

Payoff inequity reduces the effectiveness of correlated–equilibrium recommendations

Nejat Anbarcı

Department of Economics
Deakin Business School
Geelong VIC 3220, Australia
nejat.anbarci@deakin.edu.au

Nick Feltovich*

Department of Economics
Monash Business School
Clayton VIC 3800, Australia
nicholas.feltovich@monash.edu

Mehmet Y. Gürdal

Department of Economics
Boğaziçi University
Bebek, Istanbul, Turkey TR–34342
mehmet.gurdal@boun.edu.tr

June 1, 2018

Abstract

We examine theoretically and experimentally how individuals' willingness to follow third-party recommendations in 2x2 games is affected by payoff asymmetry. We consider six versions of Battle-of-the-Sexes. Recommendations imply monetary payoffs that are equal ex ante, but unequal ex post. So, although following recommendations constitutes a Nash equilibrium under standard preferences, sufficiently inequity-averse players can rationally disobey a recommendation that would lead to a very unfavourable payoff distribution, as long as the cost of doing so is not too large. Our theoretical model incorporates inequity aversion, along with level- k reasoning. Our main experimental result is consistent with the model: as either payoff asymmetry increases or the cost of disobeying an unfavourable recommendation decreases, subjects are more likely to disobey recommendations.

Journal of Economic Literature classifications: C72, D81, D83, D03, C92.

Keywords: correlated equilibrium; recommendations; inequity; other-regarding behaviour; level- k reasoning.

*Corresponding author. Financial support from Boğaziçi University (BAP 7921) and Deakin University is gratefully acknowledged. We thank John Duffy, Philip J. Grossman, Doruk Iris, Vai-Lam Mui, Aaron Nicholas, two anonymous referees, and participants at several conferences and seminars for helpful comments.

1 Introduction

Coordination problems are ubiquitous in economics. Theoretical study of these settings dates back at least to the work of Schelling (1960), and experimental study dates back almost as far (see Ochs, 1995 and Camerer, 2003, pp. 336–407 for surveys of early experiments). From the beginning, experiments have confirmed that coordination failure is a very real possibility, and a large strand of this literature has investigated ways of overcoming coordination problems. Some well-known tools are shared history (Van Huyck et al., 1991), focal points (Mehta et al., 1994), communication (Farrell, 1987), financial incentives (Brandts and Cooper, 2006) and competition (Bornstein et al., 2002).

Another tool that has received some attention is the use of external signals. While theoretically signals serve mainly to expand the set of equilibria (which in principle ought to exacerbate the coordination problem), behaviourally they often improve coordination substantially. This is especially true when signals take the form of *recommendations* – that is, when the mapping from signals to intended outcomes is so obvious that it arguably approaches common knowledge – as the literature on correlated equilibrium induced by computer-generated recommendations (Cason and Sharma, 2007; Duffy and Feltovich, 2010; Bone et al., 2013; Kurz et al., 2017) has shown.¹

As a simple example of how these recommendations can work, consider two cars approaching a four-way intersection at night, one from the south (S) travelling north and the other from the west (W) travelling east. One needs to yield or stop so that the other can pass, but both drivers prefer that the other car does so. There are two pure-strategy Nash equilibria (driver S stops to let driver W pass, or W lets S pass), but neither is likely to be focal, since neither equilibrium exhibits *label salience*, while each is *payoff salient* for only one of the two drivers.² Moreover, solving the coordination problem by communication between the drivers is likely to be difficult. Suppose, however, the intersection contains a traffic light that – as the cars approach – is equally likely to show green to either car (and red to the other). As long as both drivers understand the signals’ literal meanings, this source of third-party recommendations will solve the coordination problem. Indeed, the recommendations are self-enforcing: as long as the driver facing the red light understands that the other car (facing a green light) will continue without stopping, he prefers to stop, and similarly the driver facing the green light will prefer to go.

The experimental literature on third-party recommendations used for solving coordination problems, while fairly small, has already examined some of the factors that make such recommendations more or less effective. Cason and Sharma (2007) found that recommendations were more effective when the level of strategic uncertainty was reduced (by having subjects play against computers that always followed recommendations, rather than against other humans). Duffy and Feltovich (2010) found that recommendations were more effective when they induce a correlated equilibrium (i.e., when following them constitutes a Nash equilibrium) with expected payoffs higher for both players than in the symmetric Nash equilibrium. Bone et al. (2013) find that recommendations are less effective when there exists a cooperative outcome that maximises joint payoffs but which is not recommended by the third party. (See also Moreno and Wooders, 1998, which to our knowledge is the first experimental paper looking specifically at correlated equilibrium.)

In this paper, we examine another potentially relevant factor: the *fairness* of the recommendations. To illustrate its potential impact, return to our traffic-light example, but suppose the intersection is near a chemical plant produc-

¹See also the closely related experimental literature on sunspots (Duffy and Fisher, 2005).

²The “label salience” of an action profile is the extent to which it is focal due to the names of its component actions, while its “payoff salience” arises from the corresponding payoffs. These concepts come from Schelling (1960), though to our knowledge he did not use the terms themselves.

ing a very unpleasant smell; the driver who stops is exposed to the foul smell for the duration of the red light, while the other driver mostly avoids it.³ The worse the disutility from the odour, the greater is the difference between the payoff to the driver getting the green light and the driver getting the red light; this is true even though the drivers' expected payoffs are *ex-ante* equal.⁴ If this difference becomes sufficiently large, the driver facing the red light may well be tempted to ignore it and take his chances with the cross-traffic.

Consider the class of Battle-of-the-Sexes (BoS) games shown in Figure 1, with $0 < n < M$. Each member of

Figure 1: General form of our BoS games (with $0 < n < M$)

$G(M, n)$		Player 2
		A B
Player	A	0, 0 M, n
1	B	n, M 0, 0

this class has two strict pure-strategy Nash equilibria: (A, B) and (B, A), in addition to a mixed-strategy equilibrium that is symmetric but Pareto dominated by either of the pure equilibria. Without any way of breaking the symmetry between the players (who face an identical decision), behaviour is likely to be characterised by the mixed equilibrium. However, both players stand to benefit if a non-strategic third party recommends either of the pure equilibria with equal probability (i.e., with probability one-half, it recommends Player 1 to choose A and Player 2 to choose B, and with probably one-half, it reverses the recommendations). Not only are expected payoffs higher than in the mixed equilibrium, but following the recommendations is itself an equilibrium.

The equilibrium induced by recommendations is *ex-ante payoff equitable*, since if both players follow recommendations, each gets the same *ex-ante* (i.e., before recommendations are received) expected payoff of $(M + n)/2$. But it is not *ex-post payoff equitable*: after the recommendation is made, one player is “favoured” (with a payoff of M if recommendations are followed) and the other is “unfavoured” (with a payoff of $n < M$). If both players are *self-regarding* (own expected payoff maximising), then this *ex-post* unfairness is irrelevant; they will follow the recommendations no matter how unequal the payoffs become, as long as n remains positive.⁵ However, there is evidence (see, e.g., the examples mentioned by Fehr and Schmidt, 1999 and Bolton and Ockenfels, 2000) that a sizeable fraction of people are *inequity averse*; while they prefer more money to less (*ceteris paribus*), they also dislike outcomes where they receive less (or even, in some cases, more) than the other player. In our BoS games, an inequity-averse player will flout a B recommendation (i.e., choose the opposite action) if the associated outcome is sufficiently inequitable and if cost of flouting the recommendation is sufficiently low – that is, he will prefer to forgo a positive payoff of n with his opponent getting a higher payoff of M , in favour of both him and his opponent getting equal payoffs of 0.⁶

³Alternatively, loud music might be coming from a nearby house, or the intersection might be in a high-crime area so that the driver who stops risks being car-jacked.

⁴Kurz et al. (2017) consider *ex-ante* payoff asymmetry (in their words, “procedural fairness”) in their experiment (in our traffic-light example, how likely each player is to get the green light). They find that it has little effect on whether recommendations are followed.

⁵As Schelling (1960, p. 144) noted, “The white line down the center of the road is a mediator and very likely it can err substantially towards one side or the other before the disadvantaged side finds advantage in denying its authority.”

⁶The vast literature on ultimatum-game (UG) experiments (for surveys, see Roth, 1995 and Camerer, 2003, pp. 151–198) shows that in some cases, individuals will indeed choose a payoff of zero over a small positive payoff if it means the opponent will also get zero instead of

The current paper is a theoretical and experimental exploration of this possibility. We consider six members of the class of games shown in Figure 1; the specific games are shown in Figure 2 below. Under standard theory, there is no reason to expect any systematic differences across these games in whether recommendations are followed or in whether they are sought. However, in Section 2 we construct a behavioural model that incorporates Fehr and Schmidt’s (1999) “inequity aversion” and the “level- k reasoning” model of Nagel (1995) and Stahl and Wilson (1994, 1995). Predictions based on analysis of this model closely follow the lines of the intuitive argument above: recommendations are *less* likely to be followed, and indeed sought, as (a) payoff asymmetry *increases* (i.e., as $M - n$ increases), and as (b) the cost of flouting recommendations *decreases* (i.e., as n decreases). We chose our games so that three of them vary in payoff asymmetry while holding constant the cost of flouting recommendations, while three vary in the cost of flouting recommendations while holding payoff asymmetry constant – thus allowing separate examination of both (a) and (b).

In our lab experiment, besides the game, we vary the availability of third-party recommendations. In our *no-recommendations* (N) treatment, players simply choose actions without receiving recommendations. In our *certain-recommendations* (C) treatment, players receive computer-generated recommendations as described above, prior to making their action choices. A third *voluntary-recommendations* (V) treatment is similar, except that one of the two players is randomly selected to choose – on behalf of both players – between asking and not asking for recommendations. If she chooses to seek recommendations, they are received by both players, while if she chooses not to seek recommendations, neither player receives one.

Our main result is that the response to unfavourable B recommendations varies systematically with the game as predicted by our model: they are followed *less often* as payoff asymmetry *increases*, as the cost of flouting recommendations *decreases*, or both. By contrast, favourable A recommendations are nearly always followed (this is roughly consistent with our model, which predicts they are always followed). The different reaction to A and B recommendations is notable because it suggests that to the extent recommendations are unsuccessful, it is not due simply to the absence of an obvious focal point; rather, some subjects knowingly avoid a positive but inequitable payoff in favour of both getting a zero payoff. Finally, we do not find systematic differences in how often recommendations are sought, in contrast to the model’s predictions.

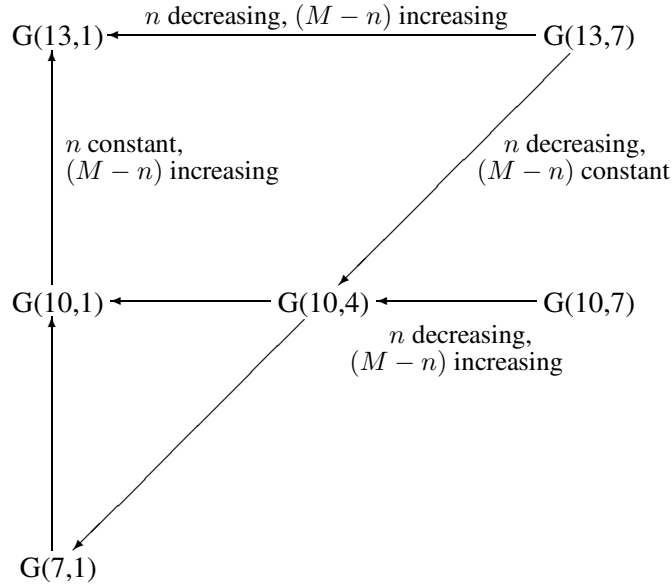
2 Theoretical background

The stage game has the form shown in Figure 1 above, with six constellations of M and n (see Figure 2), all with $M > n$. Since recommended outcomes are always either (A, B) or (B, A), they always imply a payoff of M for the player receiving an A recommendation and n for the player receiving a B recommendation; we will refer to them as the “favoured” and “unfavoured” player respectively (and by the same token, to M and n themselves as the “favoured” and “unfavoured” payoffs). Also, we call $(M - n)$ the “level of payoff asymmetry” of a particular game, and n the “cost of flouting recommendations”.

As Figure 2 shows, the particular games we chose satisfy three important three-game orderings. First, moving from G(7,1) to G(10,1) and from there to G(13,1) increases payoff asymmetry while holding constant the cost of flouting recommendations. Second, moving from G(13,7) to G(10,4) and thence to G(7,1) lowers the cost of flouting

a larger positive payoff. Obviously, the UG differs in many ways from the games we consider, not least in how this kind of decision arises (directly from an opponent choice in the UG, versus based largely on a random draw in our setting). Some experiments have found that the source of decisions – that is, “attribution” (Charness, 2004) – is an important factor underlying decision making. In our conclusion, we briefly discuss attribution in light of our own experimental results.

Figure 2: Games $G(M,n)$ used in the experiment



recommendations while holding payoff asymmetry constant. Third, moving from $G(10,7)$ to $G(10,4)$ to $G(10,1)$ simultaneously increases payoff asymmetry and lowers the cost of flouting recommendations.

Under the standard assumptions of self-regarding preferences and correct beliefs about opponent play, we consider two solution concepts. In the mixed-strategy Nash equilibrium (NE), each player chooses A with i.i.d. probability $M/(n + M)$. In the correlated equilibrium induced by our recommendations, the outcomes (A, B) and (B, A) each occur with probability one-half. Table 1 shows some implications of these solution concepts: the associated

Table 1: Implications of mixed Nash equilibrium (NE) and correlated equilibrium based on recommendations (CE)

Game (M, n)	Prob(A choice)		Prob(coordination)		Mean payoff	
	NE	CE	NE	CE	NE	CE
G(7,1)	0.875	0.500	0.219	1.000	0.875	4.000
G(10,1)	0.909	0.500	0.165	1.000	0.909	5.500
G(13,1)	0.929	0.500	0.133	1.000	0.929	7.000
G(10,4)	0.714	0.500	0.408	1.000	2.857	7.000
G(10,7)	0.588	0.500	0.484	1.000	4.118	8.500
G(13,7)	0.650	0.500	0.455	1.000	4.550	10.000

probability of choosing A, the probability of successful coordination on a pure-strategy Nash equilibrium (a realisation of either (A, B) or (B, A)) and the expected ex-ante payoff (which is equal for the two players since the game, the mixed-strategy equilibrium and the recommendation distribution are all symmetric).

Previous research (Duffy and Feltovich, 2010) found that in games similar to ours, aggregate behaviour is described well by the relevant correlated equilibrium when recommendations that correspond to a correlated equilib-

rium are available, and by the symmetric equilibrium when they are not.⁷ That result, combined with the information in Table 1, can be used to construct a very simple model of behaviour. First, when recommendations are not given (our N treatment), players play according to the mixed–strategy Nash equilibrium due to its symmetry and the lack of any better symmetric equilibrium. Second, when recommendations are given (our C treatment), the implied correlated–equilibrium outcome payoff–dominates the mixed–strategy Nash equilibrium; hence players follow recommendations. Third, since recommendations lead to higher payoffs than no recommendations, recommendations will be sought when players have this choice (as one player within each pair will in our V treatment).

2.1 A model of other–regarding preferences

In this section we examine how a distaste for disadvantageous outcomes might affect behaviour. We do so by constructing and analysing a model that, while still fairly simple and tractable, incorporates such a distaste. As we will see, this model will allow us to formulate non–trivial directional hypotheses: different behaviour across games. This is in contrast to the simpler model described in the previous paragraph, which implies no differences in behaviour across games when recommendations are received.

We depart from the standard model in two ways. First, preferences allow for aversion to unfair outcomes. Specifically, preferences are as described by an extension of Fehr and Schmidt’s (1999) “inequity aversion” model. If x_i is player i ’s monetary payoff, then her full preferences are given by

$$V_i = \delta U_i[E_p(x_i), E_p(x_j)] + (1 - \delta)E_p[U_i(x_i, x_j)] \quad (1)$$

where

$$U_i(x_i, x_j) = x_i - \alpha \cdot \text{Max}\{x_j - x_i, 0\} - \beta \cdot \text{Max}\{x_i - x_j, 0\} \quad (2)$$

with $\alpha > \beta$, $\beta < 1$ and $j = 3 - i$, and E_p is the expectation operator with respect to p , the lottery over outcomes.

The function $U(\cdot)$ in (2) is from Fehr and Schmidt (1999), and captures aversion to inequality with α being aversion to disadvantageous inequity and β being the aversion to advantageous inequity. Saito (2013) proposes (1) as an extension for settings with objective uncertainty. There, δ measures the weight that the decision maker gives to ex ante equity (equality of opportunity), whereas the remainder $1 - \delta$ measures the weight given to ex post equity (equality of outcome).

Our second departure from the standard model is that we allow beliefs about other players’ decisions to be incorrect. Beliefs are given by “level– k reasoning” (Nagel, 1995; Stahl and Wilson, 1994, 1995), in which players (with one exception) best–respond, given their preferences, to possibly–incorrect beliefs about how the opponent population behaves. The exception is the non–strategic type L0, who is assumed to blindly follow a recommendation if one is received, and play randomly if not. For $k \geq 1$, the type L k assumes the entire population is of type L($k-1$). We make the following assumption regarding the *actual* frequencies of types in the population.

⁷Similar results have been observed by others (e.g., Cason and Sharma, 2007; Bone et al., 2013; Georgalos et al., 2018). We also note here that other Nash equilibria and other correlated equilibria exist for this class of games. Our justification for appealing to symmetry when recommendations are not given comes from our use of a one–population matching protocol in the experiment. This means that players’ roles themselves cannot serve as a coordinating device for “symmetry breaking” (Crawford, 1998), and since there is no other external mechanism (such as recommendations), asymmetric equilibria are implausible. Similarly, our particular choice of correlated equilibrium arises from the distribution of recommendations given to subjects; other correlated equilibria are unlikely due to the seeming lack of any other coordination device, though arguments could perhaps be made for (i) both players flouting recommendations (i.e., choosing A if and only if receiving a B recommendation), or (ii) both players ignoring recommendations in favour of the mixed–strategy equilibrium. For the case of voluntary recommendations, a plausible alternative prediction involves the decider choosing not to seek recommendations and then choosing A in order to obtain the best possible payoff. We discuss this “signalling” alternative in the appendix.

Assumption 1 *The population frequencies of L1 and L2 are θ and $1-\theta$ respectively, with $\theta \in [0, 1)$.*

Assumption 1, which implies that all other types (L0 and those higher than L2) have zero frequency, is common in level- k models for reasons of both tractability and empirical validity (see, e.g., Crawford et al., 2008).

Even for fairly large values of β (aversion to advantageous inequity), its value does not change our analysis from the case when $\beta = 0$; therefore for tractability we set it to zero. Let $F(\cdot)$ be the cumulative distribution function for α (aversion to disadvantageous inequity, abbreviated ADI) in the population. The following assumption guarantees that recommendation A is always followed:

Assumption 2 *$F(\cdot)$ is continuous and strictly increasing over an interval containing $(0, 7/3)$ and*

$$F\left(\frac{n}{M-n}\right) > \frac{n}{M} \quad (3)$$

for all M and n .

We note that the inequality in Assumption 2 is fairly weak given the results of previous attempts to measure the distribution of α in the population. For instance, Fehr and Schmidt (1999) estimate from several experiments that $F(1/2) = 0.6$, $F(1) = 0.9$ and $F(4) = 1$, while our inequality requires only $F(1/2) > 1/3$, $F(1) > 1/2$ and $F(4) > 4/5$ (substituting (3,1), (2,1) and (5,4) respectively for (M, n) in (3)). Also, while setting n to zero (or taking the limit as $M \rightarrow +\infty$) in the inequality implies that a strictly positive fraction of players is self-regarding ($\alpha = 0$), this fraction can be made arbitrarily small if desired.

Assumptions 1 and 2 allow us to state our main result.

Proposition 1 *Suppose Assumptions 1 and 2 hold. Then*

1. *When there is no recommendation, the frequency of A choices increases as payoff asymmetry increases.*
2. *When recommendations are received, the frequency with which a B recommendation is followed increases as payoff asymmetry decreases or as the cost of flouting recommendations increases (ceteris paribus), while A recommendations are always followed.*
3. *When players can choose to ask for recommendations, the demand for recommendations weakly increases as payoff asymmetry decreases or as the cost of flouting recommendations increases.*

Proof: In Appendix.

We finish this section with some intuitive reasoning behind Proposition 1: how it derives from the two components of the behavioural model (inequity aversion and level- k reasoning) and what happens when one of these components is not present. In Part 1, according to our model, L1 types believe that L0 types will mix evenly between A and B, so they best-respond by choosing A. An L2 player thus believes that her opponent will play A for sure, and so she will play B unless her aversion to disadvantageous inequity is above a threshold level, in which case she will play A. As $(M - n)$ increases, this threshold level of ADI decreases, meaning that a higher fraction of L2 players will play A.

For Part 2, L1 types believe that L0 types will play as recommended. Therefore they will best-respond by choosing A if it is recommended, but they will follow a B recommendation only if their ADI is low enough compared to $n/(M - n)$ – that is, if the level of payoff asymmetry $(M - n)$ is high enough or the cost n of disobeying a B

recommendation is low enough – as will L2 types due to our Assumptions 1 and 2. Then, both L1 and L2 types are more likely to follow a recommendation as $(M - n)$ decreases and/or n increases.

Unlike the rest of Proposition 1, Part 3 requires Saito’s (2013) extension of Fehr and Schmidt’s (1999) inequity aversion, expressed in (1). If an L1 player with low ADI is the “decider” (the subject within the pair who was randomly chosen to decide), his utility from seeking recommendations will increase in $n/(M - n)$. So as either the payoff asymmetry $(M - n)$ decreases and/or the cost n of disobeying recommendations increases, he will prefer to seek recommendations if his ADI is sufficiently high. If an L2 player is chosen instead, he will believe that L1 types play A in case of no recommendation. Thus, an L2 player with a low ADI will always ask for recommendations and play A – knowing that a sufficient fraction of L1 types will play B when recommended – but an L2 player with a sufficiently high ADI will always play the recommended action and his expected payoff from asking for a recommendation will decrease as $n/(M - n)$ decreases. That is, the expected payoff will decrease as either payoff asymmetry increases or the cost of disobeying recommendations decreases.

In order to further understand the roles of inequity aversion and level- k reasoning, we briefly examine the implications of alternative models where either inequity aversion, level- k , or both is not present. The first part of Proposition 1 does not rely on either component of our model; the same theoretical result follows from symmetric mixed-strategy Nash equilibrium under standard preferences (see Table 1) as well as McKelvey and Palfrey’s (1995) logit equilibrium under standard preferences.

By contrast, the second and third parts of Proposition 1 do rely on inequity aversion. Analysis similar to that in the appendix but fixing $\alpha = 0$ implies that a level- k model with standard preferences would have L1 types always following recommendations (as they do in our model), so that L2 types would also always follow recommendations (unlike our model), as would L3 and higher types if any existed. Thus without inequity aversion, there would be no difference across games in how often recommendations are followed. Moreover, since ex-ante expected payoffs are higher with than without recommendations, both L1 and L2 players would seek recommendations with certainty, also irrespective of the game.

In order to identify the effect of level- k reasoning on Proposition 1, we need to specify an alternative. The closest rational-expectations analogue to our model is probably one where both players know their own value of α and the population distribution, but not the opponent’s value of α . In equilibrium, each optimises given their own preferences and beliefs about the probability the opponent follows recommendations – meaning that almost all players choose pure strategies – and beliefs turn out to be correct (so that at the *population level*, mixed-strategy play may be seen). Under this assumption, it can be shown that as long as Assumption 2 is satisfied, A recommendations are always followed, and B recommendations are followed by those players with $\alpha \leq n/(M - n)$. This is the same cut-off probability as under level- k reasoning, and therefore we obtain the same comparative statics: B recommendations are followed less often as either payoff asymmetry increases or the cost of flouting recommendations decreases. Hence, at least one way of eliminating the level- k assumption maintains Part 2 of Proposition 1. However, we cannot draw general conclusions about Parts 1 and 3, the solutions of which are intractable unless additional assumptions are made about $F(\cdot)$.

2.2 Hypotheses

Our hypotheses are listed below.

Hypothesis 1 *When recommendations are given, the frequency with which B recommendations are followed decreases as payoff asymmetry increases or as the cost of flouting recommendations (n) decreases.*

Hypothesis 2 *When recommendations are given, the frequency of A choices is higher following an A recommendation than following a B recommendation.*

Hypothesis 3 *When recommendations are voluntary, the frequency with which recommendations are sought decreases as payoff asymmetry increases or as the cost of flouting recommendations decreases.*

These hypotheses follow directly from Proposition 1 (parts 2, 2, and 3 respectively).

3 Experimental procedures

Each session involved 12 subjects playing one game (see Figures 1 and 2 for details), for a total of 27 rounds: 9 each of three different treatments (described below). The subjects in a session were further divided into three 4-person “matching groups” that were closed with respect to interaction, and each subject was matched to each of the three potential opponents three times in each 9-round block.⁸ Subjects were not informed of the matching groups, but were told that in each round, every participant is to be paired with another participant and that they will not be informed of the identity of the person they are paired with in any round. They were also told that typically the pairings will change from round to round, though it is possible to be paired with the same person for two consecutive rounds (as sometimes happened due to the random matching).

In Treatment N, subjects played the baseline game without receiving any outside recommendation and simply chose their action (A or B) simultaneously, after which their payoffs were realised. In Treatment C, prior to playing the game, subjects received a non-binding recommendation of the form “the computer suggests that you play [A or B]”. Subjects were informed in the instructions that whenever one player was recommended to play A, the opponent was recommended to play B, and they were reminded of this each time they received a recommendation. After receiving their recommendations, the players chose their actions and payoffs were realised. In Treatment V, one player in each matching pair was randomly designated as the “decider” and had the choice of whether or not to ask the computer for recommendations. If the decider chose to ask for recommendations, play proceeded as in Treatment C, with both her and her opponent receiving a recommendation before choosing actions; if she chose not to allow recommendations, play proceeded as in Treatment N. End-of-round feedback was the subject’s own action choice, the opponent’s action choice, and the subject’s payoff (there was no recapitulation of recommendations). The order of treatments in a given session was either N–C–V or C–N–V, in order to ensure that subjects have experience both with and without recommendations prior to the V treatment where they are asked to choose whether or not they want them. The disadvantage of using only these two orderings – which obviously do not represent a counter-balancing of treatments – is that one has to be extremely careful in drawing conclusions about differences between the V treatment and either of the other two treatments based on our results. However, our primary interest is not in differences between these treatments but rather in differences across games, making the lack of counter-balancing less of an issue.

The experiment was conducted at the Economics Laboratory of Boğaziçi University in Istanbul, Turkey. Subjects were primarily undergraduate students, invited by emails sent to those who had previously indicated an interest in participating in economics experiments, and who registered online for a session. No-one took part more than once; otherwise, there was no exclusion of subjects. The experiment was run on networked personal computers, and

⁸Such partitioning of sessions into smaller matching groups is common in economics experiments. The advantage of doing so is an increase in the number of independent observations and hence power, though a drawback is the potential for repeated-game behaviour if subjects understand they face each other more than once.

programmed using z-Tree (Fischbacher, 2007). Subjects sat in individual carrels in a single room and were visually isolated from each other, and were asked to turn off their mobile phones and not to communicate with each other except via the computer program.

At the beginning of each block, subjects were given separate written instructions, which were also read aloud in an attempt to make them common knowledge. (See the appendix for English translations of a sample set of instructions, the list of questionnaire questions, and sample screen-shots.) There was no instructions quiz, but subjects were given a chance to ask questions at this point and throughout the session which would be answered privately, after which the first round of that block began. There were no practice rounds.

After the last treatment was finished, subjects completed a questionnaire containing some demographic and attitudinal questions. Once this was finished, subjects were paid privately and individually in cash and then left (no other experiment was conducted during the same session). One of the nine rounds from each of the three treatments was randomly chosen for payment at an exchange rate of one Turkish lira (TRY) per point – that is, payoffs in Figure 1 are in real money units, conditional on that round being chosen – and subjects additionally received a show-up fee of 10 TRY.⁹ Total earnings averaged roughly 22.48 TRY for a session that typically lasted approximately 30 minutes (20–25 minutes for the rounds themselves, and the remainder for logistics such as instructions and payments).

4 Experimental results

A total of 288 subjects participated in the 24 sessions. Some session information is shown in Table 2.

Table 2: Session information (game and ordering of treatments)

Session	(M, n)	Ordering	Session	(M, n)	Ordering	Session	(M, n)	Ordering
1	(10, 7)	N-C-V	9	(10, 1)	N-C-V	17	(7, 1)	N-C-V
2	(10, 4)	C-N-V	10	(10, 4)	N-C-V	18	(13, 1)	N-C-V
3	(10, 4)	N-C-V	11	(10, 4)	C-N-V	19	(13, 1)	N-C-V
4	(10, 1)	N-C-V	12	(10, 7)	C-N-V	20	(7, 1)	N-C-V
5	(10, 1)	C-N-V	13	(13, 1)	C-N-V	21	(7, 1)	C-N-V
6	(10, 7)	N-C-V	14	(13, 7)	N-C-V	22	(13, 7)	C-N-V
7	(10, 1)	C-N-V	15	(13, 7)	N-C-V	23	(7, 1)	C-N-V
8	(10, 7)	C-N-V	16	(13, 7)	C-N-V	24	(13, 1)	C-N-V

4.1 Aggregate choices

Table 3 shows aggregate frequencies of following and seeking recommendations by game, treatment, and recommendation. We introduce some additional notation: Vr for those observations within the V treatment where the decider chooses to seek recommendations (and similarly, Vnr for when recommendations are not sought). The left

⁹Over the time of the experiment, 1 TRY corresponded to between USD 0.30 and 0.45 at market exchange rates. However, the lower cost of living in Turkey compared to many developed countries made the stakes correspondingly higher in real terms. For comparison, the minimum monthly wage in Turkey is 1647 TRY, and a lunch at the school cafeteria costs about 6–10 TRY.

section shows results for the cases without recommendations; we include these for the sake of completeness and do not discuss them further in the main text.¹⁰

Table 3: Descriptive statistics over all rounds: A choices, following and seeking recommendations

Game	A choice (no rec.)		Follow B rec.		Follow A rec.		Seek rec.
	N	Vnr	C	Vr	C	Vr	V
G(7,1)	0.764	0.854	0.579	0.806	0.912	0.977	0.810
G(10,1)	0.803	0.853	0.468	0.497	0.870	0.982	0.764
G(13,1)	0.817	0.840	0.407	0.473	0.917	0.917	0.782
G(10,4)	0.685	0.713	0.671	0.864	0.944	0.975	0.750
G(10,7)	0.590	0.657	0.843	0.950	0.917	0.978	0.838
G(13,7)	0.616	0.636	0.787	0.975	0.870	0.962	0.727

Significant differences? (from Jonckheere tests)

G(7,1)–G(10,1)–G(13,1)	$p \approx 0.095$	$p > 0.20$	$p \approx 0.037$	$p \approx 0.002$	$p > 0.20$	$p \approx 0.044$	$p > 0.20$
G(7,1)–G(10,4)–G(13,7)	$p < 0.001$	$p \approx 0.006$	$p \approx 0.005$	$p < 0.001$	$p > 0.20$	$p > 0.20$	$p > 0.20$
G(10,1)–G(10,4)–G(10,7)	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p \approx 0.14$	$p > 0.20$	$p > 0.20$

Notes: rec. = recommendation; Vr (Vnr) = V treatment, recommendation sought (not sought).

The two centre sections of the table show the frequency with which B and A recommendations are followed. Recommendations to choose A are nearly always followed; there appear to be small differences across games, but these are typically not significant.¹¹ By contrast, there are sharp differences across games in how often B recommendations are followed: from less than half the time in G(10,1) and G(13,1) to most of the time in G(10,7) and G(13,7). The decrease in following recommendations from G(7,1) to G(10,1) to G(13,1), and the increases from G(7,1) to G(10,4) to G(13,7) and from G(10,1) to G(10,4) to G(10,7), are significant (Jonckheere tests, $p \approx 0.037$ and $p \approx 0.002$ for G(7,1)–G(10,1)–G(13,1) in the C and V treatments respectively, $p \approx 0.005$ and $p < 0.001$ for G(7,1)–G(10,4)–G(13,7) in the C and V treatments respectively, $p < 0.001$ for G(10,1)–G(10,4)–G(10,7) in both C and V treatments).¹²

¹⁰Note that when recommendations are not given, the frequency of A choices decreases as payoff asymmetry ($M - n$) decreases (moving from G(7,1) to G(10,1) to G(13,1)), as n increases with ($M - n$) held constant (moving from G(7,1) to G(10,4) to G(13,7)), and as n increases with ($M - n$) decreasing (moving from G(10,1) to G(10,4) to G(10,7)), in both N and Vnr treatments. These patterns are consistent with our behavioural model (see part 1 of Proposition 1), mixed-strategy Nash equilibrium, and other solution concepts like logit equilibrium. We will not analyse the no-recommendations data in detail, given the similarity of these results to those already seen in the literature, and given the difficulty in distinguishing amongst the various models' predictions. Instead, we make three points. First, the qualitative equivalence between our model and these other theories when there are no recommendations is evidence that the assumptions underlying our model are sensible. Second and relatedly, an implication of our behavioural model is that the N-treatment data can be used to put a bound on the fraction of L2 versus L1 subjects in the population; see the appendix for such an analysis. Finally, the good performance of all of these models in characterising the N data is reassuring since it suggests our results match those from other studies implementing treatments like our N treatment.

¹¹All of our non-parametric tests are two-sided and use matching-group-level data, thus erring on the conservative side when drawing conclusions about significance. The Jonckheere tests used here pits the null hypothesis of equality across groups against an alternative that they are ordered (e.g., G(7,1)–G(10,1)–G(13,1) or the reverse). Since our hypotheses specify an ordering across the three games, this test is more appropriate than the Kruskal–Wallis test, for which the alternative hypothesis is simply inequality across games with no order implied. See Siegel and Castellan (1988) for descriptions of the non-parametric tests used in this paper.

¹²These significance results are robust to using only rounds 4–9 of each block (allowing subjects to gain a measure of experience in the

Comparing the two kinds of recommendation, we see that A actions are chosen more often following an A recommendation than after a B recommendation, and this difference is significant, for both C and Vr treatments, for each game individually and for the games pooled together (Wilcoxon signed-ranks test, $p < 0.001$ for each game and for the six games pooled). We thus have:

Result 1 *Subjects are more likely to choose A after receiving an A recommendation than after a B recommendation.*

Result 1 allows us to accept Hypothesis 2.

Despite the differences across games in how often recommendations are followed, the frequency of *seeking* recommendations varied relatively little across games. They were sought 77.9 percent of the time overall, with no significant changes as we move from G(7,1) to G(10,1) to G(13,1), from G(7,1) to G(10,4) to G(13,7) or from G(10,1) to G(10,4) to G(10,7) (Jonckheere test, $p > 0.20$ in all three cases). A regression analysis (details available from the corresponding author upon request) similarly found no systematic differences across games in the frequency of seeking recommendations. So we have:

Result 2 *We do not find systematic differences by game in whether subjects seek recommendations.*

Result 2 allows us to reject Hypothesis 3.

4.2 When do players follow recommendations?

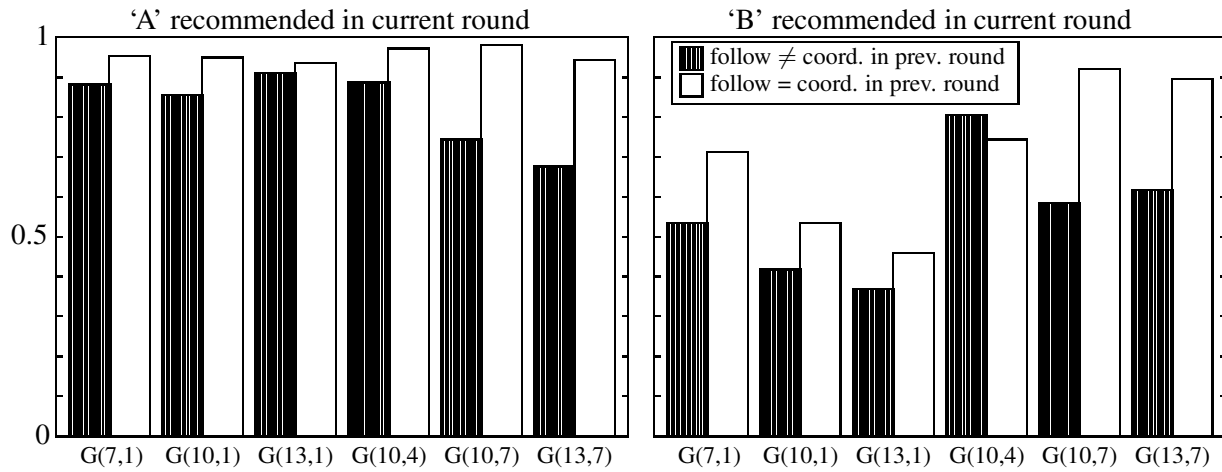
We have already seen suggestive evidence to support Hypothesis 1 in Table 3 and the accompanying non-parametric tests. However, we will see in this section that other factors also influence subjects' decisions to follow recommendations. After an examination of these factors, we will finish the section with regressions to confirm how following recommendations varies with the game after controlling for these other factors (i.e., testing Hypothesis 1).

Figure 3 shows, by game and recommendation received, how likely subjects are to follow their recommendations based on four aspects of the previous-round outcome. Shaded columns comprise two contingencies: (a) the subject followed recommendations but did not coordinate successfully, and (b) the subject did not follow recommendations but nonetheless coordinated successfully. White columns comprise the remaining contingencies: (c) following recommendations and successfully coordinating, or (d) flouting recommendations and mis-coordinating. Thus the shaded columns represent cases where following recommendations was a relatively poor tactic, and the white columns where following recommendations was a relatively good tactic – irrespective of whether or not the subject actually did follow recommendations. In eleven of the twelve game–recommendation pairs, subjects are more likely to follow a recommendation when following recommendations had been (or would have been) successful than otherwise. The difference is often small in the case of A recommendations, but more substantial for B recommendations: up to about 30 percentage points.

Figure 4 shows that decider status has little bearing on whether subjects follow recommendations in the V treatment. This figure shows scatter-plots with the frequency of A choices in the C treatment (where there are no deciders) on the horizontal axis and the frequencies of A choices for deciders and non-deciders in the V treatment – conditional on recommendations being sought – on the vertical axis, disaggregated by game and recommendation. Recall that the frequency of an A choice is exactly the frequency of following an A recommendation, and is the frequency of flouting a B recommendation (i.e., one minus the frequency of following it).

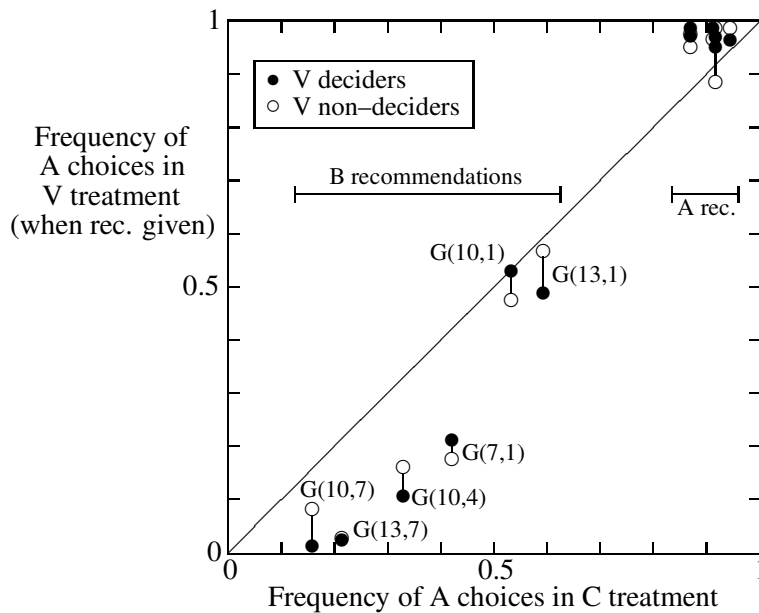
game and treatment) instead of all of the data. Details are available from the corresponding author upon request.

Figure 3: Frequency of followed recommendations by previous-round outcome (rounds 2–9 of C and V treatments, recommendations received in previous round)



The cluster of plotted points in the top-right corner of the figure reflects our previous finding that A recommendations are followed most of the time. There is little systematic difference between deciders and non-deciders, with deciders following recommendations more often than non-deciders in four of the games and less often in the other two. After B recommendations, there is similarly little systematic difference between deciders and non-deciders.

Figure 4: Frequency of followed recommendations by decider status (all rounds)



Recommendations are followed more often by deciders in four of the six games, but the differences are small and none is statistically significant (Wilcoxon signed-ranks test, matching-group-level data, $p > 0.20$ in all cases). The lack of differences between deciders and non-deciders suggests that “ownership” of a recommendation – due to being the player who sought it – has little influence on whether it is followed.

The comparison between C and V treatments is also illustrated in Figure 4. For A recommendations, nearly all of the plotted points are above the 45-degree line, indicating that recommendations are followed more often in the V treatment than in the C treatment. For B recommendations, nearly all of the points are *below* the 45-degree line, again indicating that recommendations are followed more in the V treatment. The apparent difference between C and V treatments suggests that subjects may be more predisposed toward following recommendations that were voluntary than those that were imposed. However, since the C treatment always preceded the V treatment in a session, the difference could also be due to subjects being more experienced by the time they reach the V treatment.

We move on to regressions. Table 4 shows results from four regressions with *following the recommendation* as the dependent variable. Our main independent variables are indicators for five of the six games (with G(7,1) as the baseline), an indicator for an A recommendation in the current round, and the products of this indicator with the game dummies. To allow for time-varying effects, we include the round number and its interactions with the game dummies, the A-recommendation dummy, and their product. Additional explanatory variables are indicators for decider and non-decider in the V treatment (so that the C treatment is the baseline), their interactions with the A-recommendation dummy, and an indicator for the C–N–V ordering of treatments (so that the baseline is N–C–V). Model 1 uses only these variables plus a constant term. Model 2 additionally includes a dummy whose value is one if the subject both followed recommendations and successfully coordinated in the previous round, or did neither (i.e., the contingencies represented by the white columns in Figure 3). Also included are this variable’s interactions with the games, the current recommendation, and their product. Models 3 and 4 correspond to Models 1 and 2 respectively, but additionally include the demographic and attitudinal variables we collected.¹³ All of our regressions were panel linear models, estimated using Stata v. 12, with subject random effects and with standard errors clustered by the matching group. (Panel probits without clustering produced similar results; for details contact the corresponding author.) Models 1 and 3 were estimated over the observations in which a recommendation was received, while Models 2 and 4 were estimated over the subset of this set in which a recommendation was received in the previous round as well (in particular, leaving out the first round of each nine-round block).

Table 4 shows the marginal effects of the game dummies conditional on either a B or an A recommendation, and for the case of B recommendations, p -values for the three three-game comparisons relevant to Hypothesis 1.¹⁴ According to the estimated marginals, subjects are less likely to follow B recommendations in G(10,1) than in G(7,1), and less likely still in G(13,1); a joint-significance test rejects the null of equality across the three games at the 1% level. Subjects are more likely to follow B recommendations as we move from G(7,1) to G(10,4) and thence to G(13,7), and similarly as we move from G(10,1) to G(10,4) and thence to G(10,7) ($p < 0.01$ for both three-way

¹³The variables we include are: indicator for female, age in years, number of older siblings, number of younger siblings, indicator for only child, indicators for economics student and non-economics business student, Likert response to a general risk question (“How willing are you to take risks in general?”), and mean Likert response to a set of DOSPERT risk questions (Weber et al., 2002). We include these variables as controls, rather than out of interest in their effects; hence we leave them out of our tables to save space.

¹⁴For completeness, the table also displays average marginal effects for some of the other explanatory variables. These effects are not directly connected to our hypotheses, but some are interesting in their own right. For example, subjects are significantly more likely to follow A recommendations than B recommendations, and they are significantly more likely to follow recommendations when the previous recommendation led to successful coordination. The effect of the C–N–V variable suggests that subjects were more likely to follow recommendations if their initial experience was with no access to recommendations. The decider and non-decider dummies’ effects imply more following of recommendations in the V treatment compared to the C treatment. However, their marginals are not significantly different from each other, suggesting that the source of recommendations (chosen by oneself or by the opponent) has little impact on whether they are followed. Finally, the round number is positive and significant, indicating that subjects become more willing to follow recommendations with experience, even after controlling for the previous-round outcome. However, the size of the effect (an increase of about 6 percentage points over the nine rounds) is small.

Table 4: Regression results, dependent variable = *following* recommendation (standard errors in parentheses)

	[1]	[2]	[3]	[4]
<i>Marginal effects given B recommendation</i>				
G(10,1)	-0.199** (0.083)	-0.152* (0.081)	-0.189** (0.080)	-0.146* (0.079)
G(13,1)	-0.252*** (0.050)	-0.209*** (0.070)	-0.249*** (0.048)	-0.209*** (0.067)
G(10,4)	0.075** (0.032)	0.082*** (0.029)	0.059* (0.031)	0.068** (0.029)
G(10,7)	0.202*** (0.057)	0.197*** (0.064)	0.193*** (0.056)	0.192*** (0.062)
G(13,7)	0.180*** (0.025)	0.153*** (0.023)	0.172*** (0.023)	0.147*** (0.022)
<i>Joint significance:</i>				
G(7,1)=G(10,1)=G(13,1)	$p < 0.001$	$p \approx 0.002$	$p < 0.001$	$p \approx 0.002$
G(7,1)=G(10,4)=G(13,7)	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
G(10,1)=G(10,4)=G(10,7)	$p < 0.001$	$p \approx 0.003$	$p < 0.001$	$p \approx 0.003$
<i>Marginal effects given A recommendation</i>				
G(10,1)	-0.027 (0.025)	-0.026 (0.023)	-0.019 (0.025)	-0.023 (0.023)
G(13,1)	-0.015 (0.036)	-0.027 (0.041)	-0.011 (0.040)	-0.025 (0.044)
G(10,4)	0.019 (0.025)	0.029 (0.023)	0.003 (0.027)	0.014 (0.024)
G(10,7)	0.001 (0.050)	0.001 (0.042)	-0.010 (0.051)	-0.009 (0.045)
G(13,7)	-0.031 (0.042)	-0.055 (0.046)	-0.040 (0.043)	-0.062 (0.045)
<i>Average marginal effects</i>				
'A' recommendation	0.246*** (0.020)	0.233*** (0.023)	0.245*** (0.020)	0.232*** (0.022)
Round number	0.006*** (0.002)	0.006*** (0.002)	0.006*** (0.002)	0.006*** (0.002)
Decider (V treatment)	0.089*** (0.013)	0.091*** (0.012)	0.088*** (0.013)	0.091*** (0.012)
Non-decider (V treatment)	0.080*** (0.014)	0.080*** (0.014)	0.080*** (0.014)	0.081*** (0.015)
(recommendation + coordination) or (no rec. + no coord.) in prev. round		0.085*** (0.017)		0.083*** (0.016)
C-N-V ordering	-0.045* (0.026)	-0.044* (0.022)	-0.054*** (0.022)	-0.048** (0.019)
Demographic variables?	No	No	Yes	Yes
Round x game interactions?	Yes	Yes	Yes	Yes
N	4610	3748	4610	3748
R^2	0.220	0.262	0.230	0.269

* (**,***): Marginal effect significantly different from zero at the 10% (5%, 1%) level.

comparisons). Thus, subjects are more likely to follow B recommendations as payoff asymmetry decreases, as the cost of flouting recommendations increases, or both.¹⁵

Result 3 *The frequency with which B recommendations are followed decreases as payoff asymmetry ($M - n$) increases or as the cost of flouting recommendations (n) decreases.*

Result 3 allows us to accept Hypothesis 1.

5 Summary and discussion

Our main result (Result 3 in Section 4.2) is that when subjects receive recommendations, they are less likely to follow them as payoff asymmetry increases or as the cost of flouting recommendations decreases. This does not indicate a broad, across-the-board failure of subjects to understand how to solve coordination problems. When subjects receive a “favourable” A recommendation (meaning that if recommendations are followed, this subject would receive a *higher* money payoff than the opponent), they are overwhelmingly likely to obey it irrespective of the game and whether recommendations were sought or imposed. Subjects tend to flout recommendations, instead, when they receive “unfavourable” B recommendations (which would lead to receiving a *lower* money payoff than the opponent if recommendations are followed).

Our Result 3 is in the spirit of findings like those of Crawford et al. (2008), who investigate coordination games without recommendations, in which actions are labelled so that the pure-strategy equilibria have both players choosing the same action. They find that when payoffs are symmetric, subjects are usually able to coordinate successfully, but this success nearly disappears when even a small amount of payoff asymmetry is introduced.¹⁶ They attribute this result to the interplay between *label salience* (actions are focal because of their names) and *payoff salience* (actions are focal because of the payoffs they yield).¹⁷ In their symmetric “*SL*” game (see Figure 5), both

Figure 5: Two games from Crawford et al. (2008)

		Player 2				Player 2	
		X	Y			X	Y
Player 1	X	5, 5	0, 0	Player 1	X	5, 5.1	0, 0
	Y	0, 0	5, 5		Y	0, 0	5.1, 5
Game <i>SL</i>				Game <i>ASL</i>			

pure-strategy equilibria are payoff salient, so payoff salience has no focal ability for either player, and hence they coordinate based on label salience: they recognise that action X is label salient whereas action Y is not, and so both choose X. By contrast, in their (slightly) asymmetric “*ASL*” game, Y is payoff salient for Player 1 and X is payoff salient for Player 2, and since X is still label salient for both players, label salience reinforces payoff salience for one player but opposes it for the other. So, neither action is focal, making coordination difficult.

¹⁵These conclusions are robust to estimating marginals in the last round of a block instead of using the average marginal over all rounds (results available from the corresponding author upon request).

¹⁶The importance of payoff symmetry in coordination has also been observed in recent experimental studies by López-Pérez et al. (2015) and Parravano and Poulsen (2015).

¹⁷See Dugar and Shahriar (2012) for another recent study of label salience in coordination games.

The decreasing usefulness of recommendations as payoff asymmetry increases in our experiment is at once a weaker result and a stronger result than Crawford et al.’s. Our result is *weaker* in the sense that mis-coordination in our setting does not arise as soon as any slight payoff asymmetry does: subjects typically coordinate successfully under the moderate level of payoff asymmetry present in our G(10,7) game, and are still fairly successful even in G(13,7) and G(10,4) (which correspond roughly to Crawford et al.’s “large asymmetry” game). However, our result is *stronger* in that by providing recommendations to subjects, we stack the deck in favour of coordination by “super-charging” label salience: rather than relying on a shared interpretation of linguistic cues to determine which action is label salient, giving recommendations that form pure-strategy equilibria makes that determination explicit and manifest. The fact that coordination nevertheless frequently fails under high payoff asymmetry suggests that this failure is not due to subjects failing to recognise the obvious focal point, but rather to some of them not viewing it as a focal point – perhaps because their preferences contain an other-regarding component.¹⁸

This result also complements previous experimental studies of recommendations. Duffy and Feltovich (2010) report that recommendations are typically followed when (a) they form a correlated equilibrium, and (b) the correlated equilibrium payoff-dominates the mixed-strategy Nash equilibrium, while Bone et al. (2013) report that they tend to be followed if (c) the correlated equilibrium is not in turn payoff-dominated by some other outcome. Our results go still further. These three conditions (satisfied in all of our games) are not sufficient; it is further necessary for the correlated equilibrium either to be fairly (ex post) payoff-equitable or for the cost of unilaterally disobeying recommendations to be sufficiently high. In that, our results are in the spirit of those of Georgalos et al. (2018), who investigate the related notion of *coarse correlated equilibrium* (Moulin and Vial, 1978), a further generalisation of correlated equilibrium. (See Moulin et al., 2014, for a related theoretical analysis.) For our purposes, we can think of coarse correlated equilibrium as a distribution of recommendations that players are willing not only to seek, but also to be bound by their realisation: thus like our V treatment but (a) with each player independently deciding whether to seek recommendations only for herself, and (b) without the option to flout recommendations once they are given. Georgalos et al. find that subjects typically avoid those recommendations, and conjecture that subjects dislike recommendations leading to unfair outcomes even if the recommendations themselves are procedurally fair.

As noted in Section 2.1, our behavioural model was intended as a simple implementation of inequity aversion, in order to understand how inequity-averse players’ behaviour would vary across games (see that section for details). Thus, while Result 3 broadly supports our model, we view this as a success for inequity aversion – combined with level- k reasoning – broadly, rather than for the specific way these was implemented. However, for readers interested in model building, we point out that several of our secondary results may be useful in providing additional detail about how a behavioural model should look. First, the lack of systematic differences we observe based on where recommendations come from (deciders in the V treatment versus non-deciders in the V treatment versus everyone in the C treatment) suggests that attribution (Charness, 2004) did not have a large impact in our experiment, and hence extending our model to account for intentions (Rabin, 1993; Dufwenberg and Kirchsteiger, 2004) is not necessary.¹⁹ Second, the lack of evidence supporting our model’s predictions for *seeking* recommendations suggests that in Saito’s (2013) framework, aversion to ex-ante disadvantageous inequity swamps aversion to ex-

¹⁸Crawford et al. focus on bounded rationality rather than other-regarding preferences as an explanation for their results. This is probably appropriate, since it seems implausible that other-regarding preferences could explain the results from games with slight payoff asymmetry like their Game *ASL* (Figure 5), and they do not find substantial further drops in coordination from these games to those with more asymmetry.

¹⁹To be fair, however, we should point out that because our objective was not to test attribution, our experiment was not well-designed for such a test (e.g., since recommendations are ex-ante fair, non-deciders have little to gain, even psychologically, from punishing a decider who seeks recommendations). Given this, our results would seem to have limited bearing on whether attribution is important in other settings.

post disadvantageous inequity (i.e., δ is approximately 1). Given our results for *following* recommendations, one interpretation of this might be that subjects in the experiment – when making a decision to seek recommendations – fail to foresee the distaste they may feel once the recommendations are realised. Third, the evidence we find for previous-round results influencing both seeking and following of recommendations suggests that our static theory may be incomplete, with some explicit modelling of dynamic behaviour called for. Such an extension is beyond the scope of the current paper, but may be a promising avenue for future research.

We close with some thoughts about what our results imply for the applicability of correlated equilibrium. We are reluctant to view our experiment as a test of correlated equilibrium as a solution concept (indeed, a comprehensive test would require a wider range of correlated equilibria than our 50–50 randomisation between two pure Nash equilibria); rather, we consider it to be a test of the usefulness of recommendations. This is because the usual caveat about experimental tests of game theory – regarding the distinction between monetary payments and game-theoretic payoffs – applies here. As noted above, a subject in the experiment who chooses action A upon receiving a B recommendation, anticipating that this will lead to zero monetary payments for him and his opponent instead of a positive payment for himself and a larger payment for his opponent, is not necessarily violating any principle of game theory; he may merely be showing that his utility is not identical to his money earnings. At most, we could conclude that the combination of correlated equilibrium and self-regarding preferences does not describe behaviour well as payoff asymmetry increases.

References

- Benjamini, Y. and Y. Hochberg (1995), “Controlling the false discovery rate: a practical and powerful approach to multiple testing”, *Journal of the Royal Statistical Society, Series B (Methodological)* 57, pp. 289–300.
- Bolton, G.E. and A. Ockenfels (2000), “ERC: a theory of equity, reciprocity, and competition”, *American Economic Review* 90, pp. 166–193.
- Bone, J., M. Drouvelis and I. Ray (2013), “Coordination in 2x2 games by following recommendations from correlated equilibria”, University of Birmingham Economics Discussion Paper 13–04 (12–04R).
- Bornstein, G., Gneezy, U. and R. Nagel (2002), “The effect of intergroup competition on group coordination: an experimental study”, *Games and Economic Behavior* 41, pp. 1–25.
- Brandts, J. and D.J. Cooper (2006), “A change would do you good an experimental study on how to overcome coordination failure in organizations”, *American Economic Review* 96, pp. 669–693.
- Camerer, C.F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton, NJ.
- Cason, T.N. and T. Sharma (2007), “Recommended play and correlated equilibria”, *Economic Theory* 33, pp. 11–27.
- Charness, G. (2004), “Attribution and reciprocity in an experimental labor market”, *Journal of Labor Economics* 22, pp. 665–688.
- Crawford, V.P. (1998), “A survey of experiments on communication via cheap talk”, *Journal of Economic Theory* 78, pp. 286–298.

- Crawford, V.P., U. Gneezy and Y. Rottenstreich (2008), “The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures”, *American Economic Review* 98, pp. 1443–1458.
- Duffy, J. and N. Feltovich (2010), “Correlated equilibria, good and bad: an experimental study”, *International Economic Review* 51, pp. 701–721.
- Duffy, J. and E. Fisher (2005), “Sunspots in the laboratory”, *American Economic Review* 95, pp. 510–529.
- Dufwenberg, M. and G. Kirchsteiger (2004), “A theory of sequential reciprocity”, *Games and Economic Behavior* 47, pp. 268–298.
- Dugar, S. and Q. Shahriar (2012), “Focal points and economic efficiency: the role of relative label salience”, *Southern Economic Journal* 78, pp. 954–975.
- Farrell, J. (1987), “Cheap talk, coordination, and entry”, *RAND Journal of Economics* 18, pp. 34–39.
- Fehr, E., and K.M. Schmidt (1999), “A theory of fairness, competition and cooperation”, *Quarterly Journal of Economics* 114, pp. 817–868.
- Fischbacher, U. (2007), “z-Tree: Zurich toolbox for ready-made economic experiments”, *Experimental Economics* 10, pp. 171–178.
- Georgalos, K., I. Ray and S. Sen Gupta (2018), “Coordination vs. (coarse) correlation”, working paper.
- Kurz, V., A. Orland and K. Posadzy (2017), “Fairness versus efficiency: how procedural fairness concerns affect coordination”, forthcoming, *Experimental Economics*.
- López-Pérez, R., Á. Pintér and H.H. Kiss (2015) “Does payoff equity facilitate coordination? A test of Schelling’s conjecture”, *Journal of Economic Behavior & Organization* 117, pp. 209–222.
- McKelvey, R.D. and T.R. Palfrey (1995), “Quantal response equilibria for normal form games”, *Games and Economic Behavior* 10, pp. 6–38.
- Mehta, J., Starmer, C. and Sugden, R. (1994), “Focal points in pure coordination games: an experimental investigation”, *Theory and Decision* 36, pp. 163–185.
- Moreno, D. and J. Wooders (1998), “An experimental study of communication and coordination in noncooperative games”, *Games and Economic Behavior* 24, pp. 47–76.
- Moulin, H., I. Ray and S. Sen Gupta (2014), “Improving Nash by coarse correlation”, *Journal of Economic Theory* 150, pp. 852–865.
- Moulin, H. and J.P. Vial (1978), “Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon”, *International Journal of Game Theory* 7, pp. 201–221.
- Nagel, R.C. (1995), “Unraveling in guessing games: an experimental study”, *American Economic Review* 85, pp. 1313–1326.
- Ochs, J. (1995), “Coordination games”, in *Handbook of Experimental Economics*, J.H. Kagel and A.E. Roth, eds., Princeton University Press, Princeton, NJ., pp. 195–252.

- Parravano, M. and O. Poulsen (2015), “Stake size and the power of focal points in coordination games: experimental evidence”, *Games and Economic Behavior* 94, pp. 191–199.
- Rabin, M. (1993), “Incorporating fairness into game theory and economics”, *American Economic Review* 83, pp. 1281–1302.
- Roth, A.E. (1995), “Bargaining experiments”, in *Handbook of Experimental Economics*, J.H. Kagel and A.E. Roth, eds., Princeton University Press, Princeton, NJ, pp. 253–348.
- Saito, K. (2013), “Social preferences under risk: equality of opportunity versus equality of outcome”, *American Economic Review* 103, pp. 3084–3101.
- Schelling, T. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.
- Stahl, D.O., II and P.W. Wilson (1994), “Experimental evidence on players’ models of other players”, *Journal of Economic Behavior & Organization* 25, pp. 309–327.
- Stahl, D.O., II and P.W. Wilson (1995), “On players’ models of other players: theory and experimental evidence”, *Games and Economic Behavior* 10, pp. 218–254.
- Van Huyck, J.B., R.C. Battalio and R.O. Beil (1991), “Strategic uncertainty, equilibrium selection, and coordination failure in average–opinion games,” *Quarterly Journal of Economics* 106, pp. 885–910.
- Weber, E.U., A.-R. Blais and N. Betz (2002), “A domain–specific risk–attitude scale: measuring risk perceptions and risk behaviors”, *Journal of Behavioral Decision Making* 15, pp. 263–290.

A Proof of Proposition 1

Let $r_i \in \{A, B\}$ denote the recommendation for player i and $s_i \in \{A, B\}$ denote i 's action choice. When there is no recommendation we use $r_i = \emptyset$. By construction, if $r_i = A$ then $r_j = B$, if $r_i = B$ then $r_j = A$, and if $r_i = \emptyset$, $r_j = \emptyset$.

(1) When player i is of type $L1$ and $r_i = \emptyset$, she believes that player j is of type $L0$ playing A and B with equal probability. Therefore, the expected utility for i is $M/2$ when $s_i = A$ and $n/2 - \alpha(M/2 - n/2)$ when $s_i = B$. Since $n < M$, $s_i = A$ whenever player i is of type $L1$.

When player i is of type $L2$ and $r_i = \emptyset$, the expected payoff is 0 when $s_i = A$, since i believes that her $L1$ opponent will choose $s_j = A$. On the other hand, the expected payoff is $n - \alpha(M - n)$ when $s_i = B$. Based on this, $s_i = B$ if and only if $n/(M - n) > \alpha$, otherwise $s_i = A$. The fraction of $L2$ players for whom $s_i = B$ is given by $F(n/(M - n))$ which is increasing in n and decreasing in M since $F(\cdot)$ is an increasing function over its domain by Assumption 2.

(2) Suppose that player i is of type $L1$ and receives a recommendation. She believes that player j is of type $L0$, so that $s_j = A$ when $r_j = A$ and $s_j = B$ when $r_j = B$. Now suppose that $r_i = A$, which also means that $r_j = B$ and $s_j = B$. Hence, the payoff is $M/2$ when $s_i = A$ and 0 when $s_i = B$, so i will follow the recommendation.

Next suppose that i is of type $L1$ and $r_i = B$. She believes that $s_j = A$ since $r_j = A$. Consequently, the expected utility is 0 when $s_i = A$ and it is $n - \alpha(M - n)$ when $s_i = B$. In this case, player i will follow the recommendation only if $n/(M - n) > \alpha$. To summarise, when i is of type $L1$, $s_i = r_i$ when either $r_i = A$ or both $r_i = B$ and $n/(M - n) > \alpha$ (i.e., α is sufficiently small). Given that both types of recommendations are equally likely, the frequency with which the recommendation is followed among $L1$ types is $0.5[1 + F(n/(M - n))]$ which increases as M decreases and n increases.

Now suppose that player i is of type $L2$ and $r_i = A$. In this case, i expects that $s_j = A$ with probability $1 - F(n/(M - n))$ and $s_j = B$ with probability $F(n/(M - n))$. Her expected payoff is $M \cdot F(n/(M - n))$ when $s_i = A$ and $[1 - F(n/(M - n))][n - \alpha(M - n)]$ when $s_i = B$. Based on this, expected payoff from $s_j = A$ is larger if and only if $F(n/(M - n)) > [n - \alpha(M - n)]/[M(1 - \alpha) + n(1 - \alpha)]$. Since $\alpha > 0$ and $M > n > 0$, the right hand side is less than n/M , which is itself less than $F(n/(M - n))$ by assumption 2. Therefore, $s_i = A$ when $r_i = A$.

Finally, suppose that player i is of type $L2$ and $r_i = B$. Player i also believes that $s_j = A$ since $r_j = A$. In this case, $s_i = B$ if and only if $n/(M - n) > \alpha$, otherwise $s_i = A$. Overall, the recommendation is followed among $L2$ types with a frequency of $0.5[1 + F(n/(M - n))]$ which increases as n increases and M decreases.

(3) **Case 1 (Player i is of type $L1$):** When player i is of type $L1$, $s_i = A$ in case $r_i = \emptyset$ and so her expected payoff is $M/2$ if she does not ask for recommendations. If she asks for recommendations, $s_i = A$ when $r_i = A$ or when $r_i = B$ and $n/(M - n) < \alpha$. On the other hand, $s_i = B$ when $r_i = B$ and $n/(M - n) > \alpha$.

Case 1a ($n/(M - n) < \alpha$): Here, $s_i = A$ and player i believes that player j is of type $L0$ and follows the recommendation. Since both types of recommendations are equally likely, the expected payoff to player i from asking for recommendations is $M/2$. Therefore i is indifferent between asking for recommendations or not.

Case 1b ($n/(M - n) > \alpha$): Here, $s_i = A$ when $r_i = A$ and $s_i = B$ when $r_i = B$. Based on expectations about player j and the likelihood of recommendations, the expected payoff to player i from asking for recommendations is given by

$$\delta(M/2 + n/2) + (1 - \delta)(0.5)[M + n - \alpha(M - n)],$$

which is larger than $M/2$ (the payoff from not seeking recommendations) if $n/(M - n) > \alpha(1 - \delta)$, which is guaranteed by the condition for this case ($n/(M - n) > \alpha$). So, i strictly prefers to seek recommendations.

Thus, player i is only willing to *not* ask for recommendations when $n/(M - n) < \alpha$: that is, when $n < M\alpha/(1 + \alpha)$. So, as n increases, the fraction of the population strictly preferring to ask for recommendations increases.

Case 2 (Player i is of type $L2$): When player i is of type $L2$ and $r_i = \emptyset$, she expects that $s_j = A$. In this case $s_i = B$ whenever $n/(M - n) > \alpha$, and $s_i = A$ otherwise. Her expected payoff from not asking for recommendations is $n - \alpha(M - n)$ when $n/(M - n) > \alpha$ and it is 0 when $n/(M - n) < \alpha$.

Case 2a ($n/(M - n) < \alpha$): If she asks for recommendations, $s_i = A$ both when $r_i = A$ and when $r_i = B$. Based on beliefs about player j , the expected payoff from asking for recommendations is $M \cdot F(n/(M - n))$, which is greater than 0, implying that this player will always ask for recommendations.

Case 2b ($n/(M - n) > \alpha$): Here, $s_i = A$ when $r_i = A$ and $s_i = B$ when $r_i = B$. Based on beliefs about player j , the expected payoff from asking for recommendations is given by

$$\{0.5[g(M, n)M + n] - \alpha[M/2 + g(M, n)n/2 - (M/2)g(M, n) - n/2]\}\delta + 0.5 \cdot [(M/2)g(M, n) + n - \alpha(M - n)](1 - \delta)$$

where $g(M, n) = F(n/(M - n))$.

The term in front of δ can be written as $g(M, n)M/2 + n/2 - \alpha(0.5)[1 - g(M, n)](M - n)$. By Assumption 2 and the fact that $g(M, n) \leq 1$, this term is greater than $n - \alpha(M - n)$. Also, the term in front of $(1 - \delta)$ is greater than $n - \alpha(M - n)$ by Assumption 2 and the fact that $\alpha > 0$. So, this whole expression is greater than $n - \alpha(M - n)$ (since $\delta \leq 1$), implying that the player will always ask for recommendations.

B English translation of experimental instructions – NCV ordering

[Part 1]

Welcome!

Thank you for your participation. The aim of this study is to understand how people make decisions in certain situations. From now on, talking to each other is prohibited. Please also turn off your mobile phones. If you have a question please raise your hand. We will come to you and answer your question. Please do not hesitate to ask questions, since it is very important that all participants understand the rules in this study.

The experiment will be conducted on the computer and you will make all your decisions using the computer. You will earn a monetary reward in the game you will play during the experiment. The amount you will earn depends on your decision and the decisions of other participants. This amount and the participation fee will be paid to you in cash at the end of the experiment. The experiment consists of 3 parts.

We now start describing **Part 1**:

The Game:

This part of the experiment will last for 9 rounds and every period you will be matched with another participant and play the game you see in the table below. For a given round, there's no other player in the game except you and the player you are matched with.

Let's introduce the game in more detail. Please do not hesitate to ask questions if you have any.

	The other player chooses A	The other player chooses B
You choose A	You earn: 0 TL The other player earns: 0 TL	You earn: $[M]$ TL The other player earns: $[n]$ TL
You choose B	You earn: $[n]$ TL The other player earns: $[M]$ TL	You earn: 0 TL The other player earns: 0 TL

You and the other player will choose one among options A and B simultaneously. You will make these choices separately and without the knowledge of the other person. Your earnings will depend on your choice and the choice of the other person.

If you both choose A, both of you will earn 0 TL.

If you choose A and the other person chooses B, you will earn $[M]$ TL and the other player earns $[n]$ TL.

If you both choose B, both of you will earn 0 TL.

If you choose B and the other player chooses A, you will earn $[n]$ TL and the other player earns $[M]$ TL.

Let's describe the matchings in the game.

Matchings:

In every round you will be matched with another participant. You will never learn the identity of the person that you are matched with. Similarly, the other participant will not learn your identity. Although it is possible to be matched

with the same person in two consecutive rounds, the person you are matched with will be different than the previous round, in general.

Earnings:

We will randomly choose 1 round among 9 rounds of this part, and your earnings in that round will be your actual earnings from this part of the experiment.

Your total earnings are, “your earnings from part 1” + “your earnings from part 2” + “your earnings from part 3” + “show-up fee”.

[Part 2]

We now start describing **Part 2**:

The Game:

This part of the experiment will also last for 9 rounds and every period you will be matched with another participant and play the game you see in the table below:

	The other player chooses A	The other player chooses B
You choose A	You earn: 0 TL The other player earns: 0 TL	You earn: $[M]$ TL The other player earns: $[n]$ TL
You choose B	You earn: $[n]$ TL The other player earns: $[M]$ TL	You earn: 0 TL The other player earns: 0 TL

But, different from Part 1, before you and the other player make your choices, the system will make a suggestion to both of you.

With probability 50%, the system will suggest you to choose A and suggest the other player to choose B. Or with probability 50%, the system will suggest you to choose B and suggest the other player to choose A. In other words, the system will suggest one of you to choose A and the other one to choose B. As you can see, it is equally likely that the system will suggest you to choose A or B.

Matchings:

In every round you will be matched with another participant. You will never learn the identity of the person that you are matched with. Similarly, the other participant will not learn your identity. Although it is possible to be matched with the same person in two consecutive rounds, the person you are matched with will be different than the previous round, in general.

Earnings:

We will randomly choose 1 round among 9 rounds of this part, and your earnings in that round will be your actual earnings from this part of the experiment.

Your total earnings are, “your earnings from part 1” + “your earnings from part 2” + “your earnings from part 3” + “show-up fee”.

[Part 3]

We now start describing **Part 3**:

The Game:

This part of the experiment will also last for 9 rounds and every period you will be matched with another participant and play the game you see in the table below:

	The other player chooses A	The other player chooses B
You choose A	You earn: 0 TL The other player earns: 0 TL	You earn: $[M]$ TL The other player earns: $[n]$ TL
You choose B	You earn: $[n]$ TL The other player earns: $[M]$ TL	You earn: 0 TL The other player earns: 0 TL

But, different from the first two parts, before you and the other player make your choices, the system will choose one of you randomly. At this point, it is equally likely that you or the other player is chosen.

The chosen player will determine whether to allow the system to make a suggestion or not.

If this player allows the system to make a suggestion, the system will make a suggestion as in the previous part. The way suggestions are determined the likelihood of each session is the same as the previous part. If this player doesn't allow the system to make a suggestion, the system will not make a suggestion.

Matchings:

In every round you will be matched with another participant. You will never learn the identity of the person that you are matched with. Similarly, the other participant will not learn your identity. Although it is possible to be matched with the same person in two consecutive rounds, the person you are matched with will be different than the previous round, in general.

Earnings:

We will randomly choose 1 round among 9 rounds of this part, and your earnings in that round will be your actual earnings from this part of the experiment.

Your total earnings are, “your earnings from part 1” + “your earnings from part 2” + “your earnings from part 3” + “show-up fee”.

C Questionnaire questions (English translation)

How willing are you to take risks in general? (0 lowest – 10 highest)

Attitudinal questions (all answered on a scale from 1=absolutely wrong to 6=absolutely right):

q1: I wouldn't hesitate to ask my boss for a pay rise.

q2: I might invest in risky assets.

q3: Lying is difficult for me.

q4: I could choose someone as a roommate without knowing that person well.

q5: I would vouch for the debt of a friend.

q6: I could argue with a friend who has a very different opinion on an issue.

q7: If I lose my wallet, I believe that the person who finds it will bring it to me with the things in it.

q8: I trust my friends about money issues.

q9: I can spend money without thinking about the consequences.

q10: I can admit it easily when my tastes are different from my friends'.

q11: I wouldn't hesitate to move to a different city.

q12: If I would have an unexpected money windfall, the first thing I would do would be to share with people I know.

q13: I could take a job at which I will get paid on commission only.

q14: If I were rich enough, I would lend high amounts of money to my friends.

q15: My verdicts about others' trustworthiness generally turn out to be correct.

Demographic questions:

Age: in years (integer valued).

Sex: 1=female, 0=male.

Living: living arrangement for the subject (0=student housing, 1=with family, 2= with friends, 3=alone).

Siblings: number of siblings of subject.

Older siblings: number of siblings who are older than the subject.

Major: 2=economics, 1=other business, 0=other.

Econ: number of economics classes (0, 1, 2, 3, 4+).

D Additional data analysis

D.1 Correlation, coordination and inequity

In this section, we examine how the differences we have seen in how subjects react to recommendations are reflected in other outcome variables. First, we compare the effects of recommendations on correlation between row and column player actions. Then we investigate variables associated with this correlation: coordination, payoff efficiency and payoff inequity.

Row and column player action choices are perfectly uncorrelated in a mixed–strategy Nash equilibrium, and perfectly negatively correlated in the correlated equilibrium. Table 5 shows the aggregate correlation for each game and each treatment, and within the V treatment, according to whether recommendations were sought. Also shown are p -values from significance tests across games (Jonckheere tests), and within each game, between observations with and without recommendations received (Wilcoxon rank–sum tests). Correlations tend to be closer to zero without

Table 5: Correlation between row and column player action choices

	N treatment	V, rec. not sought	C treatment	V, rec. sought	Signif. diff.? (pooled C and Vr. vs. pooled N and Vn)
G(7,1)	−0.05	+0.22	−0.29	−0.61	$p < 0.001$
G(10,1)	−0.04	−0.17	−0.20	−0.28	$p \approx 0.15$
G(13,1)	−0.07	−0.03	−0.15	−0.30	$p \approx 0.15$
G(10,4)	−0.01	−0.31	−0.41	−0.70	$p < 0.001$
G(13,7)	−0.21	−0.57	−0.53	−0.90	$p \approx 0.003$
G(10,7)	−0.22	−0.14	−0.69	−0.90	$p \approx 0.034$
<i>Significance of differences</i>					
G(7,1)–G(10,1)–G(13,1)	$p > 0.20$	$p > 0.20$	$p > 0.20$	$p \approx 0.004$	
G(7,1)–G(10,4)–G(13,7)	$p > 0.20$	$p \approx 0.08$	$p \approx 0.016$	$p \approx 0.004$	
G(10,1)–G(10,4)–G(10,7)	$p \approx 0.18$	$p \approx 0.04$	$p < 0.001$	$p < 0.001$	


* (**,***): *Significance across games based on Jonckheere test, and within games based on Wilcoxon signed–ranks test (both using matching–group–level data).*

recommendations than with them. The difference is significant in four of the six games, and approaches significance in the other two (G(10,1) and G(13,1), with p -values of 0.15 in both). Comparison across games suggests that when recommendations are received (C treatment and V with recommendations sought), correlation becomes weaker as payoff asymmetry increases, but stronger as the unfavoured payoff (and thus the cost of flouting recommendations) increases.

Next, Figure 6 compares actual levels of coordination with the levels that would be predicted based on the observed raw choice frequencies in each game and treatment, combined with the assumption that subjects' choices (and in particular, paired row and column players') are independent and identically distributed. For example, in G(10,7) under the C treatment, subjects choose action A 53.7 percent of the time. If these choices had been i.i.d., coordination should have occurred with probability $2(0.537)(1 - 0.537) = 0.497$. However the actual frequency of coordination is 0.843, substantially higher. Also shown for comparison – as horizontal segments on the right side of the figure – are the actual mixed–NE levels of coordination (which of course assume not only i.i.d. play but also equilibrium choice frequencies).

As the figure shows, when recommendations are not available (the N treatment), predicted and observed coordi-

Figure 6: Observed and predicted (from observed choice frequencies) coordination – treatment aggregates



nation are roughly similar, illustrated by the corresponding plotted points being close to the 45-degree line. The same is true when recommendations are available but not sought (the Vnr treatment), predicted and observed coordination are again roughly similar, with the exception of G(13,7), where many subjects manage to coordinate on the external choice of the decider role. It does happen in some of these cases (e.g., G(10,1) in the N treatment) that observed coordination is substantially higher than the Nash equilibrium level. However, this is due to choice frequencies being closer to 50–50 play than predicted, not to successful coordination between row and column players.

When recommendations are given (C and Vr treatments), the relationship between predicted and observed coordination depends on the game. In G(10,1) and G(13,1), the plotted points remain close to the 45-degree line. As with the no-recommendation cases, it sometimes happens that they lie well above the mixed-equilibrium level (e.g., G(13,1) in the C treatment), due to choice frequencies deviating away from equilibrium levels towards 50–50. Again, this does not imply successful coordination between players.

By contrast, both of the plotted points corresponding to G(10,7) are well above the 45-degree line, with those for G(13,7), G(10,4) and G(7,1) also substantially above the 45-degree line for at least one of the two points (C, Vr or both). Indeed, in some cases (e.g., C treatment for G(10,7) and G(13,7)), coordination levels approach the perfect coordination implied by correlated equilibrium (represented in the figure by the upper segment of the square). Thus the extent to which recommendations facilitate coordination varies across games: highest when there is little payoff asymmetry and a high cost of flouting recommendations, and less as payoff asymmetry increases, the cost of flouting recommendations decreases, or both.

Rather than providing separate analyses of joint payoffs and payoff inequity, we make use of the observation that within each game, there is a perfect correspondence between these and coordination frequencies. This follows directly from the fact that out of the four pure outcomes, two (both choose A and both choose B) involve miscoordination and imply joint payoffs and payoff inequity of zero, and two (exactly one player chooses A) have successful coordination, joint payoffs of $M + n$, and payoff inequity of $M - n$. Thus average joint payoffs (holding the game constant) are just $M + n$ multiplied by the coordination frequency, and average payoff inequity is $M - n$ times the coordination frequency, meaning that the results in Figure 6 are exactly the same if we substitute either payoff levels or payoff inequity for coordination.

D.2 Estimating the fraction of L1 versus L2 types

As noted in the first part of the proof of Proposition 1 (see above), L1 players in the N treatment will always choose A, while L2 players will choose B if and only if $\alpha < n/(M - n)$. So, the probability a given L2 player chooses B is $F(n/(M - n))$, and the probability a given player of unknown type chooses B is $(1 - \theta)F(n/(M - n))$ (from Assumption 1). If we denote the probability of B choices in game $G(M, n)$ to be $q_B(M, n)$, we thus have $(1 - \theta)F(n/(M - n)) = q_B(M, n)$, or

$$1 - \theta = \frac{q_B(M, n)}{F(n/(M - n))}.$$

From Assumption 2, $F(n/(M - n)) > n/M$, so

$$1 - \theta < \frac{q_B(M, n)}{n/M} = \frac{M}{n}q_B(M, n). \quad (4)$$

If we treat the observed frequencies of B choices in the N treatment of our six games as our $q_B(M, n)$, multiplying by M/n gives us an upper bound on $1 - \theta$, the fraction of L2 players in the population, by (4). (We assume that we were successful in randomly assigning subjects to treatments, so the population fraction should be the same for all of our games.) From the second column of Table 3, we have

$$\begin{aligned} 1 - \theta &< \text{Min} \left\{ \frac{7}{1}(1 - 0.764), \frac{10}{1}(1 - 0.803), \frac{13}{1}(1 - 0.817), \frac{10}{4}(1 - 0.685), \frac{10}{7}(1 - 0.590), \frac{13}{7}(1 - 0.616) \right\} \\ &= 0.585. \end{aligned}$$

Thus the fraction of L2 players in the population is *at most* just under 60 percent.

The second part of the proof of Proposition 1 does not provide further information about the distribution of L1 and L2 players, as the condition on α for following B recommendations is the same for both types. The third part of this proof can only help if an assumption is made about the decisions of those L1 types who are indifferent between seeking and not seeking recommendations. If we assume that these players do not seek recommendations, then we have

$$\left(1 - \frac{n}{M}\right)\theta = 1 - q_{seek}(M, n), \quad (5)$$

where the right-hand side of (5) is the observed frequency of seeking recommendations for a given game. Then a process similar to that above yields

$$\begin{aligned} \theta &> \text{Max} \left\{ \frac{7}{6}(1 - 0.810), \frac{10}{9}(1 - 0.764), \frac{13}{12}(1 - 0.782), \frac{10}{6}(1 - 0.750), \frac{10}{3}(1 - 0.838), \frac{13}{6}(1 - 0.727) \right\} \\ &= 0.592. \end{aligned}$$

This implies a fraction of L1 players of *at least* just under 60 percent (and hence L2 players would make up at most just over 40 percent of the population). This is clearly stronger than the above bound, but relies on the assumption we made about indifferent L1 players.

D.3 An aside about signalling

Our standard–theory prediction for behaviour in the V treatment, as discussed in Section 2, is that behaviour would be similar to that in the C treatment if recommendations are given, and like the N treatment if they are not. In this section we discuss an alternative prediction arising from the extra piece of information subjects receive in the V treatment – the exogenous assignment of decider status. Even without obtaining recommendations, this bit of information could be used to break the symmetry of the game.

One way this could happen, even under standard preferences, is through the following “forward induction” behaviour: the player with the choice of seeking recommendations (the “decider”) will choose *not* to seek them, forgoing the correlated equilibrium expected payoff with the intent of playing action A. This is equilibrium behaviour given standard preferences (the other player’s best response is B), and yields a higher payoff to the decider, but a lower payoff to the other player, compared to the CE expected payoff. A more complex signalling setting requires heterogeneous inequity aversion, but replaces our level- k assumption with standard rational beliefs. Players with high levels of inequity aversion could decline recommendations as a (truthful) signal that they would always prefer action A to action B, while players with low inequity aversion would seek recommendations (and follow them).

One shortcoming of both of these theoretical approaches is the usual multiple–equilibrium problem: even though it is possible to show that equilibria exist in which players signal by not seeking recommendations, there are typically other equilibria where signalling does not happen.²⁰ Here, we acknowledge that shortcoming but focus on whether signalling explains our results. First, recall from Section 4.1 that deciders in the V treatment seek recommendations 77.9 percent of the time. So, while choosing not to seek them – a necessary condition for a signalling equilibrium – is clearly not the predominant behaviour in the V treatment, it occurs often enough to be worth a closer look.


Figure 7 shows behaviour in the V treatment where the decider chose not to seek recommendations (i.e., the Vnr treatment), disaggregated by decider and non–decider. In five of the six games, deciders are more likely to choose A than non–deciders, consistent with our signalling hypothesis. The difference is at least weakly significant in three of the games (Wilcoxon signed–ranks test, $p \approx 0.062$ in G(7,1), $p \approx 0.016$ in G(10,4) and G(10,7), while it is insignificant in two ($p \approx 0.16$ in G(13,7) and $p > 0.20$ in G(10,1)), and the effect is in the opposite direction, though insignificant, in one game (G(13,1), $p > 0.20$). As a set of six individual within–game results, therefore, these six results can be considered as suggestive – though not conclusive – evidence that deciders sometimes attempt to signal and that these signals sometimes have the intended effect on the receiver.²¹

However, the relationships *across* games casts doubt on whether this is actually signalling based on inequity

²⁰E.g., consider the following strategy–belief profile. The decider always seeks recommendations, and everyone follows recommendations except for sufficiently inequity–averse people (those whose inequity–aversion parameter α is more than $n/(M - n)$) who choose A irrespective of the recommendation. Off the equilibrium path, any decider not seeking recommendations is believed to be the least inequity–averse type, all non–deciders choose A, and deciders choose B unless they are sufficiently inequity averse (with the same cut–off level as earlier), in which case they choose A. This profile forms a sequential equilibrium as long as the fraction of sufficiently–inequity–averse types in the population is not too high.

²¹Taking the six p –values as a set and applying a correction for multiple comparisons (e.g., Benjamini and Hochberg, 1995), the most significant p –value would be weakly significant at conventional levels ($p \approx 0.094$), while the other five would be insignificant. Alternatively, the probability of at least five of the six games having more A choices for deciders than non–deciders (under a null of either being equally likely to have more) is 0.109.

Figure 7: Frequency of A choices by game in V treatment when recommendations are not sought (all rounds)



aversion. Such signalling ought to be more common as the level of payoff asymmetry increases or as the cost of flouting recommendations decreases. In fact, we see just the opposite. Moving from $G(7,1)$ to $G(10,1)$ to $G(13,1)$ in Figure 7 leads to *fewer* A choices by deciders and *more* by non-deciders, even though payoff asymmetry is increasing. Moving from $G(7,1)$ to $G(10,4)$ to $G(13,7)$ leads to negligible change in A choices by deciders and *fewer* A choices by non-deciders, even though the cost of flouting recommendations is decreasing. So, to the extent that the differences in this figure are due to signalling, there is little reason to believe the signals are coming from the most inequity-averse deciders, though we cannot rule out idiosyncratic occurrences of the “forward induction” strategy described initially in this section. We would encourage future experimental research into the factors affecting signalling in these environments.

Appendix E: Sample screenshots (translated)

Action choice, N treatment:

Period 1 of 1 Remaining time [sec] 54

The earnings in the game are described in the table below:

	Other picks A	Other picks B
You pick A	You win 0 Other Person wins 0	You win 10 Other Person wins 4
You pick B	You win 4 Other Person wins 10	You win 0 Other Person wins 0

Which option do you choose: A
 B

Ready

Action choice, C treatment:

Period 1 of 1 Remaining time [sec] 49

The earnings in the game are described in the table below:

	Other picks A	Other picks B
You pick A	You win 0 Other Person wins 0	You win 10 Other Person wins 4
You pick B	You win 4 Other Person wins 10	You win 0 Other Person wins 0

Based on a random draw, the computer suggests you to play B. At the same time, the computer suggested the other player to play A

Which option do you choose: A
 B

Ready

Decider's choice to seek recommendation, V treatment:

Period: 1 of 1 Remaining time [sec]: 50

The earnings in the game are described in the table below:

	Other picks A	Other picks B
You pick A	You win 0 Other Person wins 0	You win 10 Other Person wins 4
You pick B	You win 4 Other Person wins 10	You win 0 Other Person wins 0

Based on the random draw you will decide whether to let the computer to make a suggestion or not.
Remember that, if the computer suggests you to play A then it will suggest the other to play B. If it suggests you to play B then it will suggest the other player to play A.
Both of these suggestions are equally likely to happen.

Do you let the computer to make a suggestion? No Yes

Ready

Non-decider's action choice, V treatment (decider sought recommendations):

Period: 1 of 1 Remaining time [sec]: 59

The earnings in the game are described in the table below:

	Other picks A	Other picks B
You pick A	You win 0 Other Person wins 0	You win 10 Other Person wins 4
You pick B	You win 4 Other Person wins 10	You win 0 Other Person wins 0

The other player decided to allow the computer to make a suggestion.
Based on a random draw, the computer suggests you to play A. At the same time, the computer suggested the other player to play B.

Which option do you choose: A B

Ready

Non-decider's action choice, V treatment (decider did not seek recommendations):

Period: 1 of 1 Remaining time [sec]: 57

The earnings in the game are described in the table below:

	Other picks A	Other picks B
You pick A	You win 0 Other Person wins 0	You win 10 Other Person wins 4
You pick B	You win 4 Other Person wins 10	You win 0 Other Person wins 0

The other player decided not to allow the computer to make a suggestion.

Which option do you choose: A
 B

Ready

Decider's action choice, V treatment (decider did not seek recommendations):

Period: 1 of 1 Remaining time [sec]: 41

The earnings in the game are described in the table below:

	Other picks A	Other picks B
You pick A	You win 0 Other Person wins 0	You win 10 Other Person wins 4
You pick B	You win 4 Other Person wins 10	You win 0 Other Person wins 0

You decided not to allow the computer to make a suggestion.

Which option do you choose: A
 B

Ready