure the amount to which others in the set have punished the Demanders. So how could they mutually arrange for their *average* take to be so strategically propitious? How could a group of players who never met and never observed each other produce such an effective ensemble effect?

There is a second-order tension within Sharers between those who pay the cost of punishing Demanders and those who do not, but note again that in this game they can never even detect each other's existence!

The mathematics of Will-Testing is a candidate explanation of how this coordination all gets invisibly organized. Its equilibrium predicts conflict, and quantifies it; as a consequence, the flat average earnings curve predicted by naive analysis is transformed into a curve that monotonically increases with earnings ratio. Even more dramatically, the curve of richer-player earnings is transformed from monotonically decreasing to monotonically increasing – a profound alteration, produced solely by old fashioned equilibrium reasoning.

The most games any one person played was 15 but there were 172 people playing. It's clear, therefore, that there simply were not enough games played for people to be able to survey their entire community or even 10% of it. Hence most of the behavior on display here reflects traits that were acquired *outside* the experimental environment. There is no convergence toward a negotiated equilibrium here; there is no group adaptation to a new and curious little game; what we are observing is the distribution of strategies that already exist in the human wild for dealing with everyday games such as these.

Origin of Cooperation

Much of the literature on the origin of cooperation dwells on five separate models: 1) direct reciprocity, 2) indirect reciprocity, 3) spacial selection, 4) kin selection, and 5) group selection. We feel that none of these models captures what is going on in KingPawn. The observed distribution of punishments creates a meta-incentive for everyone to cooperate, but it is not a result of any of the above models. All those models require the existence of some sort of reputation or relationship between *specific* players, but such things do not exist in the anonymous setting imposed in our experiment.

The players do not know who they are playing. In fact the only detail they have about their partner is that they have *not* played with that person before. Thus, reputation is the one feature that is *explicitly removed* from the situation. It is replaced only with the knowledge that they are playing against randomly-chosen members of a large group. That group has at least these things in common: they all live in modern states, are computer savvy, are capable of performing tasks on Mechanical Turk, have some disposable leisure time, and can read English. All that is inferable here are generic features of the ensemble; nothing of individuals is available. Reputation and relationships between our players are truly vaporous.

We believe that Will-Testing is familiar in various forms to any adult human, and that we all have an understanding and perhaps a personality trait that helps us deal with its subtle equilibrium. It seems perfectly plausible that intelligent creatures with highly honed game theoretic skills could understand the potential for and the value of cooperation in KingPawn, could acquire behavioral tactics that encourage it, and deploy those behaviors in appropriate settings.

There is nothing in this study that speaks to whether the effect is embodied in genetic or cultural or learned or intellectual form.

Existence of Sharers

We offer no quantitative model of how Decision 1 gets made. We view it as a consequence of the entangled wills of two players. Those players are exposed to the marked slope of increasing payoff with increasing earnings ratio, and hence they should be responsive to it and know where to seek higher payoffs. Although this is only a qualitative statement, it clearly agrees with the preponderance of Sharers shown in Figure 7.

Existence of Demanders

A deeper puzzle is why the Demander behavior persists, given that it pays so poorly. If the meta-game has been so well developed and deployed, why would everyone not have become a Sharer already? One explanation is simply that Demanders are merely naive players who have not grasped the meta-game that is being imposed on them, or they have not yet collected enough long-term statistics to conclude the better value of Sharing.

Another explanation revolves around the newness of the assembled group, and it might occur even if no one is naive enough to have missed the logic of Sharing elsewhere. The logic of Demanding is certainly compelling too: if my partner becomes convinced that I will not stop playing King, then it is game-theoretically rational for him to play Pawn. Very little is known about the new group; no one knows who else is a member, or how they will play; no one has a solid basis for assuming that its behavioral statistics will be what a global post hoc view of the data has shown them to be. In such a state of ignorance, it is reasonable to think that some people could become convinced of a Demander's resolve, and thereby deliver the whole of the King's premium to the Demander. We have not experimented to see what happens when a small stable group is allowed to play for a long time (playing other individuals repeatedly), so we can only speculate on whether Demanders would eventually disappear.

A third type of explanation might actually create a stable persistence of the distribution of strategies that we observed, perhaps using an explore-versus-exploit idea. Just as data must be collected (at some cost) to be able to find the equilibrium of Will-Testing, data will probably need to be collected in order for Decision 1 to be resolved too. We have not yet developed a way of quantifying such a dynamic.

Conclusion

By breaking KingPawn into subparts, we have escaped from the unsatisfying naive equilibrium analysis. Without discerning what it is that causes two players to end up with a particular earnings ratio, the ratio has simply been taken as a given – and KingPawn has thus been opened to solution by recasting it as parameterized Will-Testing.

Satisfyingly, the new analysis predicts moderate amounts of conflict – which is a widely held intuitive prediction.

The behavioral average payoffs are decidedly more like the new prediction than the average payoffs of the naive KingPawn prediction. Hence we claim some validation of the theoretical analysis, although a marked discrepancy between data and prediction still remains. All the data is on one side of the theoretical curves. The arguments about data exploration and switchover costs correctly describe which side that is.

The collective behavior of KingPawn players is effectively imposing a payoff structure that encourages sharing. It is not claimed that this is cognitively recognized or deliberative in all players. The important observation is that the latent payoff curve serves to coerce the behavior of players in a way that the prima facie rules do not.

The distribution of "capitulation times" appears to agree with the predicted distributions well enough to further validate the use of Will-Testing to model KingPawn.

We believe that people have a sophisticated understanding of these games, deriving from much practice in daily social situations. Our view is that the humans who attended the game sessions arrived after having spent decades developing techniques for dealing with conflicting situations such as these. And we suggest that within their daily social lives, individuals have adopted personal constants that specify their tolerance for disruptive behavior and for the sacrifice they will make to expunge it. It is possible that through constant exposure to other people, we each come to embody characteristic social strategies that are successful *as a society*. People do not have to know each other personally to play KingPawn. Whenever subjects are drawn uniformly from a population that successfully plays similar games, the cohort will be comprised of people with the same distribution of constants and will thus also be able to successfully play the games.

We hypothesize that if a small cohort is allowed to play KingPawn long enough, the statistics of payoffs will convince most people to become Sharers, their overall group efficiency will rise, and the distribution of earnings among people will flatten out. This prediction is an easy consequence from having measured the relative expected payoffs for Demanders and Sharers; once people have enough samples to see this effect, they will presumably succumb to its logic. The important and subtle and sophisticated behavior detected by this study is that *the newly-assembled cohort is immediately ready to deploy an ensemble of strategies that encourage this prediction to come true.*

We believe that it requires considerable cognitive skill to perform this feat, and thus success at it is likely to be limited to big-brained animals. First, each player must be able to understand the game; nematodes are not likely going to succeed at this. Players must be able to perceive that the rewards are not monotonic functions of the action space and they must have the memory and perceptual equipment to pursue strategies through a complex time sequence. Second, players will benefit from perceiving the core abstraction of the game and being able to draw upon a history of experience at similar games to select a non-trivial strategy. Squirrels might not be able to wield this amount of abstract deliberation. Third, the player cohort needs to be ready to select strategies that are approximately at equilibrium ab initio (otherwise, a lot of time will be spent in a convergence phase to get close to an equilibrium). Grizzly bears, who do not spend a lot of time in groups, may not be quick at this.

We speculate that similar results might well arise if the game is played by a group of chimpanzees, or any animal that is both intelligent and highly socialized, but most species probably do lack the cognitive and game-theoretic ability to play either game as insightfully as humans do.

References

Bearden, J. N. 2001. Ultimatum bargaining experiments: The state of the art. *SSRN-id626183*.

Bowles, S., and Gintis, H. 2003. The origins of human cooperation.

Brown, R. 1986. *Social Psychology*. The Free Press, 2nd edition.

Camerer, C. F. 2003. *Behavioral Game Theory*. Princeton University Press.

Cesarini, D.; Dawes, C. T.; Fowler, J. H.; Johannesson, M.; Lichtenstein, P.; and Wallace, B. 2008. Heritability of cooperative behavior in the trust game. *PNAS* 105(10):37213726.

Fehr, E., and Gächter, S. 2000. Cooperation and punishment in public goods experiments. *American Economic Review* 90(4):980–994.

Gurvits, L., and Judd, J. S. 2012. The social will-testing game and its solution. *arXiv.org*.

Judd, S., and Kearns, M. 2008. Behavioral experiments in networked trade. *Proceedings of the ACM Conference on Electronic Commerce*.

Kearns, M.; Judd, S.; Tan, J.; and Wortman, J. 2009. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences* 106(5):1347–1352.

Kearns, M.; Suri, S.; and Montfort, N. 2006. An experimental study of the coloring problem on human subject networks. *Science* 313(5788):824–827.

Rabin, M. 1993. Incorporating fairness into game theory and economics. *American Economic Review* 83(5):1281–1302.

Sanfey, A. G.; Rilling, J. K.; Aronson, J. A.; Nystrom, L. E.; and Cohen, J. D. 2011. The neural basis of economic decision-making in the ultimatum game. *Science* 300:1755–1758.