Focal Points Revisited: Team Reasoning, the Principle of Insufficient Reason and Cognitive Hierarchy Theory

Nicholas Bardsley† and Aljaz Ule‡

†School of Agriculture, Policy and Development, Department of Food Economics and Marketing, University of Reading, UK. Corresponding author; n.o.bardsley@reading.ac.uk

‡Center for Research in Experimental Economics and Political Decision Making (CREED), Faculty of Economics and Econometrics, Universiteit van Amsterdam

DRAFT: Please do not cite without permission

September 2012

Focal Points Revisited: Team Reasoning, the Principle of Insufficient Reason and Cognitive Hierarchy Theory

Abstract

Coordination on focal points in one shot games can often be explained by team reasoning, a departure from individualistic choice theory. However, a less exotic explanation of coordination, based on best-responding to uniform randomisation, could explain much of the same data. We test the team reasoning explanation of coordination experimentally against this alternative, using games with variable losses under non-coordination. Subjects' responses are observed when the behaviour of their partner is determined in accordance with each theory, and under game conditions where behaviour is unconstrained. The results are more consistent with the team reasoning explanation. Increasing the difficulty of coordination tasks produces some behaviour suggestive of response to randomisation, but this effect is not pronounced.

Keywords

Coordination, cooperation, focal points

Introduction

Recent evidence from one-shot coordination games has been interpreted as showing that individuals making causally independent choices actually act in concert, asking themselves "What should we do?" (Bardsley *et al*. 2010). According to this 'Team Reasoning' (TR) hypothesis, an individual identifies a profile of strategies which is optimal for her team and then performs her part in it unconditionally. This has been invoked to explain coordination by Bacharach (1999, 2006) and Sugden (1995), drawing on Schelling (1960). We report on experiments that test the TR explanation against an alternative conjecture which grounds coordination in responses to potential randomisation by the other. This alternative can be rationalised under the Principle of Insufficient Reason (PIR) or in terms of Cognitive Hierarchy Theory (CHT). The PIR/CHT hypothesis is consistent with the usual individualistic reasoning of decision and game theory, with the agents asking themselves "What should I do?". Our results are more consistent with the TR explanation.

We define a coordination game as a game with multiple, strict, pure-strategy Nash equilibria, and null payoffs for other strategy combinations. We consider one-shot games only.[i] 'Pure' coordination games are defined as ones with payoff-identical equilibria. Nothing within their payoff structure enables a particular equilibrium to be selected by standard theory. Yet people often solve tasks which seem to instantiate them with high success rates (Schelling, 1960; Mehta *et al.,* 1994; Bacharach and Bernasconi, 1997; Bardsley *et al.* 2010).

In 'impure' variants the equilibria are Pareto-ranked. These games therefore seem even simpler, perhaps trivial, for real players. But they too are puzzling to many commentators, since within the standard framework of common knowledge of rationality the theoretical problem of equilibrium selection still obtains (Regan, 1980; Hollis, 1998; Bacharach 2006). Where there is a payoff dominant equilibrium (PDE) this serves empirically as a strong attractor, but its magnetism is essentially unexplained within game theory. For it is utility maximising to choose the PDE strategy if and only if one expects the other to do so with sufficient probability; but the same holds for every other equilibrium strategy. This leaves the expectation of PDE ungrounded. The same is true for other equilibrium refinement concepts, including risk dominance (which coincides with PDE in 2x2 coordination games).

TR and PIR/CHT offer competing explanations of equilibrium selection in one-shot coordination games. We set out the existing empirical support for TR in section 1 but show that best responding to randomisation can explain much of the same data in section 2. We then test the two accounts experimentally, using a game for which they produce distinct predictions (section 3). In a second experiment we increase the cognitive difficulty of coordination on the PDE to see whether this affects the relative success of the theories (section 4). The experiments use an original design in which, in one treatment, control over one subject's decision is allocated to the computer. This enables us to model, behaviourally, responses to randomisation and to TR, for comparison to game data. Section 5 provides interpretation and discussion, and section 6 concludes.

1. Apparent Evidence for Team Reasoning

On the TR account, faced with equal best equilibria agents transform a coordination game into a suitable impure coordination game. In the impure coordination game, agents consider which set of actions would be best for them and play their part in it, which rules out all payoff-dominated equilibria. Consider for example a game in which two players share a prize if and only if they nominate the same integer. According to Schelling (1960, p94), if players consider possible decision rules that might occur to their partner, including choosing a personal favourite, a culturally significant number and so on, each will be led to conclude that the best rule is to choose the number that is most clearly unique. This rule gives the best chance of coordination if both players adopt it. Because the number 1 is rather obviously unique in being the first integer, players using that rule will tend to coordinate on the number 1. In contrast, if they were merely picking a number for no particular reason they might well select a favourite number, or a culturally prominent one.

Mehta *et al.* (1994) reported cases of pure coordination games which confirmed Schelling's conjecture about differences between coordination and mere picking. However, relatively few tasks returned data with different picking and coordination distributions, and it was possible that biased beliefs about what was psychologically salient could explain the differences.[ii] Bardsley *et al.* (2010) therefore added a 'guessing' treatment to picking and coordinating treatments. Here, subjects had to guess what another subject had chosen in a picking treatment. TR predictions were based on characteristics such as 'odd man out' status, archetypal status, and indexical properties.[iii] Impure coordination games were also studied. The authors ran two experiments, one of which produced strong evidence in line with TR predictions.

For example, in one pure coordination game, the choice set was {Ford, Ferrari, Porsche, Jaguar}. In the picking treatment, the modal choice was {Jaguar}, but the guessers' modal choice was {Ferrari} and coordinators' mode was {Ford}. This accorded with a prior expectation that the cars would be categorised according to a luxury / ordinary brand distinction. That renders the options {the ordinary brand, a luxury brand} and the PDE is for both players to choose {the ordinary brand}, since this offers certainty of coordination rather than a 1/3 chance. In one impure coordination game, the choice set was {10, 10, 10, 9, 9, 8}. In the picking and guessing treatments, most subjects selected a {10}, but in the coordinating treatment {8} was the modal choice. This case rules out that subjects generally favour equilibria offering the highest payoffs of the (untransformed) monetary game. If the transformed options are {the 8, either 9, any 10} though, the equilibrium where both players choose {8} becomes the PDE under plausible assumptions about risk aversion.

2 An Alternative Mechanism for Coordination: Best-Responding to Randomisation

Best-responding to randomisation offers an explanation of coordination within individualistic rationality. One proposal is that in an impure coordination game the agents apply PIR, and assign equal probabilities to the other player's strategies (Gintis, 2003). If strategic reasoners start from this principle, then their best response will be to choose the strategy associated with the PDE. If both players reason in this way, then the agents will coordinate on that outcome. In impure coordination games, as defined in section 1, TR will therefore coincide with the application of PIR. This account amounts to an application of Harsanyi's 'tracing procedure' (Harsanyi and Selten, 1988) with PIR providing the initial beliefs. Gintis (2003) argues on these grounds that PIR renders TR superfluous as an explanation of coordination.[iv]

A very similar idea occurs as a version of CHT. CHT posits a population structured by different levels of rationality, and has been formalised in 'level-k' theories (Stahl, 1993; Stahl and Wilson, 1995; Camerer *et al.*, 2004).[v] Level 0 players are the least rational and choose nonstrategically. Level 1 players optimise based on their beliefs about level 0 players' behaviour. Level 2 players optimise based on their beliefs about the distribution and behaviour of level 0 and level 1 players, and so on. Agents in each level apart from level 0 optimise based on beliefs about the rest of the players, who are assumed to belong to lower tiers. If one assumes uniform randomisation for level 0 players, and unbiased expectations about lower tiers' behaviour, CHT makes the same equilibrium prediction as PIR. However, CHT is also capable of generating a richer set of predictions than PIR, based on auxiliary hypotheses about bounded rationality. We exploit this point in experiment 2 below.

It is important to note that coincidence between TR and PIR/CHT predictions actually occurs in both pure and impure coordination games. This is demonstrated by Bacharach and Stahl (2000)'s CHT-based framework 'Variable Frame Level-*n* Theory' (VFLNT). VFLNT invokes the same process of partition of the available strategies using a categorisation rule, or 'frame', as TR does. The 'variable frame' terminology reflects that more than one frame might apply, and that players at each level judge how probable different frames are to occur to players at lower levels. In the pure coordination game, with choice set {Ford, Ferrari, Porsche, Jaguar}, according to the TR argument expounded in section 1, the strategies are re-categorised as the options {the ordinary brand, a luxury brand}. In the re-categorised game coordination on {a luxury brand} yields 1/3 of the payoff from coordination on {the ordinary brand}. At that point, VFLNT models level 0 players as uniformly randomising over these two options, and the best response is {the ordinary brand}.

The availability of two explanations which both invoke players' unobserved re-descriptions of strategies threatens to seriously confound data interpretation in coordination studies. In both Blume and Gneezy (2010) and Crawford *et al.* (2008) for example, subjects had to coordinate on segments of partitioned discs, one of which is identified as unique by a framing involving shading. In each case, the prediction of coordination on this segment can be derived from either VFLNT or TR. Consequently, essentially the same behaviour is interpreted in Blume and Gneezy's design as evidence of the former, and Crawford *et al.'*s as evidence of the latter. The alternative readings seem equally justified, but invoke very different modes of reasoning.

3. Experiment 1: Game Play versus Response to Randomisation

Gintis (2003) describes a variation on a coordination game in which TR and PIR/CHT make clearly distinct predictions. This introduces risk in the sense of potential losses for coordination failure. If losses are variable, precautionary play can be separated from the PDE. In Gintis's example, each player has to choose an integer in the interval [1, 10]. If each selects the same integer, each wins that number of monetary units. If different integers are chosen, each loses the larger of the two numbers. This gives rise to the normal form game matrix shown in Figure 1. The game is doubly symmetric: both players either win or lose the same amount in each cell.

	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
1	1	-2	-3	-4	-5	-6	-7	-8	-9	-10
2	-2	2	-3	-4	-5	-6	-7	-8	-9	-10
3	-3	-3	3	-4	-5	-6	-7	-8	-9	-10
4	-4	-4	-4	4	-5	-6	-7	-8	-9	-10
5	-5	-5	-5	-5	5	-6	-7	-8	-9	-10
6	-6	-6	-6	-6	-6	6	-7	-8	-9	-10
7	-7	-7	-7	-7	-7	-7	7	-8	-9	-10
8	-8	-8	-8	-8	-8	-8	-8	8	-9	-10
9	-9	-9	-9	-9	-9	-9	-9	-9	9	-10
10	-10	-10	-10	-10	-10	-10	-10	-10	-10	10

Figure 1: Risky matching game

Here, choosing larger numbers increases the magnitude of prospective losses given uncertainty about the other's selection. Standard theories of choice under uncertainty, including Expected Utility theory and Prospect Theory (Kahneman and Tversky, 1979), predict that an agent responding to uniform randomisation should choose either {2} or {3}; this prediction carries over to PIR and CHT (proof: Appendix 1). The TR prediction is for both to choose {10}. Gintis (2003) proposes that in this game TR fails comprehensively, but does not cite empirical evidence. To the best of our knowledge such evidence does not yet exist. We therefore test the conjecture experimentally.

Subjects played the risky coordination game shown in Figure 1. The strategy set for each player consisted of integers in the interval [1,10]. In one treatment ('human computer,' or 'HC'), control over the actions of one player in each pair was taken away. Their strategy was determined by computer with uniform probability. The other player in each pair was told that this was how her partner's action would be determined, and had to choose an integer normally. In the second treatment ('human human,' or 'HH'), the same subjects played under standard game conditions, with each player freely choosing her integer.

Having a computer choose on behalf of a person seems to us better controlled than having subjects 'play against a computer'. For, although the determination of one player's action was shifted to the computer, a social choice situation was maintained, in the sense that each strategy selection affects the payoff of a pair of human subjects.

If coordination proceeds via responses to uniform randomisation, we should observe in HH the same pattern of choices as in HC, since HC implements randomness. According to CHT, any level 0 players will randomise, whilst players in level 1 *best* respond to randomisation, in both treatments, choosing from {2, 3}. Higher level players best respond to randomness in HC and to mixtures of lower level play in HH, but still choose from {2,3} (Appendix 1). If, alternatively, TR is the correct explanation of coordination, we would expect, in contrast, that players choose {10} in HH.

To summarise, in experiment 1, we test the following predictions:

i) TR predicts {10,10} in HH

ii) PIR/CHT predicts {2,3} in HC and HH

iii) PIR/CHT predicts there is no difference between distributions of choices in HC and HH

Minor caveats apply to predictions ii) and iii). Under the CHT account (but not PIR), there should be some unsophisticated players in the population, that is, level 0 players, who actually randomise uniformly over strategies. Thus, prediction ii) can be stated more precisely for CHT as a modal strategy choice of {2,3} with other choices uniformly dispersed. The proportion of level 0 players is often modelled as vanishingly small (Camerer *et al.*, 2004). Concerning prediction iii) PIR/CHT allows, only, for some switching from {2} in HC to {3} in HH depending on risk attitudes (Appendix 1).

3.1 Experiment 1: Procedures

Experiment 1 was conducted at the CREED laboratory at the University of Amsterdam (UvA) in June 2006, with 44 subjects. Each was given a show-up fee of 15 euros, in 30 experimental currency units, from which potential losses could be deducted. The design was counterbalanced, with half of the subjects playing HC before HH, and half the opposite order, to control for potential order effects. Treatment HC was divided into two tasks. In the first task the computer made the choice for one subject in each pair, and in the second task it made the choice for the other subject. Thus, there were three tasks per subject pair, two in HC and one in HH, from which two choices were observed per subject. The experiment lasted approximately 30 minutes including instructions, comprehension questions and a single sequence of the three tasks. No feedback was given on outcomes or earnings before the end of the experiment. The instructions are given in Appendix 2.

3.2 Experiment 1: Results

Choices are shown in Figure 2 below. Prediction i) is supported in the sense that the majority of subjects (66%) chose the TR prediction, {10}, in HH.

Prediction ii) is rejected since only a small minority of subjects (7% in HH and 11% in HC) chose a strategy from {2,3}. The modal choice in HC is {1}, which is stochastically dominated. If one therefore interprets choices of {1} in HH as flawed attempts to best respond to randomisation, counting {1,2,3} as consistent with PIR/CHT, this would only increase the proportion to 16% of subjects.

Prediction iii) is that there is no difference of any kind between choices in HH and HC. An appropriate nonparametric test is the chi-square test of independence. Since the test requires expected cell frequencies of at least 5 (Agresti, 1996), this requires combining response categories into bins. A simple method is to determine the bins from the data as follows. The mode is identified of HH and HC choices combined, and bins comprise the mode, integers below it and integers above it. (All data partitions and χ^2 tests in this paper, following this approach, are detailed in Table 1, section 4.2.) Here {10} is the overall mode and bins comprise {10} and {[1,9]}. We therefore test the null hypothesis of no difference between HC and HH using a chi-square test with one degree of freedom. The null hypothesis is rejected ($\chi^2(1)=23.2$; $p<0.01$). Thus, we find strong evidence against prediction iii).

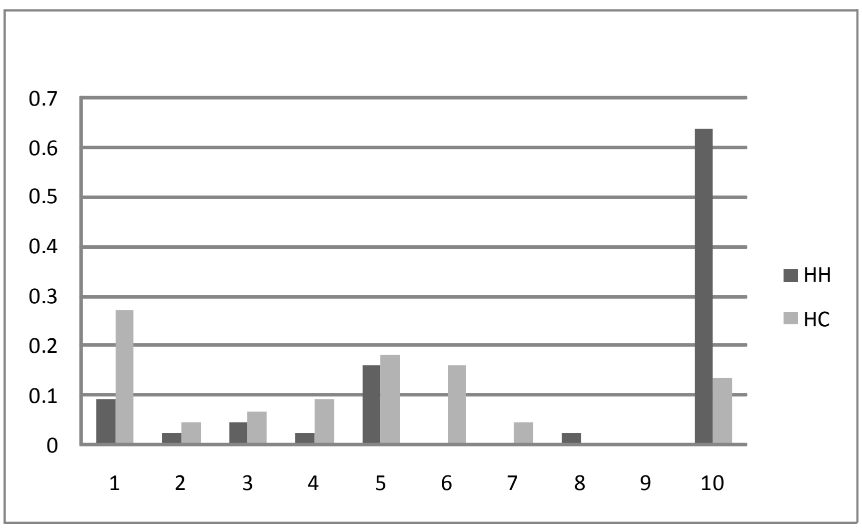

Figure 2: Frequency Distribution of Strategy Choices in Experiment 1

3.3 Interpretation of Experiment 1

The main result of experiment 1 is that TR strongly out-performs best-responding to randomisation in the game of Figure 1. The very different shapes of the distributions in HH and HC make it highly unlikely that HH choices are based on responses to randomisation. Subjects seem unable to optimise in a one-shot game, since the modal choice in HC, {1}, is stochastically dominated. As {1} is the lowest integer, participants were probably attempting to minimise exposure to loss. However, this description is incomplete, since HC choices suggest a doubly censored normal distribution with an interior mode at roughly the mid-point of the strategy space. It is therefore not obvious how best to characterise behaviour in the HC treatment overall. We also note that around 1/3 of subjects violate the TR point prediction in HH.

4. Experiment 2: Game Play versus Response to Team Reasoning

Experiment 1 returned evidence favourable to the team reasoning interpretation of coordination, and inconsistent with the PIR/CHT accounts based on best responding to randomisation. A strict falsificationist might advocate stopping there, so far as the latter are concerned. However, falsificationism has lost ground to views which see empirical work more as theory-developing than theory-refuting (Pawson and Tilley, 1997). The idea of best-responding to random behaviour seems strategically plausible to many and seems to have empirical support in some experimental contexts (Nagel, 1995). We therefore conducted a further test on the premise that there are some settings in which best responding to randomisation will operate and some settings more conducive to team reasoning. The aim of experiment 2 was to gain insight into the conditions under which the PIR/CHT account, and TR, either succeed or fail, with the goal of informing theory development in this area.

Experiment 2 attempted to undermine TR, and boost consistency with PIR/CHT, by increasing the cognitive difficulty of the coordination problem. The rationale for this is as follows. In a task as computationally easy as the risky coordination game of experiment 1, it is perhaps unrealistic to expect there to be a cognitive hierarchy. The 'hierarchy' of CHT, it seems to us, is dependent on the cognitive difficulty of the decision problem. Sufficient easiness will lead, in effect, to cognitive equality, but higher levels of difficulty should give rise to a ranking of abilities. Our auxiliary hypothesis on bounded rationality is that the parameters describing a cognitive hierarchy are endogenous to the choice problem. This is in line with Camerer *et al*.'s (2004, p863 n1) suggestion that the frequency distribution of player types may be sensitive to the costs and benefits of thinking harder.

We suggest specifically that in harder tasks the perceived net benefits of deliberation compared to randomisation are diminished, resulting in an increase in the proportion of level 0 type players. Further, actors should be more likely both to anticipate unpredictable behaviour, and responses to unpredictable behaviour, as difficulty increases. We should, then, be more likely to observe responses to randomisation, and less likely to observe TR, in harder tasks. We therefore aimed to induce a cognitive hierarchy by manipulating the difficulty of calculating the TR choice. This was done not to test subjects' maths ability, but to see whether behaviour is more consistent with the PIR/CHT account when we depart further away from the common knowledge of the game payoffs.

In experiment 1, coordination game decisions were compared to play against simulated randomisers so that the treatment comparison tests CHT. Experiment 2, in contrast, uses the HC condition to simulate team reasoners, so that the treatment comparison tests TR. This was done so as not to favour one theory over the other via an asymmetry in testing. In treatment HC, then, the integer of one of the paired players was predetermined according to the TR prediction. The other player was told that the computer had been programmed to enter the number which gives the highest joint earnings if both participants choose it. In HC, therefore, the choosing subject has to respond to TR. If TR is the only non-random process at work in HH, choices in HH and HC should be realisations of the same underlying distribution. According to CHT, in contrast, there will be some agents who can solve the TR computational problem but lack confidence that others can. So CHT predicts choices of lower integers in HH than in HC.

Three doubly symmetric games were used. They shared the feature with Experiment 1, that if the paired subjects chose different integers (again in the [1,10] range), they would both lose the larger number in currency units. If their chosen integers matched they would earn positive amounts. The winning amounts may, however, differ from the face value of the chosen integers, as set out below:

a. 'Low' difficulty. Matches on prime numbers pay their face value, while matches on other integers pay half their face value.
b. 'Medium' difficulty. A match on x pays its face value, where x = 8!/7!, while matches on all other integers pay half their face value.
c. 'High' difficulty. A match on x pays its face value, where $\sqrt[x]{59049} = 9$ while matches on all other integers pay 4.

As the labelling indicates, the tasks were constructed to increase difficulty of TR across a-c. In High, the lower-ranked equilibria each offered the same payoffs, one token less than the dominant equilibrium, to reduce further the perceived benefits of TR deliberations. The subject recruitment was not restricted to courses with mathematical content. We therefore expected that there would be considerable variation in participants' problem solving ability, and, therefore, good prospects of observing responses to randomisation in HH. PIR/CHT predicts low number choices for HH, with the exact prediction varying slightly between games as specified below.

These tasks give rise to the normal form game matrices shown in Figure 3 below.

Low

	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
1	1	-2	-3	-4	-5	-6	-7	-8	-9	-10
2	-2	2	-3	-4	-5	-6	-7	-8	-9	-10
3	-3	-3	3	-4	-5	-6	-7	-8	-9	-10
4	-4	-4	-4	2	-5	-6	-7	-8	-9	-10
5	-5	-5	-5	-5	5	-6	-7	-8	-9	-10
6	-6	-6	-6	-6	-6	3	-7	-8	-9	-10
7	-7	-7	-7	-7	-7	-7	7	-8	-9	-10
8	-8	-8	-8	-8	-8	-8	-8	4	-9	-10
9	-9	-9	-9	-9	-9	-9	-9	-9	4.5	-10
10	-10	-10	-10	-10	-10	-10	-10	-10	-10	5

Medium

	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
1	0.5	-2	-3	-4	-5	-6	-7	-8	-9	-10
2	-2	1	-3	-4	-5	-6	-7	-8	-9	-10
3	-3	-3	1.5	-4	-5	-6	-7	-8	-9	-10
4	-4	-4	-4	2	-5	-6	-7	-8	-9	-10
5	-5	-5	-5	-5	2.5	-6	-7	-8	-9	-10
6	-6	-6	-6	-6	-6	3	-7	-8	-9	-10
7	-7	-7	-7	-7	-7	-7	3.5	-8	-9	-10
8	-8	-8	-8	-8	-8	-8	-8	8	-9	-10
9	-9	-9	-9	-9	-9	-9	-9	-9	4.5	-10
10	-10	-10	-10	-10	-10	-10	-10	-10	-10	5

High

	1	2	3	4	5	6	7	8	9	10
1	4	-2	-3	-4	-5	-6	-7	-8	-9	-10
2	-2	4	-3	-4	-5	-6	-7	-8	-9	-10
3	-3	-3	4	-4	-5	-6	-7	-8	-9	-10
4	-4	-4	-4	4	-5	-6	-7	-8	-9	-10
5	-5	-5	-5	-5	5	-6	-7	-8	-9	-10
6	-6	-6	-6	-6	-6	4	-7	-8	-9	-10
7	-7	-7	-7	-7	-7	-7	4	-8	-9	-10
8	-8	-8	-8	-8	-8	-8	-8	4	-9	-10
9	-9	-9	-9	-9	-9	-9	-9	-9	4	-10
10	-10	-10	-10	-10	-10	-10	-10	-10	-10	4

Figure 3: Risky Matching Games in Experiment 2

An additional motivation for experiment 2 was to eliminate the possibility that subjects in HH are coordinating on salient features of the strategy space in something other than the team reasoning sense. For example, in experiment 1, it is conceivable that 10 is simply a salient number. To exclude this possibility, the strategy space is the same in each game, so number salience is held constant, whist TR selects a different integer in each case.

To summarise, in experiment 2, we test the following predictions:

iv. TR predicts {7} in Low, {8} in Medium and {5} in High in HH
v. TR predicts identical distributions of choices in HC and HH, in each case
vi. PIR/CHT predicts a mode of {2,3} in Low , {2} in Medium and {1, 2} in High, in HH
vii. PIR/CHT predicts strategic switching to lower choices in HH compared to HC
viii. TR will do progressively worse across Low, Medium and High in HH;
ix. PIR/CHT will do progressively better across Low, Medium and High in HH.

4.1. Experiment 2: Procedures

Experiment 2 was conducted at the CREED laboratory at the University of Amsterdam, in June 2010 and June 2011. Each subject was given a show-up fee of 15 euros, in 30 experimental currency units. Separate samples were drawn from the same student population for Low, Medium and High. Sample sizes were 30, 28 and 32 respectively. All subjects played treatment HH first and HC second in order to avoid biasing HH decisions in favour of TR. As in experiment 1, treatment HC was divided in two tasks and each subject played HC once actively, and once passively with the computer making her decision. The computer chose according to TR. Thus, there were three tasks per subject, two of which involved decision making. The experiment lasted approximately 30 minutes including instructions, comprehension questions and one sequence of the three tasks. No feedback was given on outcomes or earnings before the end of the experiment. Instructions are given in Appendix 2.

4.2. Experiment 2: Results

Subjects' choices are show in Figure 4 below. Concerning prediction iv), the TR point prediction is strongly modal for choices in HH in each game, with 46%, 50% and 50% of subjects making this choice in Low, Medium and High respectively. Prediction v) is tested with a chi-squared test. This is not significant at the 5% level for any of the three tasks, but is significant at the 10% level for Medium and High ($\chi^2(2) = 1.1$, p=0.57; $\chi^2(2) = 5.9$, p = 0.05, $\chi^2(2) = 5.4$, p = 0.07 respectively). However, combining data from the three games results in a test statistic of $\chi^2(2) = 9.5$, p<0.01, a strong rejection of the null hypothesis. Thus, prediction v fails.

Prediction vi) fares poorly in comparison to prediction iv), with relatively few subjects in HH choosing according to the PIR/CHT point prediction. 7% of subjects conform to this prediction in Low, 4% in Medium and 25% in High. However, as in experiment 1, one might interpret choices of {1} in Low and Medium as flawed attempts at PIR/CHT. This would alter the proportions to 13% and 29% respectively.

For prediction vii), a binomial test across the three games can be used to ascertain whether subjects who change their choice between HC and HH do so randomly. 35 subjects changed their decisions, with 26 of these choosing a lower number in HH. The null hypothesis that switches to higher and lower numbers are equi-probable is rejected (2-tailed binomial test, p < 0.01).

Prediction viii) is not supported by the data, since the proportion choosing consistently with TR does not differ significantly across the games.

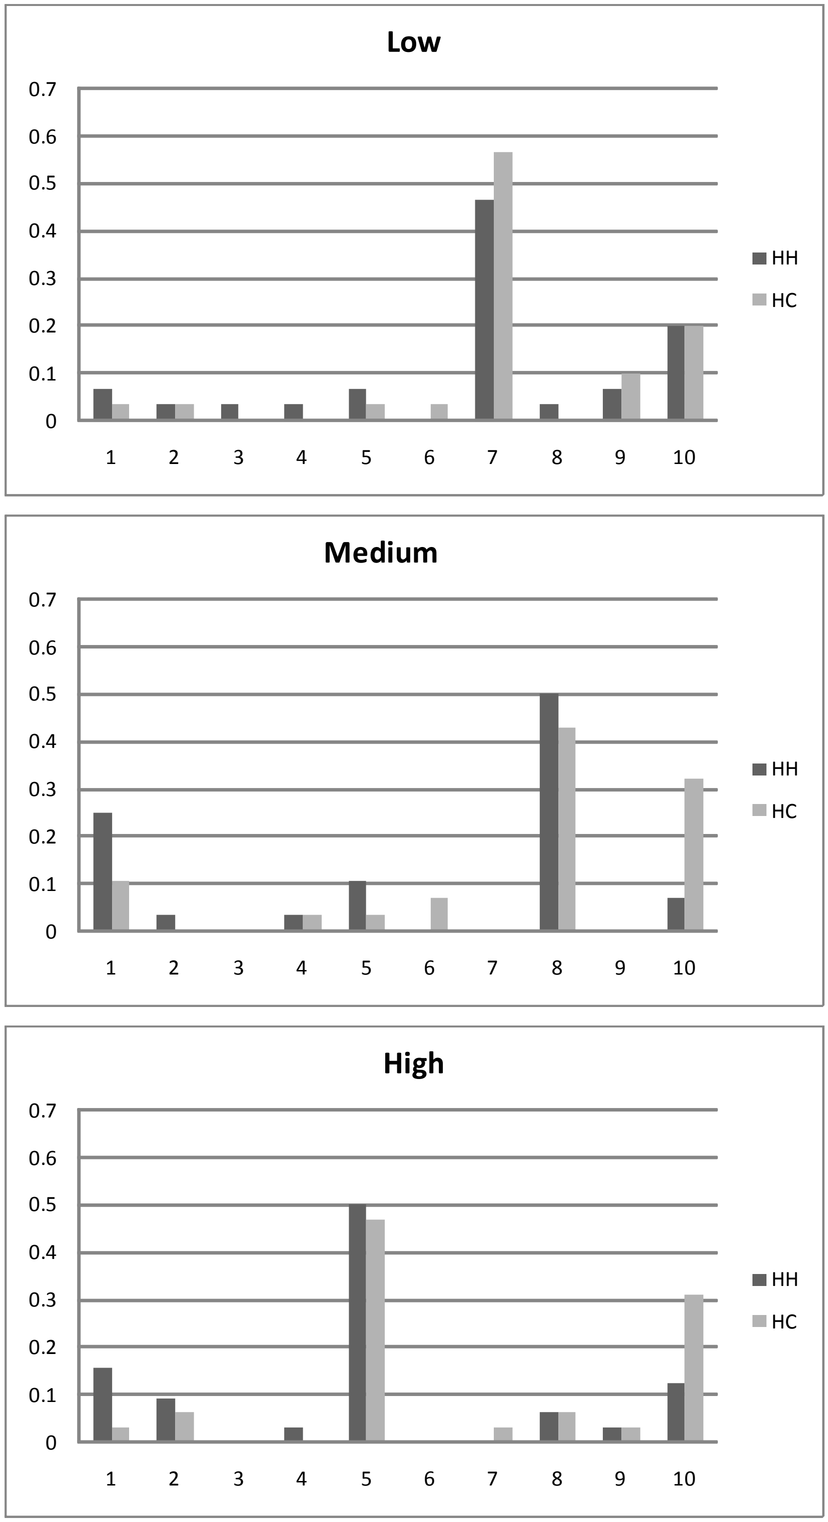

Figure 4: Relative Frequency Distributions of Strategy Choices in Experiment 2

Sample	Bins	Partitioned Distribution HH	HC	χ^2 Statistic (df)	P-value
Experiment 1	<10 10	16 28	38 6	23.2 (1)	<0.01
Experiment 2: Low	<7 7 >7	7 14 9	4 17 9	1.1 (2)	0.57
Experiment 2: Medium	<8 8 >8	12 14 2	7 12 9	5.9 (2)	0.05
Experiment 2: High	<5 5 >5	9 16 7	3 15 14	5.4 (2)	0.07
Experiment 2: Combined	<5 5 >5	23 21 46	9 17 64	9.5 (2)	<0.01

Table 1: Partitioned Distributions of Choices in Experiments 1 and 2, with Chi-squared tests

Note: In each case, HH and HC choices were combined to determine the overall mode of the distribution. The bins were then set as integers below, equal to and above this value in HH and HC separately. The requirement of the χ^2 test that expected cell frequencies are at least 5 precludes general use of a finer partition.

Prediction ix) can be assessed both in relation to the point predictions of PIR/CHT and in terms of its prediction of a treatment effect. Concerning the former, we compare the proportion of subjects behaving consistently with PIR/CHT in Low versus Medium, Low versus High and Medium versus High, using a 2-tailed Z test with Bonferroni correction for multiple comparison. The third test is significant at the 5% level (Z = 2.54; p = 0.02). The second is significant at the 10% level (Z = 2.05; p = 0.08). However, this analysis is dependent on *not* viewing choices of {1} in Low and Medium as attempts at best responding to randomisation. If instead we view choices of integers [1,3] as cohering with PIR/CHT in each game, as seems natural, there is no significant difference at the 10% level in each case. Concerning the latter, we judge whether changes of decisions to higher and lower integers are equi-probable for each game separately, using a 2-tailed Binomial test. In Low, 63% of switches were to higher integers in HC (p = 0.29), in Medium this fraction was 75% (p = 0.02) and in High 82% (p = 0.01). This pattern of results is supportive of prediction ix).

4.1. Interpretation of Experiment 2

Overall the results of experiment 2 favour TR over PIR/CHT in all three games, despite our attempt to make things difficult for TR. However, the manipulation of cognitive difficulty, as measured by the proportion of correct choices in HC, seems not to have been as effective as intended. In HC, the proportion played the TR prediction in High and in Medium is not significantly different. In spite of this, there is evidence of a tendency towards PIR/CHT type behaviour as difficulty increases, though it is not pronounced. It is clear from the failure of prediction v) that TR cannot be the only non-random process at work generating the observed data. The support for predictions vii) and ix) is consistent with the strategic anticipation of unpredictable behaviour in the manner envisaged by the PIR/CHT account.

5. Discussion

The main result of this study, which is consistent across both experiments, is that the TR predictions fare much better than the predictions of response to uniform randomisation in risky coordination games. It therefore seems implausible that PIR/CHT could account for the evidence that has been claimed for TR. When we simulate randomising players, we find differences in modal choices between HH and HC. When we simulate TR, we do not.

We conjectured that responding to randomisation was a plausible behavioural strategy where a cognitive hierarchy is likely to exist, and that this is more probable when tasks are more demanding. Therefore experiment 2 sought to increase the cognitive difficulty of the games. This resulted in some divergence between HH and HC. However, responding to randomisation did not become a very pronounced feature of the data as cognitive difficulty increased. This suggests that CHT with uniform randomisation at level 0 may have little behavioural significance for coordination games. This conclusion is drawn tentatively, as the manipulation of cognitive difficulty was not as effective as expected. As noted in the previous section there is nonetheless support in the data for a relatively weak tendency towards PIR/CHT, when the tasks became more difficult. An observation here on theory is that, in contrast to current CHT models, if level 0 frequency is close to zero it is likely to have little behavioural impact, especially if alternative modes of reasoning such as TR offer determinate advice.

A further reason that the PIR/CHT account performed relatively badly may be that uniform randomisation is not a good representation of what people do when a particular decision problem is beyond their ability to solve. This is suggested in particular by the pattern of HC choices in experiment 2, shown in Figure 4, in which the incorrect choices occurred with greater frequency *above* than *below* the correct answer. There, participants knew that their partner's strategy would be computationally correct, regardless of its difficulty. It may therefore become defensible to choose a high number, if a subject knows the solution is not a low number. For example, if a subject in High believed the answer to be 8, 9 or 10, with equal probability, they would expect equi-probable payoffs of 8, -9, or -10 from choosing {8}, -9, 9 or -10 from {9} and -10, -10 or 10 from {10}. Choices in the interval [1, 7] would be seen as dominated, {8} as stochastically dominated, and a risk neutral subject would choose from {9, 10}.

It therefore seems that actual behaviour in games when people are cognitively challenged is a complex matter. For example in High, people who were not able to spot the solution may have nonetheless have known that it was a number greater than, say, 3, if they understood the mathematical notation. In HH, they then also have to weigh the probability that their partner regards the problem as easy. This aspect of their decision is not currently represented in CHT, since CHT agents do not consider that others may be more sophisticated than themselves.

Regarding uniform randomisation as a characterisation of level 0, there is also evidence from "Buridan tasks" that it is difficult in practice to get people to randomise with uniform probabilities. Here, options are constructed so that there is no reason to choose one option rather than another. The tasks are named after the ass in the fable, which starved to death unable to choose between two equivalent piles of hay. Bacharach (2001) reports experiments on such problems, including pure coordination tasks against randomising devices. Subjects seemed to latch onto any available distinguishing features of options, rather than choosing at random.

Such considerations make it challenging to provide a tractable version of CHT capable of generating clear predictions for any possible game. Specifying level 0 behaviour *ex ante* is a key difficulty here. In the context of coordination games, an alternative strategy has been to adopt an empirical specification of CHT as, in effect, in Lewis (1969), Mehta *et al*. (1994) and Bardsley *et al*. (2010), focussing on predictions across treatments. Further behavioural research may determine whether an intermediate approach is possible, that is, one which organises insights across classes of games.

6. Conclusions

We subjected the TR explanation of coordination to strong experimental tests against PIR/CHT, which it largely withstood. Our data therefore support the view that TR is a key mechanism responsible for coordination. The alternative explanatory mechanism, best responding to randomisation, failed to organise the data in a class of games that was specifically devised to elicit it. Although we tried to make the alternative work by increasing the cognitive difficulty of the coordination problem, in the spirit of CHT, this had only limited success. The poor performance of responding to randomisation in the experiments reported here suggests that neither PIR nor CHT (with uniform randomisation at level 0) are probable explanations of the existing coordination game data. It seems rather that the explanatory mechanisms for focal points with empirical support are i) TR and ii) CHT with label salience at level 0, with ii) being necessary for 2x2 pure coordination games.

We suggest that behavioural economics could contribute to CHT by further observation of what people do when they are cognitively unable to optimise. Also, an empirically-supported account seems still to be wanting of the circumstances in which TR and CHT-type reasoning processes obtain, though our data suggest difficulty of the coordination task may play a role. Finally, we believe that it is interesting and important to conduct further robustness tests of TR given its radical break from received versions of methodological individualism, which rational choice theorists typically take as axiomatic (Elster 1982, 1985). One suggestion is as follows. The empirical research to date, including this report, has not sought to establish directly what is going on in game players' heads, preferring to work with choice data alone. We believe there is therefore a role for qualitative and possibly neurological research in future, to probe the TR hypothesis more directly.

References

Agresti, A. (1996). *An Introduction to Categorical Data Analysis*. Wiley: Chichester.

Bacharach, M. (2006). N. Gold and R. Sugden (Eds.) *Beyond Individual Choice. Teams and Frames in Game Theory.* Princeton University Press: Princeton, N.J.

Bacharach, M. (2001). Choice without preference: a study of decision making in Buridan problems. University of Oxford Discussion Paper.

Bacharach, M. (1999). Interactive team reasoning, a contribution to the theory of cooperation *Research in Economics*, 53, 117-147

Bacharach, M. and Bernasconi, M. (1997). The variable frame theory of focal points: an experimental study. *Games and Economic Behavior*, 19, 1-45

Bacharach, M. and Stahl, D.O. (2000). Variable-frame level-*n* theory. *Games and Economic Behavior*, 33, 220–46.

Bardsley, N., Mehta, J. Starmer, C. and Sugden, R. (2010). Explaining focal points. *The Economic Journal*, 120, 40-79

Bergstrom, T., Blume, L. and Varian, H. (1986). On the private provision of public goods. *Journal of Public Economics.*

Bjerring, A. (1978). The tracing procedure and a theory of rational interaction. In C.A. Hooker, J. Leach, & E. McClennen (Eds.), *Foundations and Applications of Decision Theory*. Reidel: Dordrecht.

Blume, A. and Gneezy, U. (2010). Cognitive forward induction and coordination without common knowledge: An experimental study. *Games and Economic Behavior*, 68, 488-511

Brandts, J. and Schram, A. (2001). Cooperation and noise in public good experiments: applying the contribution function approach. *Journal of Public Economics*, 79, 399-427.

Camerer, C.F., Ho, T. and Chong, K. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119, 861-898.

Crawford, V.P., Gneezy, U. and Rottenstreich, Y. (2008). The power of focal points is limited. Even a degree of payoff asymmetry may yield large coordination failures. *American Economic Review*, 98, 1143-1158.

Elster, J. (1982). The case for methodological individualism. *Theory and Society*, 11, 453-482.

Elster, J. (1985). *Making Sense of Marx.* Cambridge University Press: Cambridge.

Gintis, H. (2003). A critique of team and Stackleberg reasoning. *Behavioral and Brain Sciences*, 26, 160-161.

Harsanyi, J. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press: Cambridge, MA.

Hollis, M. (1998). *Trust within Reason*. Cambridge University Press: Cambridge

Kahneman, D. and Tversky, E. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, 47, 263-292.

Lewis, D. (1969). *Convention: a Philosophical Study.* Harvard University Press: Cambridge, MA.

Mehta, J. Starmer, C. and Sugden, R. (1994). The nature of salience: an experimental investigation. *American Economic Review*, 84, 658–673.

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

Pawson, R. and Tilley, N. (1997). *Realistic Evaluation*, Sage: London.

Regan, D.H. (1980). *Utilitarianism and Cooperation*. Clarendon Press: Oxford.

Schelling, T. (1960). *The Strategy of Conflict*. Harvard University Press: Cambridge, MA.

Skyrms, B. (1989). Correlated equilibria and the dynamics of rational deliberation. *Erkenntnis*, 31, 347–364.

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

Stahl, D.O. and Wilson, P. (1995). On players' models of other players. *Games and Economic Behavior*, 10, 218–254.

Sugden, R. (1984). Reciprocity: the supply of public goods through voluntary contributions. *The Economic Journal*, 94, 772-787.

Sugden, R. (1993). Thinking as a Team: Towards and Explanation of Nonselfish Behaviour. *Social Philosophy and Policy*, 10, 69-89.

Sugden, R. (1995). A theory of focal points. *The Economic Journal*, 105, 533-550.

Appendix 1

Proof of PIR/CHT Predictions for Experiment 1

Let j denote the opponent's chosen integer. The difference in utility, defined over experimental tokens, from choosing {i+1} over {i} is:

0 if $j > i+1$

U(i+1) – U(-i-1) if $j = i+1$

U(-i-1) – U(i) if $j = i$

U(-i-1) – U(-i) if $j < i$

If U′(i) > 0 for all i and the player evaluates equally the probabilities that its opponent chooses any strategy {j} then we can ignore probabilities and probability weights. It follows that

$$\{i+1\} \begin{smallmatrix}\succ \\ \sim \\ \prec\end{smallmatrix} \{i\} \leftrightarrow U(i+1)-U(i)+(i-1)[U(-i-1)-U(-i)] \begin{smallmatrix}> \\ = \\ <\end{smallmatrix} 0 \qquad (1)$$

For i=1 this reduces to U(2) - U(1) > 0, thus strategy {2} is always preferred to strategy {1}. Strategy {1} is in fact stochastically dominated by strategy {2}. For strategies {2},...,{10}, (1) implies {i+1} is weakly preferred to {i} if and only if

$$U(i+1) - U(i) \geq (i-1)[U(-i) - U(-i-1)] \qquad (2)$$

Consider $i \geq 3$. Under EUT with either risk aversion or risk neutrality, and also under Prospect Theory, U(i+1) - U(i) ≤ U(-i) - U(-i-1). Therefore (2) is not satisfied, and strategy {i} is preferred to strategy {i+1}. Hence, under standard models of choice under risk, strategies {2} and {3} are preferred to all other strategies.

Next, consider i=2. Under risk neutrality (2) holds with equality because of the assumption that U′=k, so {2}~{3}. Under risk aversion U″<0, and under Prospect Theory U′(x) < U′(-x). Either assumption implies that (2) does not hold, so {2} is strictly preferred to {3}.

Finally, as under risk neutral EUT, if each player believes that the other applies PIR, then from an interim conclusion that {2}~{3}, it follows that {3} is preferred, since {3} is the best response to a 50/50 chance that j={2} and j={3}. Under CHT, if for level 1 players {2}~{3} then for levels 2 and above {3} is preferred, if agents at those levels infer equi-probable choices from indifference at lower levels. The distribution should therefore have a single mode at {3}, with the relative frequencies of {2} and {3} depending on those of level 1 and higher-level players.

Parallel derivations can be given of CHT predictions in experiment 2.

Appendix 2

Instructions for Experiment 1

[*Our explanatory comments, not shown in the instructions, are shown between square brackets* [ ]. *Instructions are shown for the order HH-HC*]

Welcome to this experiment on decision making. In this experiment you can earn money. You have been given 30 points initially so that your points total cannot be negative. At the end of the experiment your points will be converted to cash, according to the exchange rate:
2 points = 1 euro.

The experiment consists of 3 independent tasks; what you earn in one task does not affect what you can earn in another. How much you will earn depends on your decisions and the decisions of one other participant, who we shall call "your paired participant". Your paired participant is randomly selected at the start and remains paired with you for all three tasks.

Throughout the experiment you will receive no information about the decisions of your paired participant or any other participant. At the end of the experiment you will learn the decisions of your paired participant, and will see your earnings. You will then be paid your earnings in private.

Next, you will be given a general description of the tasks. More detailed instructions will be given at the start of each task.

Please do not communicate with other participants at any time. If you have a question, please raise your hand. We will then come to your desk to answer it.

General Description

In each task two numbers will be determined between 1 and 10 (1 and 10 included), one for you and one for your paired participant.

- If these two numbers are the same, you will both win that number of points.
- If the numbers are different, each of you will lose the larger number of points.

The way in which the numbers are determined is different in each task.

To check your understanding, you will now have to answer some questions about the above procedure. You will receive further instructions when all participants correctly answer all questions.

Control Questions

To start the control questions please enter two **different** integer numbers between 1 and 10 in the two spaces below.

You have entered numbers X and Y. We will use these two numbers in the questions below.
If your number is X and the number of your paired participant is X, how many points does each win?
If your number is Y and the number of your paired participant is Y, how many points does each win?
If your number is X and the number of your paired participant is Y, how many points does each lose?
If your number is Y and the number of your paired participant is X, how many points does each lose?

Task 1 [HH]

In this task you and your paired participant each choose a number between 1 and 10 (1 and 10 included). Remember, if both numbers are the same, each of you wins that number of points. If the numbers are different, each of you loses the larger number of points. Both of you have been given exactly these instructions.
[*Both participants choose a number*.]

Task 2 [HC]

In this task you choose your number but the number of your paired participant is determined by computer. The computer has been programmed to enter any number from 1 to 10 with equal probability (1 and 10 included). Remember, if both numbers are the same, each of you wins that number of points. If the numbers are different, each of you loses the larger number of points.
[*The participant chooses the number, while the number for the paired participant is chosen according to the TR prediction.*]

Task 3 [HC]

In this task your number will be determined by computer, so you do not have to do anything. The other number will be chosen by your paired participant.

[*The participant's number is chosen according to the TR prediction, while the paired participant chooses her number herself.*]

END OF INSTRUCTIONS FOR EXPERIMENT 1

INSTRUCTIONS for Experiment 2

[*Our explanatory comments, not shown in the instructions, are shown between square brackets* [ ]. *Instructions are shown for the order HH-HC*]

Welcome to this experiment on decision making. In this experiment you can earn money. You have been given 30 points initially so that your points total cannot be negative. At the end of the experiment your points will be converted to cash, according to the exchange rate:

2 points = 1 euro.

The experiment consists of 3 independent tasks; what you earn in one task does not affect what you can earn in another. How much you will earn depends on your decisions and the decisions of one other participant, who we shall call "your paired participant". Your paired participant is randomly selected at the start and remains paired with you for all three tasks.

Throughout the experiment you will receive no information about the decisions of your paired participant or any other participant. At the end of the experiment you will learn the decisions of your paired participant, and will see your earnings. You will then be paid your earnings in private.

Next, you will be given a general description of the tasks. More detailed instructions will be given at the start of each task.

Please do not communicate with other participants at any time. Please also do not talk or give any comments during the experiment. During the experiment you are not allowed to ask us, or other participants, any questions. If something is wrong with your computer, please raise your hand. We will then come to your desk to check the problem.

General description

In each task two numbers will be determined between 1 and 10 (1 and 10 included), one for you and one for your paired participant.

-- If both of you choose the same number, [*in Low*:] and if it is a prime number, then you will both win that number of points. [*in Medium*:] and if it solves the equation [$x = 8!/7!$] you see written on the whiteboard, then you will both win that number of points. [*in Hard*:] and if it solves the equation [$\sqrt[x]{59049} = 9$] you see written on the whiteboard, then you will both win that number of points.

-- If both of you choose the same number, [*in Low*:] and it is not a prime number, then you will both win one half of that number of points. [*in Medium*:] and if it does not solve the equation you see written on the whiteboard, then you will both win one half of that number of points. [*in High*:] and if it does not solve the equation you see written on the whiteboard, then you will both win four points.

-- If the numbers are different, then each of you will lose the larger of these two numbers in points.

The way in which the numbers are determined is different in each task.

Control questions

You will now have to answer four questions [*with Yes/No*] about the above procedure. You are not allowed to ask us or other participants any questions about the above rules. You will receive further instructions when all participants correctly answer all questions.

1) If you and your paired participant get the same number, then you always earn points.
2) If you get a different number than your paired participant, then you lose points.
3) If you get a smaller number than your paired participant, then you lose his/her number in points.
4) You and your paired participant always earn or lose the same number of points.

Task 1 [HH]

In this task you and your paired participant each choose a number between 1 and 10 (1 and 10 included). Remember, if both numbers are the same, each of you earns points. If the numbers are different, each of you loses points. Both of you have been given exactly these instructions.
[*Both participants choose a number*.]

Task 2 [HC]

In this task you choose your number but the number of your paired participant is determined by computer. The computer has been programmed to enter the number which gives the highest earnings if both participants choose it. Remember, if both numbers are the same, you and your paired participant each earns points. If the numbers are different, each of you loses points.
[*The participant chooses the number, while the number for the paired participant is chosen according to the TR prediction.*]

Task 3 [HC]

In this task your number will be determined by computer, so you do not have to do anything. The other number will be chosen by your paired participant.
[*The participant's number is chosen according to the TR prediction, while the paired participant chooses her number herself.*]

END OF INSTRUCTIONS FOR EXPERIMENT 2

Supplementary Material

Experimental Design

Experiment 1

2 treatments were enacted in a within subject design. N=44 participants.

1 experimental token was worth E0.5; an endowment of 30 electronic tokens was pre-distributed.

Order of treatments was counterbalanced.

Paired subjects chose integers in the interval [1,10], winning the specified number of tokens in case of matching choices, but losing the larger number in case of non-matching choices. The game was one-shot.

In treatment 1 each participant decided their integer choice freely. In treatment 2 each player (in turn) had to choose knowing that their partner's integer was determined with uniform probability by the computer.

No feedback was given on outcomes or earnings before the end of the experiment.

The experiment was conducted in English.

Length: approximately 30 minutes.

Date: June 2006. Mean earnings: €8.30

Experiment 2

The design was as for experiment 1 with only the following differences:

3 games were studied with separate samples for each game. N = 30, 28 and 32 respectively.

Prizes for matching choices depended on specified mathematical properties of the integers in the interval [1,10]. We vary the payoff dominant equilibrium, and the difficulty of calculating it, across the three games.

Treatment 1 took place before treatment 2; the design was, intentionally, not counterbalanced.

In treatment 2 each player (in turn) had to choose knowing that their partner's integer was determined via the computer program to be the integer yielding the highest payoff for each player in case of matching. This was known in advance of treatment 2.

Date: 50% of observations were collected in June 2010, and 50% in June 2011. Mean earnings: €10.80

Instructions

The instructions are given in Appendix 2 of the main text.

Selection and Eligibility of Participants

Participants were students at the University of Amsterdam, recruited by free enrolment into sessions, following the standard recruitment procedures at the Center for Experimental Economics and Political Decision-Making (CREED). CREED maintains a database of several thousand past and prospective participants. The database keeps track of all past individual participation. Recruitment for all sessions is organized via public announcement to all individuals in the database that had not participated in a similar experiment before. Individuals subscribe via a dedicated webpage www.creedexperiment.nl, where they select at most one session of the experiment, and are permitted to subscribe only if they never before participated in a similar experiment. The database is regularly enlarged through public advertisments to the students of various disciplines at the University of Amsterdam and the Free University in Amsterdam, who comprise the majority of participants. The experiment was open to students regardless of degree program or year. Every participant participated in only one session. It was specified in advance that the experiments would be conducted in English.

---

[i] Repeated games provide additional means of coordination via signalling which would confound the study of the issues we are investigating.

[ii] Lewis (1969) had proposed an account of coordination based on psychological salience, meaning that certain options just have attention-grabbing properties which serve as tie-breakers when there are equal-best options.

[iii] An indexical property of a linguistic item is one that is relative to the circumstances of its occurrence. So, for example, in an experiment conducted in 2005, 2005 might be focal in a choice set consisting of items labelled {2004, 2005, 2006, 2007} because it is the current year.

[iv] As an account of rational coordination, this is controversial. For the assumption used to derive the players' beliefs is apparently contradicted in the agents' chain of reasoning (Bjerring, 1978). Initially there is a stipulation of uniform probabilities, but the players conclude that a particular strategy will be played with probability 1. Regardless of this issue, however, it seems that PIR may still have promise as an empirical account of coordination for imperfectly rational actors. The players may, for example, treat implications of their initial assumptions as new information, as in Skyrms (1989).

[v] For simplicity we do not observe a distinction in the text between CHT and level-k theory, since they coincide predictively for the games we study. The theories actually differ in that in level-k theory, each level optimises in response to behaviour of the next lowest level. Whereas in CHT each level optimises in response to a finite mixture distribution defined over perceived player types and frequencies at lower levels.